

LARGE SCALE NEURONAL RECORDING

by

Tim Blanche

BSc. (Hons) University of Technology, Sydney, 1994

A thesis submitted in partial fulfilment of
the requirements for the degree of

Doctor of Philosophy

in

The Faculty of Graduate Studies
Neuroscience

The University of British Columbia
February 2005

© Timothy John Blanche, 2005

abstract

Deciphering the neural basis of brain function will require a significant departure from the reductionistic status quo that has dominated the neurosciences over the past 50 years. In the domain of neurophysiology, this means supplanting single unit recording with electrodes capable of monitoring the activity of hundreds, and ultimately thousands, of neurons simultaneously. Conducting experiments on one or a few neurons at a time and then making elaborate conclusions or models based on these piecewise experiments is not sufficient. Therefore, the objectives of my dissertation were to develop a variety of multisite silicon-based electrode arrays, or *polytrodes*, and establish a set of analytical tools to realise their unique recording capabilities.

Chapter 1 describes the design and testing of high density, 54-site polytrodes, and their use in multiunit studies of cat visual cortex. These polytrodes were able to monitor the activity of more than 100 well-isolated neurons spanning an entire cortical column, a milestone for future experimental studies of cortical circuits. I also describe a continuous data acquisition system designed to cope with the high bandwidth of polytrodes, and techniques for precise electrode positioning. The benign nature of polytrodes was evident both histologically and in prolonged experiments where it was possible to maintain stable recordings from the same neuronal ensemble, even when the polytrode was repeatedly moved.

Polytrodes present significant challenges for conventional spike detection and sorting methods that must be solved before physiological studies are possible. In chapter 2, ideal bandlimited interpolation with sample-and-hold delay correction is shown to accurately reconstruct spike shapes and facilitate spike detection and sorting by reducing waveform variability. Optimal methods of spike detection and sorting were explored in chapter 3 using real and simulated data. A new sorting algorithm that combines unsupervised template generation with multisite template matching was accurate for signal to noise ratios as low as one, and resilient to partial spike overlap. Unlike most existing sorting algorithms, this one is suitable for large contiguous electrode arrays, and is computationally feasible for extended recordings comprising millions of spikes.

Extracellular electrodes do not usually provide accurate information about recorded neuron location, nor any indication of cell type. Chapter 4 describes an

algorithm that capitalises on the fixed, closely-spaced site geometry of polytrodes to localise neurons in 3D cortical space. The algorithm was based on a mixed monopole-dipole field model of extracellular spike potentials and was able to generalise to arbitrary neuron orientation, tissue anisotropies, and cell morphology. Estimated neuron locations emerged as non-overlapping spherical clusters within 150 μ m of the polytrode. Cluster locations moved concordantly with polytrode movements, making the algorithm a useful method for spike sorting unperturbed by electrode drift. Field potential spreads were consistent with the spike shapes and firing patterns of pyramidal cells and interneurons. These results suggest it is eminently possible to identify both the cortical location and type of neurons recorded extracellularly with high density polytrodes.

Chapter 5, the concluding chapter, considers a number of outstanding questions in visual neurophysiology that polytrodes are ideally suited to explore – questions such as the organisation of micro-scale cortical maps, specific cell types responsible for intracortical mechanisms of receptive field tuning, and the identification of precise temporal codes across neural populations. Utilising the parallel recording capabilities of polytrodes, it was possible to characterise the response properties of a large number of neurons to a wide range of visual stimuli, instead of just a few neurons to a narrow selection of stimuli. Since the data were derived from the same neural population, the hope is that a fuller, more unified understanding will be gained of primary visual cortex function, beyond that possible by combining findings from independent experiments.

table of contents

abstract.....	ii
table of contents.....	iv
list of tables	viii
list of figures	ix
list of abbreviations.....	xi
acknowledgements	xiii
 chapter 1 Polytrode development	1
1.1 Summary.....	1
1.2 Introduction.....	2
1.2.1 Silicon based polytrodes	4
1.3 Methods	7
1.3.1 Novel high density polytrodes	7
1.3.2 Polytrode site impedance tester.....	10
1.3.3 Surgery and recording procedures	10
1.3.4 Instrumentation and acquisition software	13
1.3.5 Online CSD analysis.....	15
1.3.6 Acute histological procedures.....	17
1.4 Results	18
1.4.1 General recording properties	18
1.4.2 Improved single unit isolation.....	22
1.4.3 Polytrode reusability	24
1.4.4 Track reconstruction and tissue damage.....	24
1.5 Discussion.....	30
1.6 Next steps... ..	34
 chapter 2 Signal processing.....	35
2.1 Summary.....	35
2.2 Introduction.....	36

2.3 Methods	38
2.3.1 Bandlimited interpolation	38
2.3.2 Sample and hold delay correction.....	39
2.3.3 Noise and artefact suppression.....	40
2.4 Results	42
2.4.1 Waveform reconstruction	42
2.4.2 Reliability of event detection.....	45
2.4.3 Cluster variability	46
2.4.4 Artefact rejection and denoising.....	48
2.5 Discussion.....	50
2.5.1 Avoiding aliasing.....	53
2.5.2 Noise reduction	56
2.6 Conclusions	57
chapter 3 Spike detection and sorting.....	58
3.1 Summary.....	58
3.2 Introduction.....	59
3.3 Methods	62
3.3.1 Spike detection	64
3.3.2 Automated thresholding	66
3.3.3 Partitioning of spike events.....	67
3.3.4 Unsupervised template generation.....	68
3.3.5 Spike classification: multisite template matching.....	71
3.3.6 Evaluating spike detection and sorting reliability	72
3.4 Results	74
3.4.1 Spike detection	74
3.4.2 Spike sorting.....	80
3.5 Discussion.....	84
3.5.1 Detection and sorting of low SNR spikes.....	84
3.5.2 Sorting of overlapping spikes	87
3.5.3 Refining template generation.....	89

3.5.4 Improving sorting efficiency	92
chapter 4 Neuron localisation & classification.....	93
4.1 <i>Summary</i>	93
4.2 Introduction.....	94
4.2.1 Spike-related field potentials	94
4.2.2 Extracellular classification of cell type.....	98
4.2.3 Foundational work on neuron localisation.....	100
4.3 Methods	102
4.3.1 Model overview	102
4.3.2 A mixed monopole, dipole field potential model.....	104
4.3.3 Coordinate transforms	106
4.3.4 Extracellular tissue model	107
4.3.5 The neuron localisation algorithm	107
4.4 Results	109
4.4.1 Neuron localisation	109
4.4.2 Model precision and stability.....	113
4.4.3 Spatial clustering for spike sorting.....	115
4.4.4 Neuron classification.....	116
4.5 Discussion.....	118
4.5.1 Model validation and interpretation.....	119
4.5.2 Model refinements.....	123
chapter 5 Prospective applications	125
5.1 A new experimental paradigm.....	125
5.2 Cellular-scale organization of the visual cortex.....	128
5.3 Timing and temporal coding in the visual cortex.....	132
5.4 <i>in vivo</i> spike timing dependent plasticity	134
5.5 Concluding remarks.....	136
appendix A Impedance measurements.....	137

appendix B Instrument calibrations.....	143
appendix C Validation of neuron localisation.....	150
references.....	155

list of tables

Table 1.1 Polytrode specifications.	8
Table 2.1 Cost-benefit analysis of upsampling.	52
Table 3.1 Spike sorting parameters.....	62
Table 3.2 Comparative performance of different spike detection schemes.....	80
Table 4.1 Neuron localisation model parameters and constraints.....	109
Table 5.1 Visual stimuli and purpose.....	127
Table B.1 Comparison of stimulus displays.....	149

list of figures

Figure 1.1 Experimental techniques across scales of brain organisation.	3
Figure 1.2 <i>in vivo</i> electrophysiology techniques.	4
Figure 1.3 Silicon substrate multielectrode arrays.	6
Figure 1.4 54 site polytrodes.	9
Figure 1.5 Electrophysiology instrumentation.	14
Figure 1.6 Polytrode data acquisition software.	16
Figure 1.7 Recording of neuron ensembles in cat visual cortex.	19
Figure 1.8 Diversity of spike waveforms.	21
Figure 1.9 Polytrodes provide better single-unit isolation than tetrodes.	23
Figure 1.10 Long term stability of recording sites.	25
Figure 1.11 Polytrode track reconstruction.	26
Figure 1.12 CSD analysis.	27
Figure 1.13 Estimation of damage around the polytrode.	28
Figure 1.14 Neuron passage study.	29
Figure 2.1 The case for upsampling.	37
Figure 2.2 Spike waveform reconstruction.	43
Figure 2.3 Sample & hold delay correction.	44
Figure 2.4 Interpolation improves event detection reliability.	46
Figure 2.5 Interpolation reduces cluster variability.	47
Figure 2.6 Robust artefact rejection.	48
Figure 2.7 PCA-based denoising.	50
Figure 2.8 Transition bandwidth, sampling ratio and aliasing.	55
Figure 3.1 Spike sorting flow chart.	63
Figure 3.2 Spike detection algorithms.	65
Figure 3.3 Spatiotemporal partitioning of spike events.	68
Figure 3.4 Real and simulated spike trains.	74

Figure 3.5 Noise estimates for automated thresholding.	76
Figure 3.6 Spatiotemporal spike partitioning.	77
Figure 3.7 Low amplitude spike detection.	79
Figure 3.8 Unsupervised template generation.....	81
Figure 3.9 Sorting of low SNR and overlapping spikes.	83
Figure 3.10 Automated exit criteria and sampling issues.	91
Figure 4.1 Spike field potentials.....	96
Figure 4.2 Neuron location by spike amplitude triangulation.....	101
Figure 4.3 Components of the neuron localisation model.	103
Figure 4.4 Representative model fits.....	111
Figure 4.5 Ensemble neuron localisation.....	113
Figure 4.6 Neuron localisation for spike sorting.	116
Figure 4.7 Classification of cell type.....	117
Figure 5.1 Visual stimuli.	126
Figure 5.2 Cortical micromapping.....	131
Figure A.1 Impedance meter/electroplater circuit diagram.	138
Figure A.2 Impedance meter calibration.	139
Figure A.3 Impedance spectroscopy of polytrode recording sites.	140
Figure A.4 Site impedance changes following polytrode insertion.	141
Figure A.5 The brain is a resistive medium.	142
Figure B.1 Filter properties of polytrode sites and amplifiers.....	144
Figure B.2 PSD of spikes, neural noise and the recording system.....	145
Figure B.3 Sony 200sf monitor calibration.....	147
Figure B.4 Phase locking of cortical neurons to the 100Hz screen refresh.	148
Figure C.1 Histological validation of the neuron localisation algorithm.	151
Figure C.2 <i>in vitro</i> validation of the neuron localisation algorithm.....	153
Figure C.3 <i>in vivo</i> validation of the neuron localisation algorithm.	154

list of abbreviations

2p	2 photon
AC	Alternating current
ACG	Auto-correlogram
ADC	Analog to digital conversion
AHP	After hyperpolarisation
BP	Blood pressure
BPAP	Back-propagating action potential
CCG	Cross-correlogram
CH	Chattering
CSD	Current source density
CSF	Cerebrospinal fluid
DC	Direct current
dil	1,1'-dioctadecyl-3,3,3',3'- tetramethyl- indocarbocyanine perchlorate
DTI	Diffusor tensor imaging
ECG	Electrocardiogram
ECS	Extracellular space
EDTA	Ethylenediaminetetraacetic acid
EEG	Electroencephalogram
EIB	Electrode interface board
ENOB	Effective number of bits
FEF	Frontal eye fields
FIR	Finite impulse response
FS	Fast spiking
GUI	Graphical user interface
I/O	Input/output
IB	Intrinsic bursting
ICA	Independent component analysis
IIR	Infinite impulse response
ISI	Inter-spike interval
LFP	Local field potential
LTD	Long term depression

LTP	Long term potentiation
MRI	Magnetic resonance imaging
MST	Medial superior temporal
MT	Mediotemporal
MTM	Multisite template matching
MUX	Multiplexer
NEO	Non-linear energy operator
PBS	Phosphate buffered saline
PCA	Principal component analysis
PI	Propidium iodide
PMLS	Posterior mediolateral sulcus
PSD	Power spectral density
RC	Resistor-capacitor
RF	Receptive field
RMS	Root mean squared
RS	Regular spiking
SEM	Standard error of the mean
SHD	Sample-and-hold delay
SNR	Signal to noise ratio
SR	Sampling ratio
STDP	Spike timing dependent plasticity
VLSI	Very large scale integration
WT	Wavelet transform

acknowledgements

I wish to thank my supervisor, Nicholas Swindale, for supporting me financially and academically over the course of my PhD. He is an exceptional, multifaceted thinker, and a much respected mentor. I am especially grateful for being given the freedom to pursue my own scientific goals and interests, a privilege rarely afforded to graduate students. Thanks must also go to my unofficial supervisor, Rob Douglas, for many invaluable discussions and for expanding my appreciation of the science of art.

Phil Hetherington provided the foundational work for much of my PhD, not to mention helpful programming advice well after he had left the Swindale lab. Your confidence in me and relentless positive attitude towards problem solving was inspiring, and I thank you. My lifelong friend and mentor Chris Rennie is always there to help me with his lucid scientific explanations, analytical eloquence, and general encouragement. I am indebted to you. In the context of this thesis, he is to be credited for deriving the field potential equations and coordinate system components of the neuron localisation algorithm (Chapter 4).

I thank Martin Spacek for technical support and enduring both the long overnight experiments and my foul company when such experiments were not running smoothly. I have every confidence that you will discover great things in your PhD. Thank you Keith Godfrey, programmer extraordinaire, for your help testing the spike detection and sorting algorithms. A newcomer to the lab, I can only imagine how much further down the road we'd be had you arrived earlier. Thanks to Tara Stewart for helpful advice with perfusion and histological procedures, and Howard Meadows for his sharp ophthalmological surgical tools and even sharper wit. Virginia Booth and Kris Gillespie provided exceptional animal care and surgical training, for which I am very grateful.

Finally, I extend a special thanks to Jamie Hetke and David Anderson from the Department of Electrical Engineering, University of Michigan. Without your collaboration and ready supply of polytrodes this project would not have been possible. I look forward to rewarding your support and patience with many more exciting publications over the coming years.

for Nathalie...
...who loves and cares for me

chapter 1

Polytrode development

1.1 Summary

A variety of 54-channel high density silicon electrode arrays (polytrodes) were developed for recording large numbers of neurons spanning millimetres of brain. In cat visual cortex it was possible to make simultaneous recordings from more than 100 well-isolated neurons. Using standard clustering methods, polytrodes provide a quality of single-unit isolation that surpasses that attainable with tetrodes. Guidelines for successful *in vivo* recording and precise electrode positioning are described. I also describe a high-bandwidth continuous data acquisition system designed specifically for polytrodes, and an automated impedance meter for testing polytrode site integrity. Despite having smaller interconnect pitches than earlier silicon-based electrodes of this type, these polytrodes have negligible channel crosstalk, comparable reliability, low site impedances, and are capable of making high fidelity multiunit recordings with minimal tissue damage. The relatively benign nature of planar electrode arrays is evident both histologically and in experiments where the polytrode was repeatedly advanced and retracted hundreds of microns over periods of many hours. It was possible to maintain stable recordings from active neurons adjacent to the polytrode without change in their absolute positions or neurophysiological and receptive field properties.

1.2 Introduction

Why pursue large scale neuronal recording? So much has been learned about brain function from recording single neuron responses over the past half-century, yet many of the 'big picture' questions in systems neuroscience are not amenable to such reductionism. To understand why, a useful analogy can be drawn between brain function and music. Whereas the complex polyphony of a symphony emerges from the collective activity of the musicians in the orchestra, likewise the collective interaction of neural networks give rise to brain function. In the case of a symphony, it could be argued that the sum (music) could be understood by exhaustively studying the parts (musicians, instruments) in isolation. However, this strategy would never provide a full understanding of the nature of jazz music, as Buzsaki explains:

"[the] independent 'single-cell' approach has yielded significant progress in neuroscience. However, this method would fail when applied to a jazz ensemble where the tune is created by the dynamic interactions among the musicians 'on the fly' and which interactions vary from performance to performance. It also fails when applied to central brain circuits where myriad ensembles are at work at multiple temporal and spatial scales." (Buzsaki 2004)

A similar analogy can be drawn with the operation of a desktop computer. Although the wiring (synaptic connectivity) may be rigid, the state of the memory registers (neuronal activity) and what they represent at any given instant depends on the software (perceptual, motor, cognitive task) the computer happens to be running. Moreover, the connectivity of the brain is anything but rigid (Fregnac et al. 1988; Fu et al. 2002; Trachtenberg et al. 2002; Dan and Poo 2004). Without a wiring schematic and prior knowledge of the *modus operandi* of von Neumann architecture, a chip designer would be hard pressed to reverse engineer a modern CPU given the activity of a single isolated transistor. Similarly, a music historian would not be able to restore the full glory of a Mahler symphony from a remnant score containing the melody for only one instrument. Analogies aside, the contention here is that since neurons do not exist or function in isolation of other

neurons, brain function will never be fully understood by studying its constituent parts in isolation. This is true even for primary sensory cortices, where there is ample of evidence to suggest that the context of a stimulus (Rizzolatti and Camarda 1977; Hammond and MacKay 1981; Vinje and Gallant 2002), recent neural activity (Azouz and Gray 1999), state of arousal (Worgotter et al. 1998), and possibly even attention (Lamme et al. 2000; Roelfsema et al. 2004), all modulate the response characteristics of cortical neurons.

The neurosciences are fortunate to have at their disposal a wide range of experimental techniques with which to study the brain at different scales of structure and function (Figure 1.1). Silicon substrate multisite electrode arrays, or *polytrodes*, are somewhat unique in that they can monitor extracellular action potentials (spikes) and local field potentials (LFPs) not just within, but across multiple spatial and temporal scales. For this reason polytrodes are ideally placed to study the 'symphony' of the brain.

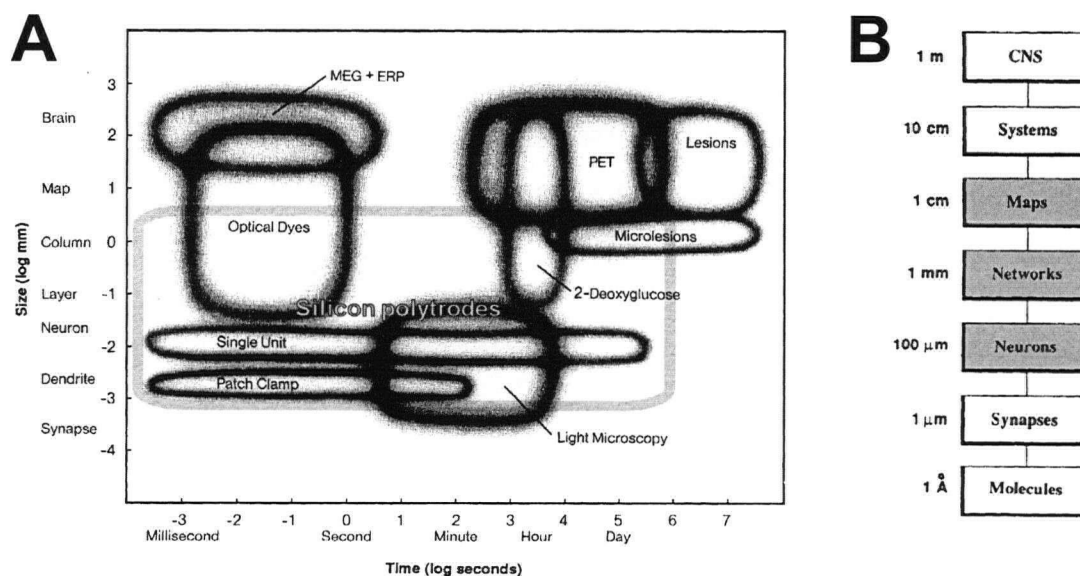


Figure 1.1 Experimental techniques across scales of brain organisation.

(A) Different experimental modalities are useful for investigating different spatial and temporal scales of brain structure and function, encompassing seven orders of magnitude in size, and ten orders of magnitude in time. Present generation polytrodes cover nearly all temporal scales of function and four orders of magnitude spatially, (B) from single neurons, to networks of neurons, to whole cortical maps (*shaded boxes*). Note that some of the more recently developed techniques such as tetrodes, multiphoton imaging and functional magnetic resonance imaging are not portrayed here. Reproduced from Churchland and Sejnowski (1988), with additions.

1.2.1 Silicon based polytrodes

Polytrodes (Drake et al. 1988; Campbell et al. 1991; Wise and Najafi 1991; Kovacs et al. 1994; Kewley et al. 1997; Ensell et al. 2000; Yoon et al. 2000; Norlin et al. 2002; Spence et al. 2003) provide electrophysiological recording capabilities beyond those of conventional electrode technologies (Figure 1.2). For *in vivo* recording in the intact brain, polytrodes possess the same temporal and single-neuron spatial resolution of single unit, intracellular, and patch clamp electrodes. As polytrodes are extracellular electrodes they cannot measure sub-threshold membrane potentials, but by the same token they are not limited to recording from only one or two neurons simultaneously. Moreover, with polytrodes it is possible to maintain stable recordings from the same neurons, not for minutes or hours, but for days, months or even years in awake behaving animals.

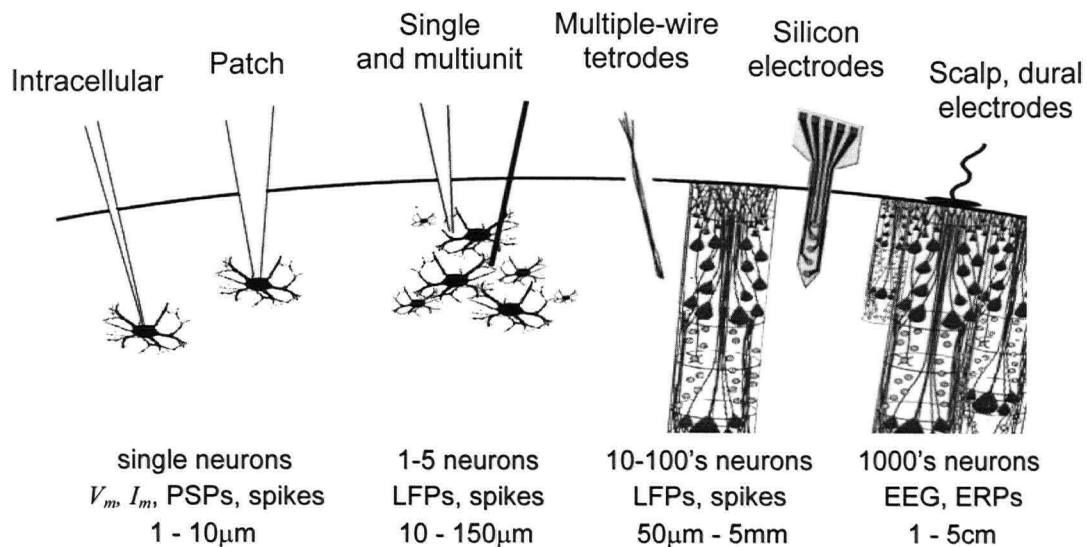


Figure 1.2 *in vivo* electrophysiology techniques.

Currently available electrophysiological methods for *in vivo* recording of brain activity. The number of neurons, suitability for measuring different neural electrical phenomena, and approximate spatial resolution are indicated for each modality.

Polytrodes have electrode site properties suitable for multiunit recording of adjacent active neurons. Although single and multiple-wire multiunit electrodes can record from several neurons over extended periods of time, only polytrodes with high recording site densities can do so from local populations of neurons

across contiguous brain regions such as whole cortical columns. And while epidural or cranial EEG electrodes may be able to monitor larger neuronal populations, they cannot resolve the activity of individual neurons[†].

Stereotrodes (McNaughton et al. 1983) and tetrodes (Wilson and McNaughton 1993) are discrete, 8-12 μ m nichrome multiunit electrodes twisted together to form a single, tightly packed wire bundle electrode. As extracellular spike amplitudes attenuate rapidly over distances of a few tens of microns, these electrodes improve single unit isolation by providing differential spike amplitude 'signatures' that are more or less neuron specific due to the unique distances between individual neurons and the electrode tip (Gray et al. 1995). Polytrodes with closely spaced recording sites afford the improved unit isolation of tetrodes, and like tetrodes can record from up to three times as many neurons as electrode sites (Gray et al. 1995; Maldonado and Gray 1996; Maldonado et al. 1997; Hetherington and Swindale 1999; Harris et al. 2000). In addition, the precise lithographic process by which polytrodes are defined (Najafi et al. 1985) ensures consistent recording properties, and makes possible arbitrary electrode shapes (single or multiple shank) and configurations of recording sites tailored for specific brain structures or applications (Figure 1.3). The number and density of sites that can be etched onto a minute piece of silicon (as narrow as 15 μ m wide and 1-15 μ m thick) far exceeds that of wire electrodes (Najafi et al. 1990), effectively increasing neuronal yield while minimising tissue displacement and potential for damage. Polytrode materials are biocompatible (Niparko et al. 1989) and are suitable for chronic implantation (Hetke et al. 1994; Hoogerwerf and Wise 1994; Rousche and Normann 1998; Mensinger et al. 2000; Vetter et al. 2004), and cortical microstimulation (Anderson et al. 1989; Rousche and Normann 1999; Weiland and Anderson 2000). Polytrodes with on-chip integrated circuitry for buffering, multiplexing, amplification and signal processing (Takahashi and Matsuo 1984; Najafi and Wise 1986; Hoogerwerf and Wise 1994; Bai and Wise 2001; Csicsvari et al. 2003) can minimise noise, channel cross-talk, and movement-related artefacts in chronically implanted devices (Figure 1.3G). Micromachined fluidic channels

[†] it is somewhat meaningless to compare polytrodes with EEG electrodes as the latter are thought to record solely from large scale synchronous synaptic events, not action potentials. Nevertheless, contrasting the temporal and spatial scales of both techniques is relevant to the current discussion, especially as polytrodes can also record high resolution LFPs.

(Figure 1.3E) can also be incorporated into the silicon substrate for cellular-scale chemical and transmitter delivery (Chen et al. 1997; Rathnasingham et al. 2004).

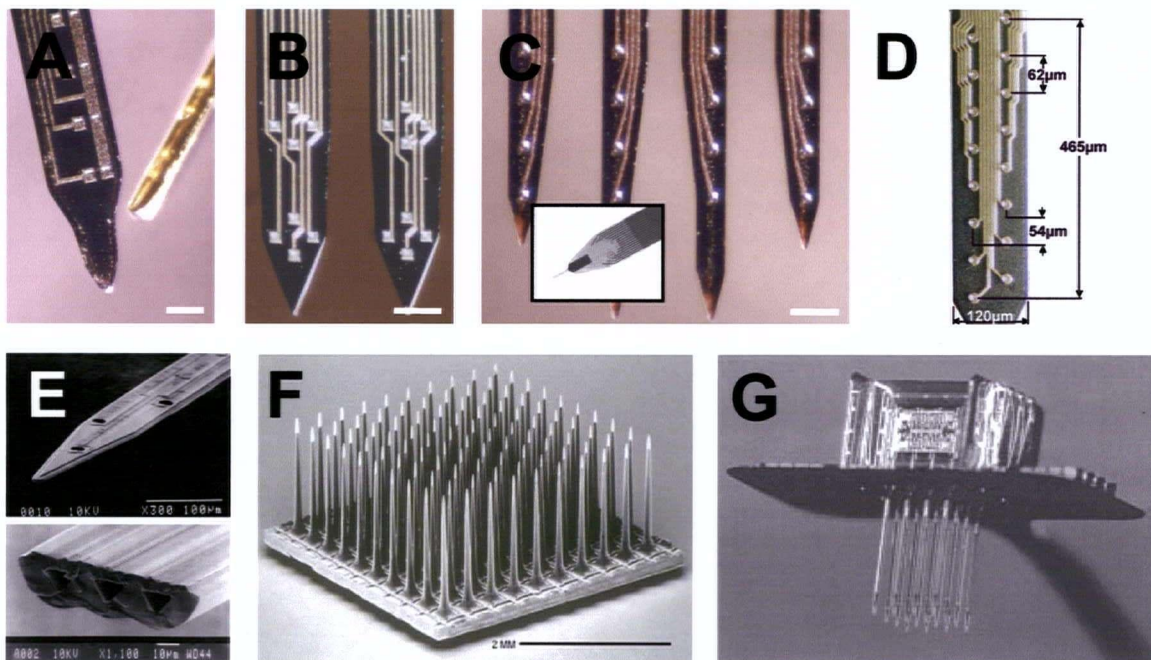


Figure 1.3 Silicon substrate multielectrode arrays.

Silicon lithography enables polytrodes with virtually unlimited numbers of recording sites, built into a variety of 2D planar or 3D configurations. (A) A single shank polytrode with 5 recording sites, shown here next to a 12μm wire tetrode. (B) and (C) 16 site, multi-shank polytrodes. The *inset* shows the four-shank polytrode bonded to a printed circuit board ready for use. Scalebars = 50μm. (D) The predecessor of the 54 site polytrodes described in this thesis (Hetherington et al. 1999). (E) Microfluidic channels for drug and chemical delivery (Chen et al. 1997). (F) A 100 site polytrode with sites spaced 200μm apart in a 10x10 array (Campbell et al. 1991). (G) A 3D multishank polytrode with integrated electronics, assembled from a 4x4 array of 2D polytrodes (Bai and Wise 2001).

This flexibility of design has seen polytrodes used successfully in a diversity of species, brain areas, and applications, from multiunit studies of neocortical plasticity (Fu et al. 2002) to hippocampal recordings in awake-behaving animals (Buzsaki et al. 1992); from mapping of auditory cortex discharges evoked by new generation cochlear-implants (Bierer and Middlebrooks 2002), to the relationship between single units and local field potential (LFP) activity during sleep (Kandel and Buzsaki 1997). Other innovative applications, such as *in vivo* studies of backpropagating action potentials (BPAPs) (Buzsaki and Kandel 1998), deducing intracellular parameters from extracellular waveforms (Henze et al. 2000), and

three dimensional spatial neuron localisation (Hetherington et al. 1999; Blanche et al. 2003), would arguably not be possible with other contemporary electrode technologies.

My interest in polytrodes was motivated by persistent questions in visual neurophysiology (chapter 5) that could only be addressed with an electrode capable of recording simultaneously from large numbers of neurons both within and across whole cortical columns (see also Olshausen and Field 2005). Furthermore, to establish techniques for 'cortical micromapping', that is to determine the precise location of recorded neurons and classification of cell type based entirely on extracellular voltage distributions (chapter 4), required high resolution spike field potential measurements.

In this chapter I describe five novel high-density 54 site polytrodes developed for these studies. I demonstrate their use, handling, and recording characteristics in acute anaesthetised cat visual cortex experiments. The customised hardware needed to record from polytrodes is described, including PC-based continuous data acquisition software, an automated impedance meter for testing recording site viability, and complementary techniques for precise electrode positioning.

1.3 Methods

1.3.1 Novel high density polytrodes

The 54-site polytrodes described in this thesis (Table 1.1) are single-shank planar electrode arrays designed for acute *in vivo* recordings. They are passive devices (*i.e.* no on-chip electronics), fabricated and packaged by the Center for Neural Communication Technology (CNCT) at the University of Michigan. The lithographic process for producing silicon substrate electrodes (Najafi et al. 1985) was pushed close to the laboratory's standard manufacturing limits by halving the usual interconnect conductor width and spacing from 3 μ m to 1.5 μ m. The overall width of the shank was dictated by the number of recording sites and associated conductors, so this refinement was necessary to keep the shank width as narrow as possible. Previous work (Drake et al. 1988) had shown that 15 μ m diameter sites

were ideal for high signal-to-noise ratio (SNR) multiunit recording, so this size was adopted for these polytrodes. The other design consideration was the geometric configuration and spacing of the recording sites. A number of co-linear and staggered site arrangements with different inter-site spacing (Figure 1.4A) were tested to achieve a good trade-off between adequate sampling and isolation of individual neurons (which requires spikes to appear on multiple sites), and traversal of as much brain as possible with a finite number of sites (Table 1.1).

Each configuration has specific advantages. The polytrode with the most closely spaced sites was designed to make high-resolution measurements of spike field potentials, with individual spikes appearing on 12 or more sites. Data derived from this polytrode were needed to 'bootstrap' the field potential model for three dimensional (3D) spatial neuron localisation and classification of cell type (chapter 4). The two-column staggered polytrodes were designed to maintain good single-unit isolation, yet be sufficiently long to record from all layers in a cortical column, whereas the hexagonal three-column designs were a compromise of both these requirements.

Table 1.1 Polytrode specifications.

Polytrode design	Site configuration	Site spacing (μm)	Extent (μm) *	Recording span (μm) †
54 μmap1a	3 column hexagonal	65	1138	1400
54 μmap1b	3 column collinear	50 _v / 46 _h	850	1100
54 μmap1c	3 column hexagonal	75	1313	1600
54 μmap2a	2 column staggered	65	1723	2000
54 μmap2b	2 column staggered	50	1325	1600

Comparison of the five polytrodes. Each has 54 recording sites distributed in a variety of configurations and inter-site spacings. * distance between the most vertically disparate sites. † approximate distance along the long axis of the polytrode over which neurons can be recorded, assuming isolatable units are recordable up to ~150 μm beyond the top and bottom recording sites (Blanche *et al.* 2003).

Each polytrode has additional silicon substrate (4 – 5mm) interposed between the recording sites and the bond pads to allow it to be inserted into the craniotomy cavity without obstruction from the skull. The polytrodes were ultrasonically bonded to a purpose-built, commercially available (Neuralynx, Tucson AZ) printed circuit board (Figure 1.4B). The electrode interface board (EIB) supports the headstage preamplifiers in close proximity to the polytrode, thereby minimising electrical and radio-frequency interference. It also serves as a point of attachment for the micromanipulator, and provides electrical access to recording sites for cortical microstimulation, cortical lesioning, or electrolytic track marking.

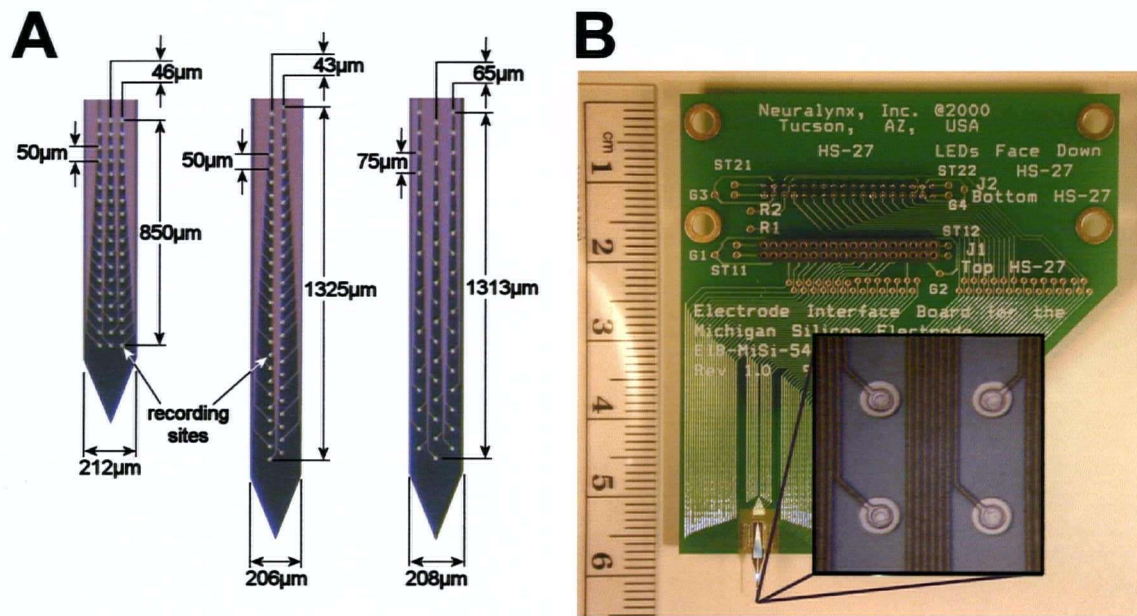


Figure 1.4 54 site polytrodes.

(A) Photomicrographs of three of the five polytrode designs. The polytrode shanks have planar recording sites spaced 43–75µm apart in two or three columns. Polytrodes with denser site spacing (*54µmap1b*) provide highly detailed field recordings, and span ~1mm. The longer two column polytrodes (*54µmap2b*) traverse 1.3–1.7mm, enabling simultaneous recordings of units from entire cortical columns in cat visual cortex. Another design (*54µmap1c*) has a slightly larger intersite spacing in a hexagonal array, and spans ~1.3mm (dimensions for all polytrode designs are given in (Table 1.1)). (B) One of the polytrodes, shown here bonded to the headstage interface board, has closely spaced co-linear sites arranged in 3 columns. The exposed recording sites and 3µm pitch (track plus space) interconnects are visible (*inset*). All shanks are 15µm thick, 199–212µm wide (depending on the design), with 15µm diameter sites made of pure iridium.

1.3.2 Polytrode site impedance tester

Prior to an experiment it was important to identify any faulty recording sites that were shorted together or electrically open. If open 'floating' sites were not grounded they tended to either saturate the amplifiers, producing noise on adjacent functional sites, or else exhibited spurious spike-like signals by capacitive coupling to adjacent conductors. Manually testing every site is laborious and error prone, so an automated multisite impedance tester was made specifically for this purpose. The device utilises software-controlled analog multiplexers to automate switching between recording sites, and is able to test an entire 54-site polytrode in less than a minute. An online graphical display provides a report of impedance magnitude and phase for each site, highlighting any open or shorted sites. The circuit schematic (Figure A.1), calibration and operational details are described in appendix A. Unless otherwise stated, all impedance measurements were made in 0.9% phosphate buffered saline (PBS) against a saturated calomel reference electrode (Accumet cat # 13 620 52).

Stray capacitive coupling of spike signals across independent recording sites compromises the effectiveness of spike sorting methods that rely upon differential amplitude measurements, and is also a recognised source of artefactual synchrony (Zhu et al. 2002a). The possibility that the closer spacing and narrower conductor widths of these polytrodes might make them susceptible to excessive channel cross-talk was tested by injecting a $100\mu\text{V}_{\text{rms}}$, 1kHz sine wave into individual electrode sites via the electrode interface board. The polytrode tip was immersed in PBS with a common reference electrode in the saline. Evidence of coupling in adjacent non-signal sites was looked for in the edge-triggered average of a few hundred cycles of the test signal.

1.3.3 Surgery and recording procedures

Adult cats or rats of either sex were prepared for acute electrophysiological recordings in accordance with guidelines established by the Canadian Council for Animal Care. For the initial surgery cats were anaesthetised with an i.v. bolus of sodium thiopental (2.5% w/v) to effect, with booster injections administered as needed. Either intubation or a tracheotomy was performed, and the cat was

placed in a stereotaxic frame and connected to temperature, BP, ECG, EEG, pO₂, and end-tidal CO₂ monitors. Pressure points and wounds were infiltrated with the local anaesthetics Lidocaine (5%) and Marcaine (bupivacaine hydrochloride, 0.25%), respectively. Dexamethasone (0.3 mg i.m.) was given to prevent brain oedema. Intravenous injections of anaesthetics were discontinued, and surgical anaesthesia was maintained by artificial ventilation with a mixture of 70% N₂O and 0.25 – 1.5% isoflurane in oxygen. Core body temperature was maintained near 38°C with a thermostatically controlled heating pad, and end-tidal SpCO₂ and pO₂ were stabilized at 40 mmHg and 99 – 100%, respectively, by varying the respiration rate. A 5 × 10 mm craniotomy was made over cortical areas 17 and 18. With the aid of a surgical microscope a small area of dura was carefully reflected. At this point paralysis was induced with pancuronium bromide (2 mg/kg) and maintained throughout the experiment by continuous i.v. infusion of pancuronium (0.2 mg/kg/hr) dissolved in lactated Ringers with 5% dextrose, delivered at a rate of 3 mls/kg/hr. Pupils were dilated with topical atropine (5%) and nictitating membranes were retracted with phenylephrine eye-drops (10%). Using reverse ophthalmoscopy, rigid gas-permeable (Metha et al. 2001) contact lenses (Harbour City Contact Lens Services, Nanaimo BC) of appropriate radius of curvature and power were selected to focus both eyes on the stimulus display monitor positioned 50cm in front of the cat. Additional drops of phenylephrine and atropine were applied as needed.

Adult rats were anaesthetised with ketamine/xylazine (50/10 mg/kg i.p.) and placed in a rodent stereotaxic frame. A small craniotomy/durotomy was made over the visual areas of one hemisphere. Vital signs including ECG, pO₂, and core temperature were monitored throughout the procedure, the latter regulated with a small DC heating pad.

Polytrodes are extremely flexible and cannot penetrate the *dura mater* without fracturing, nor usually the *pia mater* without excessive dimpling of the brain. It was thus necessary to reflect the *dura* and make a tiny incision (~300µm long) in the *pia* using ultra-fine micro-dissection scissors (Fine Science Tools, Vancouver BC) or a 32 gauge needle, bent at the end to create a micro-sized hook. Alternatively, angled slit knives intended for ophthalmic microsurgery (ClearCut™ 3.2mm, Alcon Surgical) were also ideal for opening the pia. The polytrode could then be inserted without bending or dimpling of the brain. Tight

physical and electrical coupling of the recording sites to the cortical tissue is essential for robust recordings (Starr et al. 1973). Subjective experience suggests that fluid on the surface of the brain, if allowed to seep down between the silicon substrate and neuropil, acts as a low-impedance shunt to ground, heavily attenuating or even abolishing spike amplitudes. For this reason cerebrospinal fluid (CSF) was routinely wicked away during insertion. While viewing the exposed surface of the brain through a surgical microscope, the polytrode was slowly advanced into the cortex with a micromanipulator (Narishige MW-1, East Meadow NY) until the top sites were $\sim 200\mu\text{m}$ below the surface. The usual practice was to record at a single fixed position per penetration, either traversing a cortical column by inserting vertically in the crown of the lateral gyrus, or down the medial bank of the lateral gyrus for trans-columnar recordings. Only when addressing specific technical questions relating to neuronal damage and stability of unit isolation was the polytrode repeatedly advanced and retracted. After insertion the craniotomy was filled with agar (2.5% in artificial-CSF) to diminish brain movements.

Spike amplitudes are often attenuated or even abolished following movement of the polytrode, presumably due to loss of electrical coupling. For this reason it was desirable to wait half an hour for the polytrode position to stabilise and 'recouple' to the brain tissue (Figure A.4). During this period spike amplitudes were usually restored.

Spikes were evoked with a wide range of visual stimuli (Table 5.1), including drifting bars and sinusoidal gratings, white, pink, and m-sequence noise, flashed stimuli, and natural scene movies. Stimuli were presented on a calibrated (Figure B.3) display monitor (Sony 200sf) with a 100Hz refresh rate and software-linearised gamma correction (mean luminance 55 cd/m^2). All data presented here were recorded from cortical neurons in primary visual areas 17 and 18.

Recordings in cat were typically made for 3 – 8 hours in each penetration. At the completion of the experiment the polytrode tip was carefully withdrawn and immediately cleaned with a jet of de-ionised water from a squirt bottle. Long-term storage was in air. Polytrodes used for acute experiments and cleaned in this way were successfully re-used more than once. The reusability of the polytrodes was assessed by monitoring site impedances and recording performance in successive experiments over several years.

1.3.4 Instrumentation and acquisition software

The electrophysiology system was assembled from commercially available and custom-built hardware (Figure 1.5). Extracellular electrical activity, referenced to a platinum wire loop positioned around the craniotomy, was buffered by two 27 channel unity-gain headstage pre-amplifiers (HS-27s, Neuralynx) mounted on the EIB. Polytrode signals were then passed to a 64 channel amplifier (Multichannel Systems FA-I-64, PS-20 low noise power supply, ALA Scientific Instruments, Westbury NY). All amplifier channels had a factory-fixed gain of 5000. Fifty-four were band-pass-filtered for recording units (500 – 6KHz), the remaining ten for recording LFPs (0.1 – 150Hz). A custom-made patch box (Multichannel Systems) enabled a selection of the polytrode sites to be wired into the LFP channels. The patch box also relayed power to the headstage pre-amplifiers and was used to ground any faulty polytrode sites. A custom-made multiplexer (MUX-80) received parallel outputs from the amplifier for monitoring any of the analog signals on a digital oscilloscope (Kikusui 7310A, Japan) and audio monitor (Grass AM-8, West Warwick, RI). The monitor channel could be selected either manually (push-buttons on the front panel) or in the acquisition software (by entering the site number or simply clicking on the appropriate waveform). Amplified signals were connected to a couple of 32 channel data acquisition cards (DT3010s, Data Translation, Marlboro MA) via two screw terminal blocks (EP307s, Data Translation). These were also used to link the clocks and triggers of the acquisition cards, and access the digital input/output (I/O) lines for controlling the MUX-80, polytrode impedance meter, or acquiring stimulus display related information. Prior to sampling polytrode signals were further amplified $2 - 8 \times$ (corresponding to analog-to-digital converter ranges from $\pm 1\text{mV}$ to $\pm 250\mu\text{V}$ full-scale) and digitized with 12-bit resolution at 25kHz/channel.

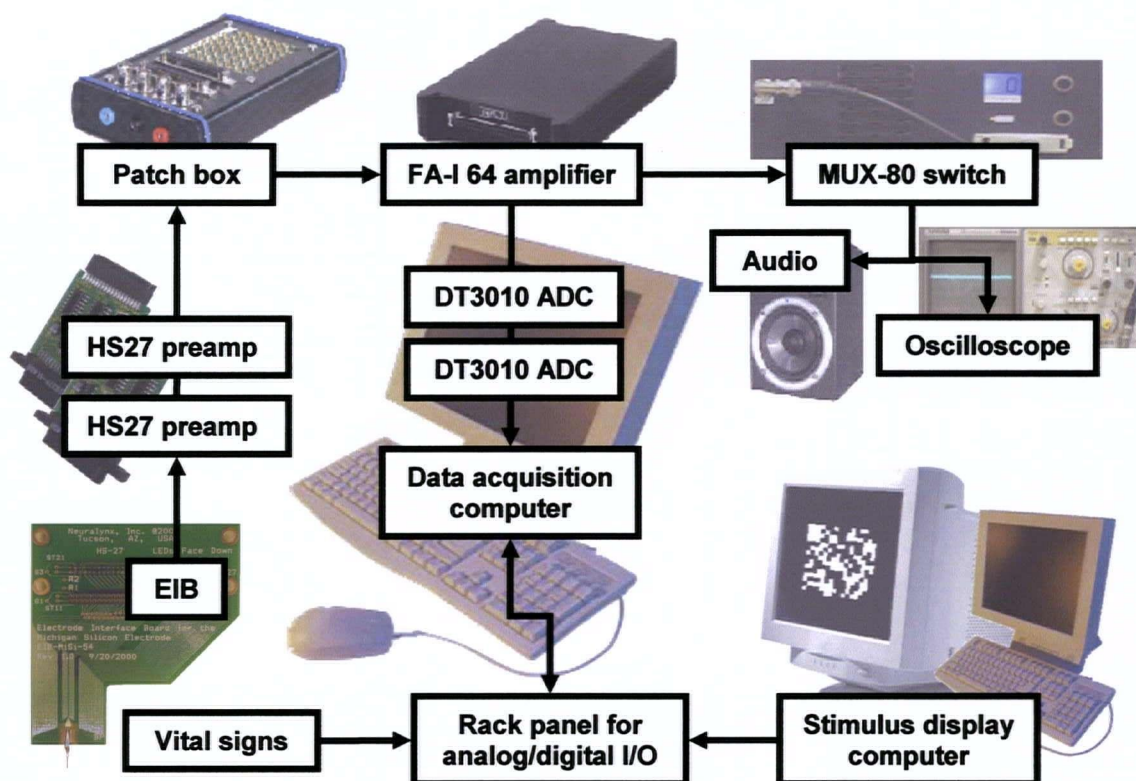


Figure 1.5 Electrophysiology instrumentation.

Schematic showing the major components of the polytrode recording system. The power supply for the headstage pre-amplifiers and amplifier is not represented in this diagram, nor are the screw terminal blocks for connecting input/output signals to the data acquisition cards. See text for details.

During recording, waveform and stimulus display-related data were displayed online and continuously streamed to hard disk using specialised acquisition software (Figure 1.6). Spike waveforms were displayed in one millisecond epochs in the same layout as the recording sites to provide a meaningful display of activity across the polytrode (Figure 1.6A). The software was programmed with the site configurations of the five 54-site polytrodes, in addition to other 16 channel polytrodes made by the CNCT. It can accommodate arbitrary numbers of polytrodes, tetrodes, and single-channel electrodes, treating each as a separate entity with respect to gain, sample rate, and display (Figure 1.6B). This 'electrode-centric' rather than channel-based approach to signal acquisition goes beyond the aesthetics of online display. Knowledge of the precise geometry of recording sites was used throughout subsequent stages of spike

detection and sorting (chapter 3), and physiological analyses (chapters 4 and 5). The polytrode graphical user interface (GUI) highlights where activity is on the polytrode, and allows the user to select a subset of sites for display triggering and/or selective recording to file. LFP and EEG signals were viewable on chart-like scrolling displays (Figure 1.6C). The 1D current source density (CSD) profile (Figure 1.6D) derived from the LFPs provided online feedback as to the depth and alignment of the polytrode in relation to the cortical laminae (see section 1.3.5). In addition, the software was capable of displaying an ongoing record of EEG spectrograms and heart rate to monitor brain state and depth of anaesthesia.

Stimulus related information, digitally encoded in a pre-stimulus data header, was sent by the stimulus display computer and updated on every frame refresh (Figure 1.6E). This information was collected by the acquisition software, removing the need to transcribe details of every visual stimulus. If desired, recordings could be automatically started, paused or stopped in response to the stimulus sequence. The use of an optimised compiler (Delphi, Borland), low-level assembler code where appropriate, and PCI bus-mastered transfer of waveform signals made the system efficient and scalable (64 channels uses ~10% load on an AMD Athlon 1800+ CPU). In practice, the maximum number of recording channels is only limited by the number of amplifiers and acquisition cards installed. Compressed data (lossless) were archived onto 4.9GB DVDs. Each DVD was capable of storing approximately one hour's worth of continuous recording.

1.3.5 Online CSD analysis

Generalised activation of the optic pathway via direct thalamic or photic stimulation evokes a characteristic laminar activation pattern in the primary visual cortex that can be revealed by CSD analysis (Nicholson and Freeman 1975; Mitzdorf and Singer 1978). Although CSD analysis is a legitimate electrophysiological tool in its own right (*e.g.* Mitzdorf 1985; Heynen and Bear 2001), it is exploited here solely to confirm the depth of the polytrode in the cortex.

Evoked responses to brief flashed stimuli (full frame, 10ms duration) were averaged, and the one dimensional CSD was computed from the second spatial derivative of LFPs from vertically aligned (translaminar) sites:

$$\frac{\delta^2 \phi}{\delta z^2} \approx \frac{\phi_{(z+n\Delta z)} - 2\phi_{(z)} + \phi_{(z-n\Delta z)}}{(n\Delta z)^2} \quad (\text{Nicholson and Freeman 1975}) \quad (1.1)$$

where ϕ was the average field potential,

z was the electrode site coordinate perpendicular to the layers,

Δz was the sampling interval (100-150 μm depending on the polytrode), and

$n\Delta z$ was the differentiation grid (typically $n=2$).

The differentiation grid is equivalent to spatial smoothing and reduces high spatial frequency noise. To aid visualization of CSD profiles, colour-mapped time series were generated using cubic spline interpolation (Press et al. 1994) along the depth axis.

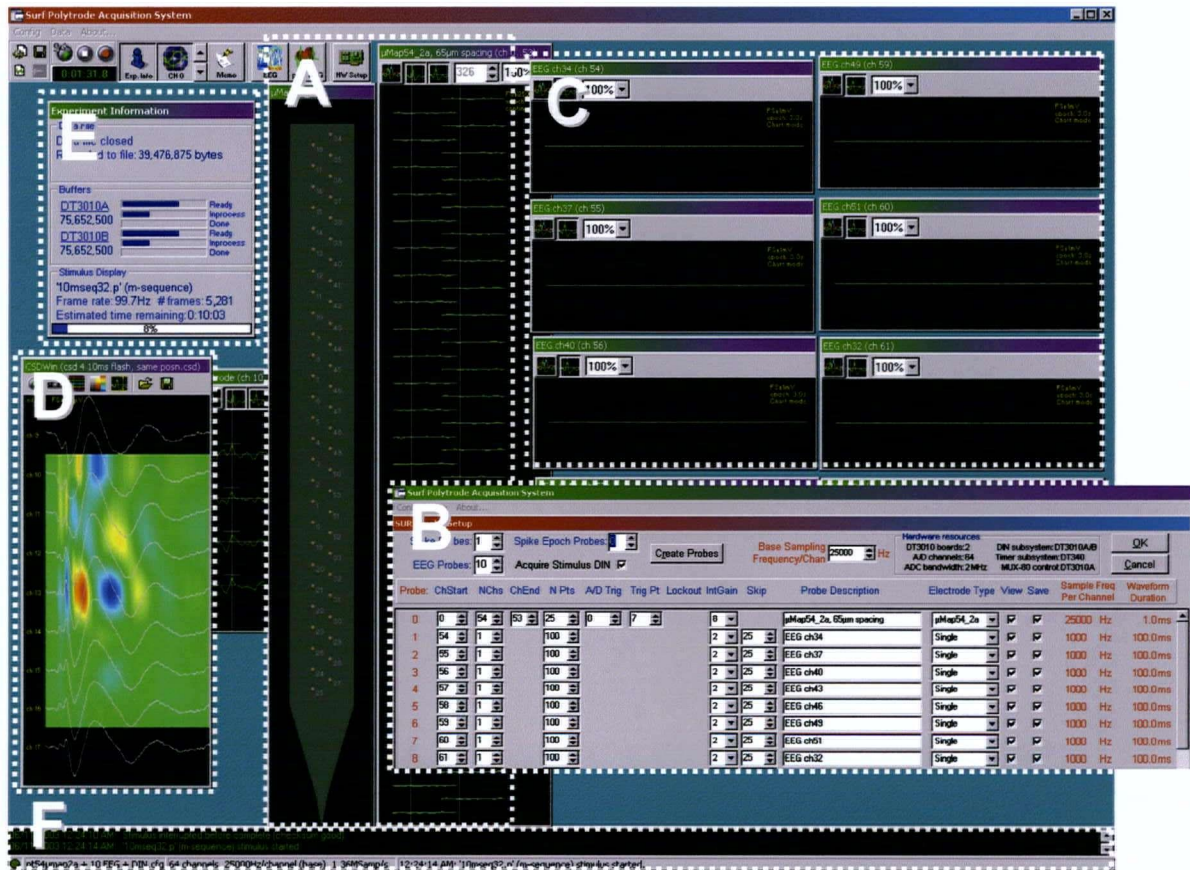


Figure 1.6 Polytrode data acquisition software.

Screenshots of the main program window showing (A) the polytrode GUI and spike waveform windows; (B) the configuration menu; (C) chart windows for displaying LFP and EEG signals; (D) online CSD profile and signal averager; (E) experiment information dialog, indicating the current visual stimulus, stimulus time remaining, and current file statistics; and (F) the experiment log, with provision for entering time-stamped notes.

1.3.6 Acute histological procedures

To determine more precisely the position of the polytrode in the cortex, acute histological procedures were established that avoid tissue shrinkage usually associated with fixatives. These methods were developed and tested in rat for use in the cat experiments, but are equally applicable to track identification in other species. They provide an independent validation of polytrode depth derived from the CSD measurements.

Prior to insertion, the rear of a polytrode shank (side opposite the recording sites) was painted with the fluorescent dye diI (1,1'-dioctadecyl-3,3,3',3'-tetramethylindocarbocyanine perchlorate, ~10% in ethanol, Molecular Probes, Eugene OR) (DiCarlo et al. 1996). As this dye is a lipophilic neuronal tracer, uptake into adjacent neurons and processes also allowed assessment of the level of structural damage caused by the polytrode. The recording properties of the polytrode did not appear to be affected by the dye, but as a cautionary measure the actual recording sites were not coated. The polytrode was then inserted into the cortex so that the top row of electrode sites was ~200µm below the surface of the brain. CSD profiles were evoked by photic stimulation and saved to file for later registration with the histology.

Immediately post-euthanasia the cortical region of interest was blocked *in situ*, removed and bathed in chilled PBS. After carefully removing the *pia mater*, 300–400µm thick coronal sections were cut on a tissue vibratome and counterstained with green fluorescent Nissl stain (Neurotrace 500/525, Molecular Probes). Nissl substance is abundant in the rough endoplasmic reticulum of neuronal cells, and Neurotrace is the fluorescent analog of conventional chromophoric Nissl stains such as cresyl violet. Briefly, the counterstain procedure involved permeabilising the tissue with Triton X-100 (0.1% w/v in PBS) for 10 minutes, followed by two 5 minutes washes with PBS. The brain slices were incubated for 30 minutes in a 50-fold dilution of the supplied stock solution. After repeating the permeabilisation and washing steps, the slices were transferred to glass slides with 90% w/v glycerol + DABCO (an antifade agent) in PBS, and coverslipped. The polytrode track, clearly demarcated by the diI against the Nissl-stained cortex, was then visualised on either a standard wide field or confocal fluorescent microscope (Zeiss LSM-510, Göttingen, Germany).

The degree of structural damage to the neural tissue surrounding the polytrode was investigated further in the same series of rat experiments. Instead of diI, a polytrode coated with propidium iodide (PI, ~10% w/v in dH₂O, Molecular Probes) was inserted into the cortex. PI is a non-toxic polar compound that can be used as an indicator of cell membrane integrity and viability (Vornov et al. 1991). The dye enters damaged or necrotic cells with leaky or otherwise compromised cell membranes, binds to nucleic acids, and becomes brightly red fluorescent. After an hour the polytrode was removed, and as before 400µm thick coronal slices were cut from the fresh, unfixed brain tissue. The brain slices were counterstained with the fluorescent Nissl stain, mounted, and coverslipped as previously described. Neurons surrounding the polytrode track were then reconstructed from serial optical sections imaged with the confocal microscope. An objective quantification of the extent of cellular damage was made by counting the proportion of Nissl-stained cells (predominantly neurons, in green), to Nissl-stained cells co-localized with PI (damaged neurons, in yellow), to PI-positive cells (non-neuronal cells, in red).

1.4 Results

1.4.1 General recording properties

Recording site impedances were $1.17\text{M}\Omega \pm 150\text{k}\Omega$ at 1kHz ($n = 1002$, \pm stdev., from 20 assorted polytrodes), with average phase angles of $-75.8^\circ \pm 3.4^\circ$ (*i.e.* largely capacitive). Such minor variations in site impedance had no measurable effect on the sensitivity of the site, nor the amplitude of recorded spikes. However, sites with significantly higher impedance tended to be slightly noisier. Impedance spectroscopy for a typical recording site is given in appendix A. On average individual polytrodes had 3.9 faulty sites (median 4, range 0–8), which were usually open circuit but occasionally shorted together. Exclusion of this number of faulty sites did not significantly compromise recordings because of the high spatial sampling of these polytrodes. In the multiunit bandpass (500–5kHz), noise was typically $3 - 4\mu\text{V}_{\text{rms}}$ ($20 - 30\mu\text{V}_{\text{pp}}$), depending on the site, measured in saline. Noise in the LFP/EEG bandpass (0.1–150Hz) was also around $3\mu\text{V}_{\text{rms}}$,

predominantly 60Hz line interference. All designs had negligible channel cross-talk. Even sites with adjacent interconnects showed less than 0.5% coupling. Taken together, the consistency of recording site properties indicates that the fabrication process for these polytrodes was as reliable as that for 'standard' 16 site polytrodes (Najafi et al. 1985), despite smaller features and higher site densities.

Representative multiunit recordings are shown in Figure 1.7 and Figure 1.8.

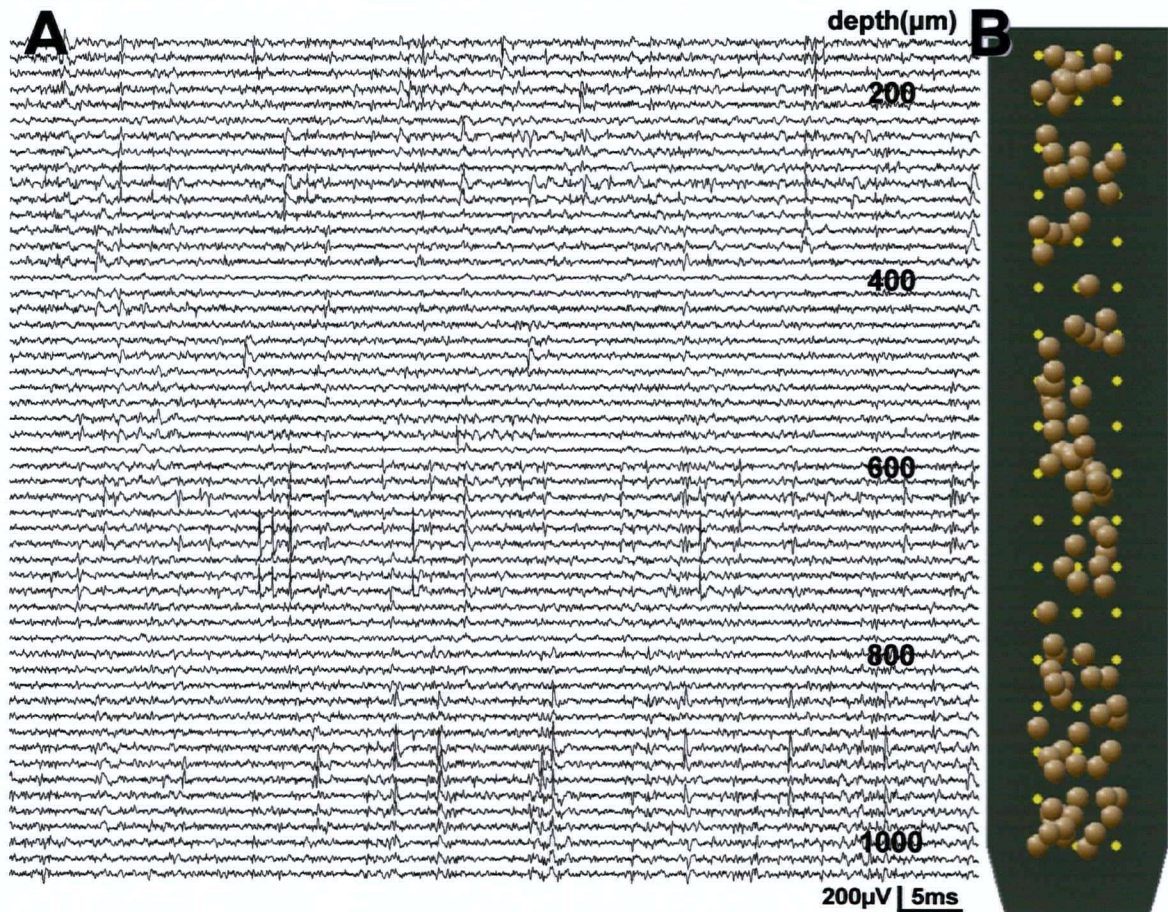


Figure 1.7 Recording of neuron ensembles in cat visual cortex.

(A) 100ms segment of multiunit activity from a 54μm x 1mm polytrode (Table 1.1) inserted perpendicularly to the surface of the cortex, traversing the cortical depth indicated. This particular recording comprised more than 100 spontaneously active and visually responsive neurons distributed across the length of the polytrode. Individual action potentials spanned several sites. (B) Approximate neuron positions from the same recording. The relative positions were estimated from the mean weighted spike amplitudes across channels.

As with any electrophysiological recording, neuronal yield is determined by many factors, including the number of nearby active neurons, depth and type of

anaesthesia, quality of visual stimulation, and cortical location (e.g. granular vs. agranular layers). Identification of well-isolated single units is also dependent on good SNR and the efficacy of spike detection and sorting. Acute recordings in anaesthetised cat visual cortex usually yielded between 20 and 50 simultaneously recorded units at a given location. One of the best recordings to date contained more than 100 clearly distinct neurons (Figure 1.7A), isolated using the multisite template-based clustering procedure described in chapter 3. Active neurons were distributed along the full extent of the polytrode (Figure 1.7B).

In a random sample of 255 neurons from 8 penetrations (Figure 1.8F), spike amplitudes ranged from noise up to 1.2mV_{pp} (mean $144 \pm 118\mu\text{V}$), presumably due to differences in the size, morphology and proximity of the neuron to the electrode sites (Fatt 1957; Bishop et al. 1962a; 1962b; Rosenthal et al. 1966; Humphrey 1976; Drake et al. 1988; Henze et al. 2000). Given an aggregate noise level of $30 - 40\mu\text{V}_{\text{pp}}$ (including thermal and biological noise from distant neuronal activity), this translates into a SNR of up to 30:1. For the majority of recorded neurons with peak-channel spike amplitudes of around $130\mu\text{V}_{\text{pp}}$, a 4:1 ratio was attained, more than adequate for unit identification. These amplitudes are comparable to those recorded extracellularly with conventional multiunit electrodes and tetrodes (Gray et al. 1995; Hetherington and Swindale 1999). Given the similar impedances and surface area of recording sites this result was expected but not guaranteed, since polytrodes comprising planar electrode sites only record neurons in front of the insulating shank (Drake et al. 1988), not the vicinity of the tip (Henze et al. 2000).

With respect to field potential spread, on the highest density three-column polytrode (1b design) up to 16 sites detected the action potentials from (presumed) pyramidal neurons with large 'open' fields (Figure 1.8A, E). Individual spikes on the polytrode with sites spaced $65\mu\text{m}$ apart in a hexagonal layout (1a design) were recorded by up to 9 sites (Figure 1.8B). Fast spikes from smaller 'closed-field' neurons, most likely inhibitory interneurons (Humphrey 1976; Henze et al. 2000; Blanche et al. 2003; Bartho et al. 2004), showed appreciable signal on only one or two sites irrespective of the intersite spacing (Figure 1.8C). In contrast, some neurons infrequently discharged spikes with current dipoles moving 100s of microns (Figure 1.8D), with a velocity ($0.7 \pm 0.15\text{m.s}^{-1}$) and direction consistent with a BPAP travelling up the apical dendrite of a large pyramidal neuron (Johnston et al. 1996; Buzsaki and Kandel 1998).

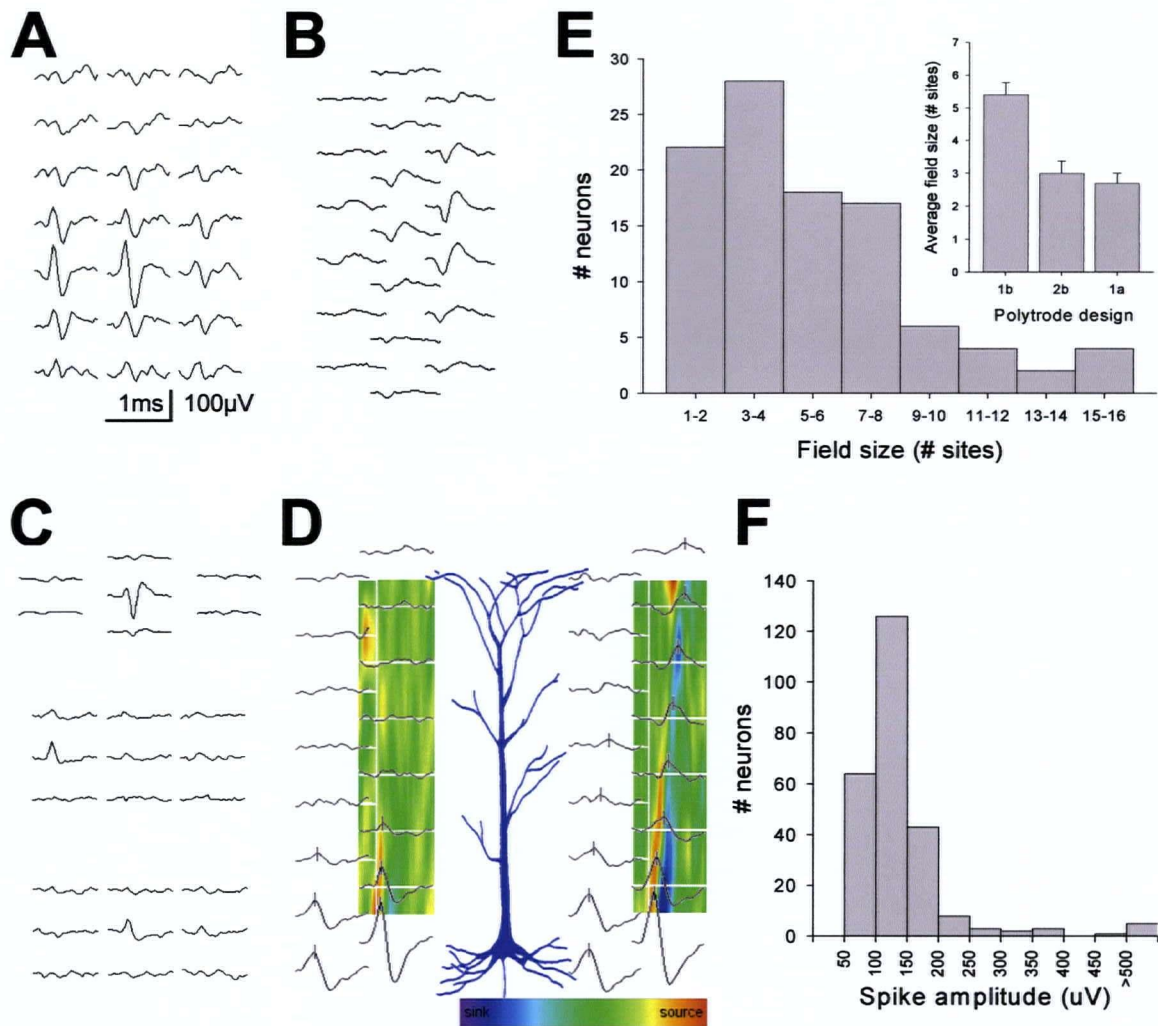


Figure 1.8 Diversity of spike waveforms.

Unaveraged spike waveforms, arranged according to site layouts, recorded with (A) a $54\mu\text{map}1\text{b}$ polytrode, showing a triphasic spike across 12 sites; (B) a biphasic spike on 5 of the more widely separated sites of a $54\mu\text{map}1\text{a}$ polytrode; (C) three small field, fast-spiking neurons; and (D) a $54\mu\text{map}2\text{a}$ polytrode showing the putative location of a layer-5 pyramidal cell relative to the polytrode. Regular somatic action potentials (*left panel*) were typically restricted to sites within a $\sim 150\mu\text{m}$ radius of the maximum amplitude site, and had instantaneous peak times (*dashes*). In contrast, action potentials presumed to be back-propagating (*right panel*) had current dipoles travelling $500\mu\text{m}$ up the apical dendrite (CSD colour maps), with increasing latency to peak amplitude. (E) Histogram showing the approximate electric field spread of the neurons portrayed in Figure 1.7, defined as the number of sites with spike waveforms over $50\mu\text{V}_{\text{pp}}$. The *inset* compares the average number of sites per neuron for polytrodes with different site configurations. Error bars are standard error of the mean (SEM). (F) Distribution of peak-channel spike amplitudes from multiple penetrations.

1.4.2 Improved single unit isolation

To assess the recording performance and capacity of these polytrodes to isolate single units compared with more established techniques, 'virtual tetrodes' (all sets of four adjacent polytrode sites excluding faulty ones) were constructed from the 54 μ map1b recording shown in Figure 1.7A. This polytrode has a site spacing comparable to that of real wire tetrodes (Jog et al. 2002). Including any neuron with a spike amplitude over 60 μ V_{pp} on at least one site of each virtual tetrode, between 7 and 24 neurons (mean = 17.1 \pm 4, n = 101) were counted, roughly double that typically reported for wire tetrodes (Gray et al. 1995; Maldonado and Gray 1996; Maldonado et al. 1997; Hetherington and Swindale 1999; Harris et al. 2000). Virtual tetrodes with this number of neurons were common in other recordings. This raises an important question: in a real tetrode recording, how many neurons identified as single-units are actually multiple neurons? A case study addressing this issue is presented in Figure 1.9. In this example 10 neurons were identified on a virtual tetrode (Figure 1.9A) using multisite template-based clustering (chapter 3). The first two principal components (PCs) derived from these templates were used to both manually and automatically [Klustakwik, (Harris et al. 2000)] cluster the data. The Mahalanobis distance was used to quantify the distance between the cluster centroids of any cluster pair $x_i = (x_1 \dots x_n)$ and $y_j = (y_1 \dots y_n)$ in n -dimensional feature space, from which an average cluster separation index S was derived:

$$d_{ij} = \sqrt{(x_i - y_j)^T C^{-1} (x_i - y_j)} \quad (\text{Mahalanobis 1936}), \quad S = \overline{d_{ij}} \text{ for } i, j = 1 \dots N, i \neq j \quad (1.2)$$

where C was the covariance matrix of the cluster pair, and N was the number of clusters. Both clustering methods could identify 7 of the neurons using the tetrode-derived PCs ($S = 14.5$), but neither method was able to separate the 3 other neurons ($S = 3.7$) due to the close similarity of their waveforms on the tetrode sites (Figure 1.9B). However, when the PCs from the surrounding polytrode sites were included (Figure 1.9C), both clustering methods could now accurately delineate these single units ($S = 12.1$), in addition to improving the overall cluster separation of the 7 other units ($S = 16.3$).

So an actual tetrode positioned at this location would have produced 8 clusters; 7 that represented valid single units, and an 8th supercluster comprised of

three neurons, indistinguishable from a valid single-unit cluster, except perhaps by other criteria [e.g. an autocorrelogram without a 1ms refractory period, although this test is unsuitable for fast spiking (FS) neurons (Nowak et al. 2003)].

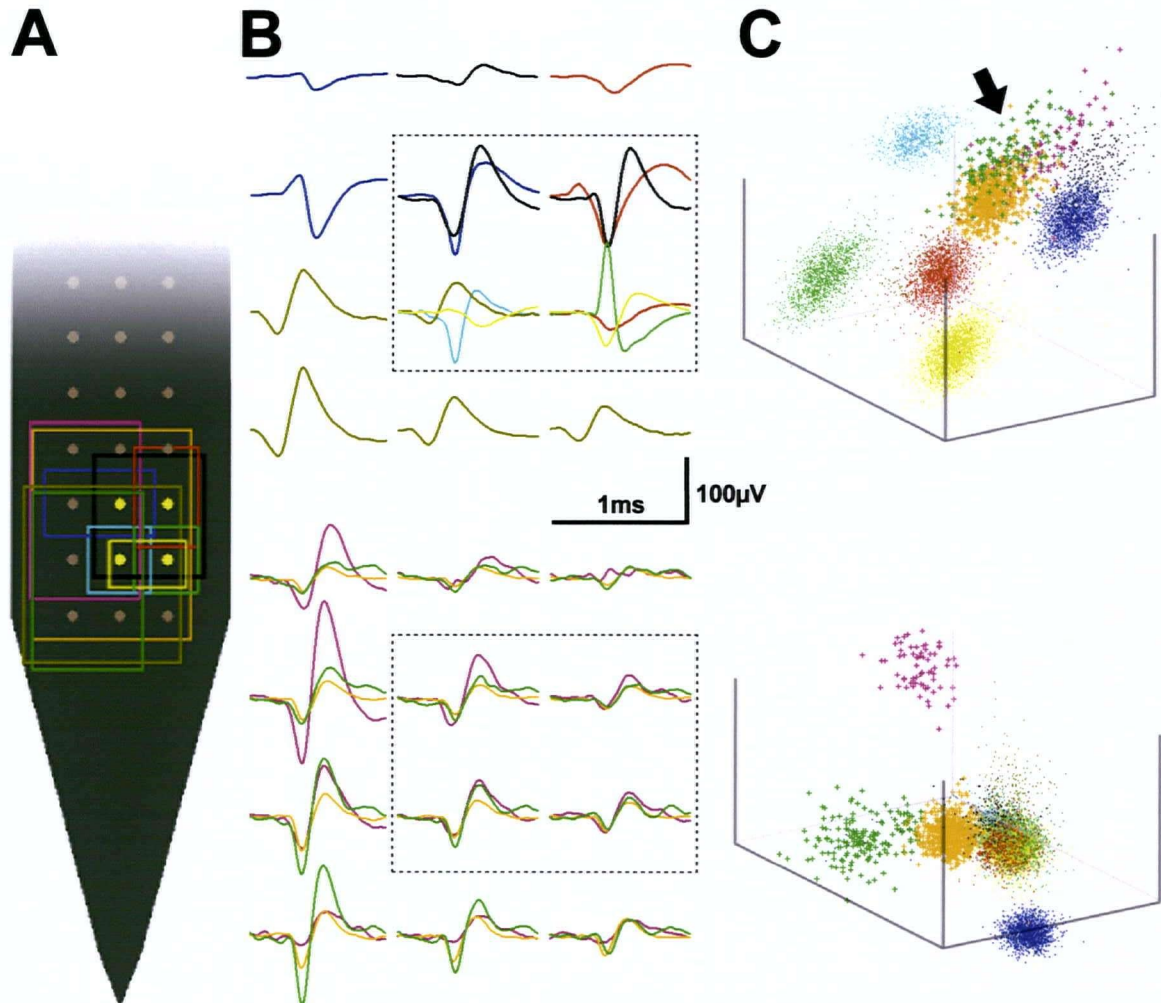


Figure 1.9 Polytrodes provide better single-unit isolation than tetrodes.

(A) Breaking the polytrode into 'virtual tetrodes' revealed that 7 - 24 neurons were recorded on any one virtual tetrode of a 54µm x 1b polytrode. In the example shown here (*highlighted sites*) 10 neurons were identified, their approximate field sizes represented by the coloured boxes. (B) Spike waveforms of these 10 neurons, 7 of which were isolatable using the tetrode sites alone (*upper panel boxed region*), and 3 that could only be isolated by including spike waveforms from the surrounding recording sites (*lower panel*). Note the similarity of the spike shapes and amplitudes on the tetrode sites. (C) A tetrode-derived cluster plot (*upper panel*) shows 5 of the 7 well-isolated neurons. Whereas 2 other neurons (*black and gold dots*) were separable along other dimensions, 3 neurons (*arrow: green, orange and pink '+' symbols*) were not distinguishable on any projection or hyperplane of the tetrode. Only by incorporating spike waveform indices from the surrounding sites was it possible to cluster and isolate these neurons (*lower panel*).

1.4.3 Polytrode reusability

Polytrode site impedances were measured following long-term storage and repeated use in acute 3-5 day experiments run several months apart (Figure 1.10). Although there was a slight increase in average site impedances from $1.1\text{M}\Omega$ to $1.3\text{M}\Omega$, this was without concomitant deterioration of recording performance. The number of faulty sites did not increase, and neuronal yield and spike amplitudes remained qualitatively unchanged. Inert iridium recording sites, when washed immediately following use and stored in air, can thus be repeatedly re-used for hundreds of hours under these conditions. The polytrodes are eventually broken during handling but have not yet been discarded due to degradation of site impedances or recording characteristics. If necessary, site impedances could be restored to pre-use levels by soaking the polytrode tip overnight in 0.25% w/v trypsin-EDTA (Invitrogen), followed by thorough rinsing in distilled water (Figure 1.10C, D). Most accumulated proteinaceous material is removed by this simple cleaning procedure.

1.4.4 Track reconstruction and tissue damage

Without direct staining or electrolytic lesions it is difficult to discern polytrode penetrations in histological sections. This is encouraging from the perspective of tissue displacement, but an obstacle for determining the location of recorded cells. The diI track staining method (Figure 1.11), is a straightforward and effective way of determining polytrode depth and alignment, avoiding the problem of tissue shrinkage associated with histological processing. Since the positions of the recording sites on the polytrode shank are known, it is possible to infer the precise cortical location of every site by simply imaging the outline of the polytrode (Figure 1.11A, E). CSD analysis (Figure 1.12) provides a complementary measure of polytrode depth.

Regarding polytrode-induced tissue damage, there was no evidence of extensive tearing or distortion of neurites or pericytes (Figure 1.11B). The track was approximately $20\mu\text{m}$ thick, and surrounding microvasculature appeared undamaged (Figure 1.11C). Neuronal somata immediately around the track had ostensibly normal morphologies (Figure 1.11D). Staining with PI shows that

damaged neurons and glia were restricted to a region immediately surrounding the penetration (Figure 1.13A). The percentage of damaged neurons decreased exponentially with distance from the polytrode (Figure 1.13C, D). Of the total recordable volume in front of the polytrode (0.0126 mm^3 for the field of view shown in Figure 1.13B), less than 2% of the neurons were damaged or necrotic. It is important to emphasise that this is likely an overestimate of the neuronal damage. Some of the Nissl stained cells were probably glia, and staining for PI does not necessarily indicate cell death, only that an axon or dendrite has been sufficiently compromised to permit uptake of the dye.

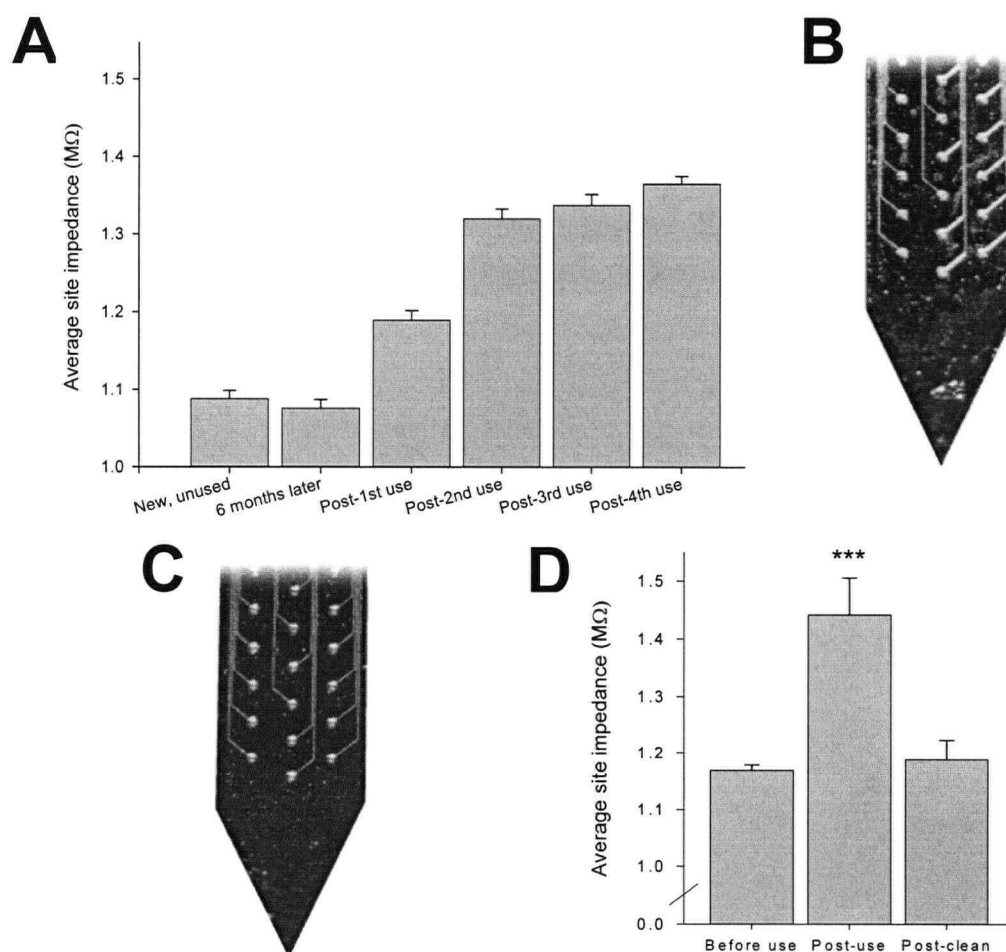


Figure 1.10 Long term stability of recording sites.

(A) Post fabrication site impedances were highly consistent and remained unchanged for six months when stored in air. Over the course of several cat experiments with multiple penetrations in each experiment, site impedances increased slightly, however recording properties remained similar. Another polytrode (B) showing residue accumulated from five recording sessions, and (C) following overnight cleaning in trypsin. (D) Site impedances of this polytrode before and after cleaning *** $p < 0.0005$. Error bars are SEM.

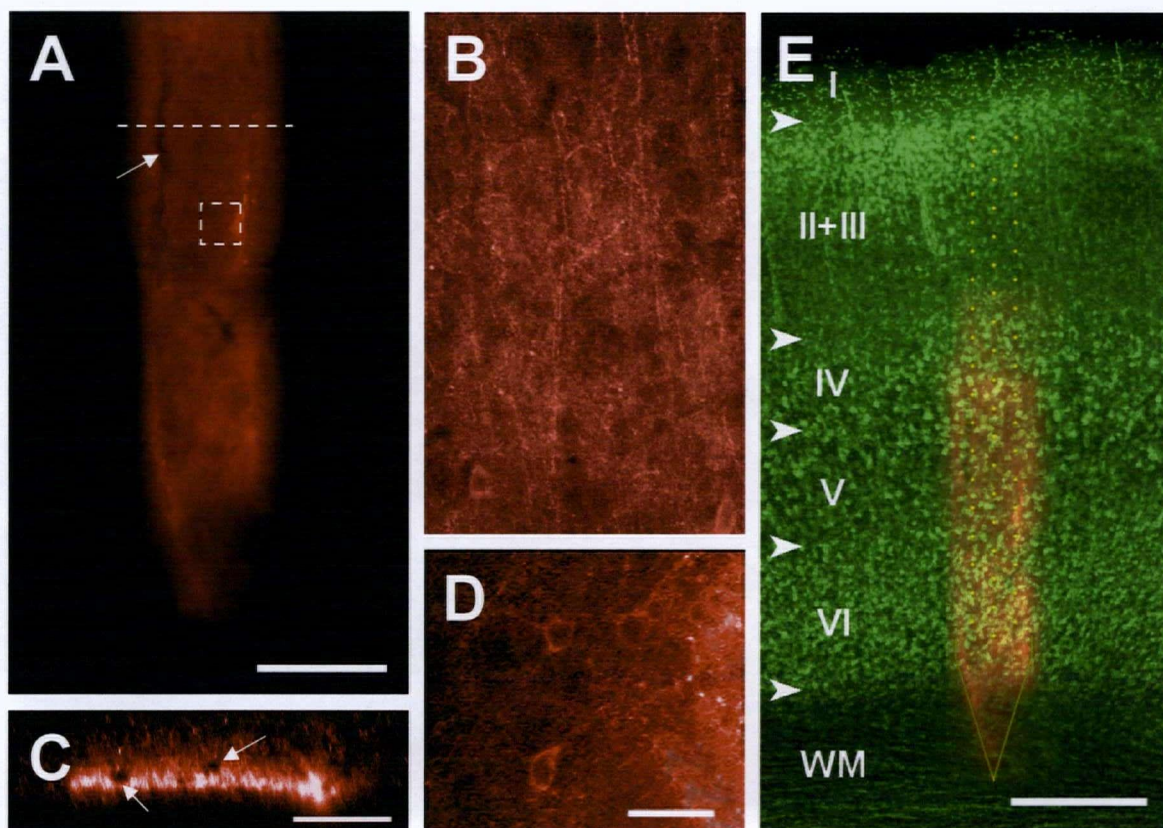


Figure 1.11 Polytrode track reconstruction.

(A) A coronal rat brain section stained with fluorescent diI deposited by the polytrode, showing the faint outline of the tapered tip. A blood vessel (*arrow*) is visible running along the plane of the track. Scalebar = 200µm. (B) Reconstruction of a 50µm stack of confocal images directly in front of the penetration. Stained processes are neurites and pericytes indicative of normal tissue morphology. (C) Cross-section of the polytrode track at the level of the *dashed line* in (A). The thickness of the track is indicated by the region of intense diI fluorescence. Note the proximity of the intact microvessels (*arrows*). Scalebar = 75µm. (D) Enlarged single plane confocal image 15µm anterior to the penetration site (*boxed region in A*), shows three neurons and their processes. Scalebar = 50µm. (E) The same electrode penetration, counterstained with fluorescent Nissl stain to delineate the boundaries of the cortical layers (*arrows*). Since the depth and orientation of the polytrode is clearly defined (*tip outlined in yellow*), the lamina location of the electrode sites can be precisely determined (*yellow circles*). The faded region in the top right was photobleached. Scalebar = 300µm.

There was no indication of any difference in the prevalence, quality, or amplitude of neurons recorded at the top, bottom or central recording sites. For example, in the recording portrayed in Figure 1.7B, there were 32, 36, and 33 neurons distributed on the top, middle, and bottom recording sites, respectively.

On numerous occasions it was possible to monitor and record neuron ensembles over successive advancements of the polytrode over hundreds of microns for many hours. Moreover, the same neurons – as determined by their spike shapes and field potential distribution across the polytrode, relative spatial location, firing pattern, and distinctive receptive field properties – could be recorded upon *retraction* of the polytrode, without noticeable deterioration of these properties (Figure 1.14). Together with the histological results, it is reasonable to conclude that tissue damage caused by the polytrode is restricted and relatively minimal compared to conventional electrodes. This is presumably because polytrodes have an ultra-thin silicon substrate (10–15 μm) that, unlike tungsten in glass or tetrode wire bundles, does not create a bore hole as it penetrates the brain.

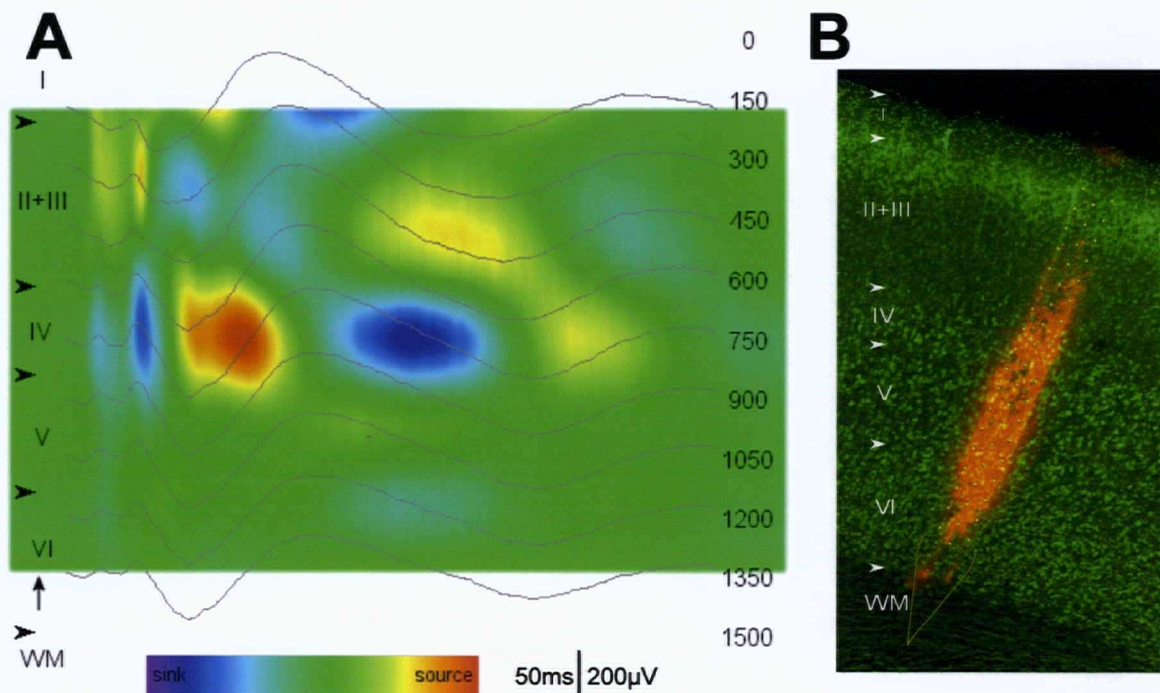


Figure 1.12 CSD analysis.

(A) Photic stimulation evokes characteristic LFP responses in rat visual cortex (*average traces*, $n = 400$), from which the CSD profile was derived (*colour map*). Prominent current sinks and sources in upper layer II/III and IV can be used to identify the depth of the polytrode with respect to the cortical layers (*arrows*). In this example the top and bottom polytrode sites were 150 μm and 1350 μm deep, respectively (*cortical depth in μm is indicated to the right of the figure*). The flash (*vertical arrow*) duration was 10ms, LFPs were sampled (Δz) every 150 μm , and for the differentiation grid $n = 2$. (B) Histological reconstruction of the electrode track confirmed the vertical alignment and depth of the polytrode. The tip was visible in the white matter of the adjacent brain section.

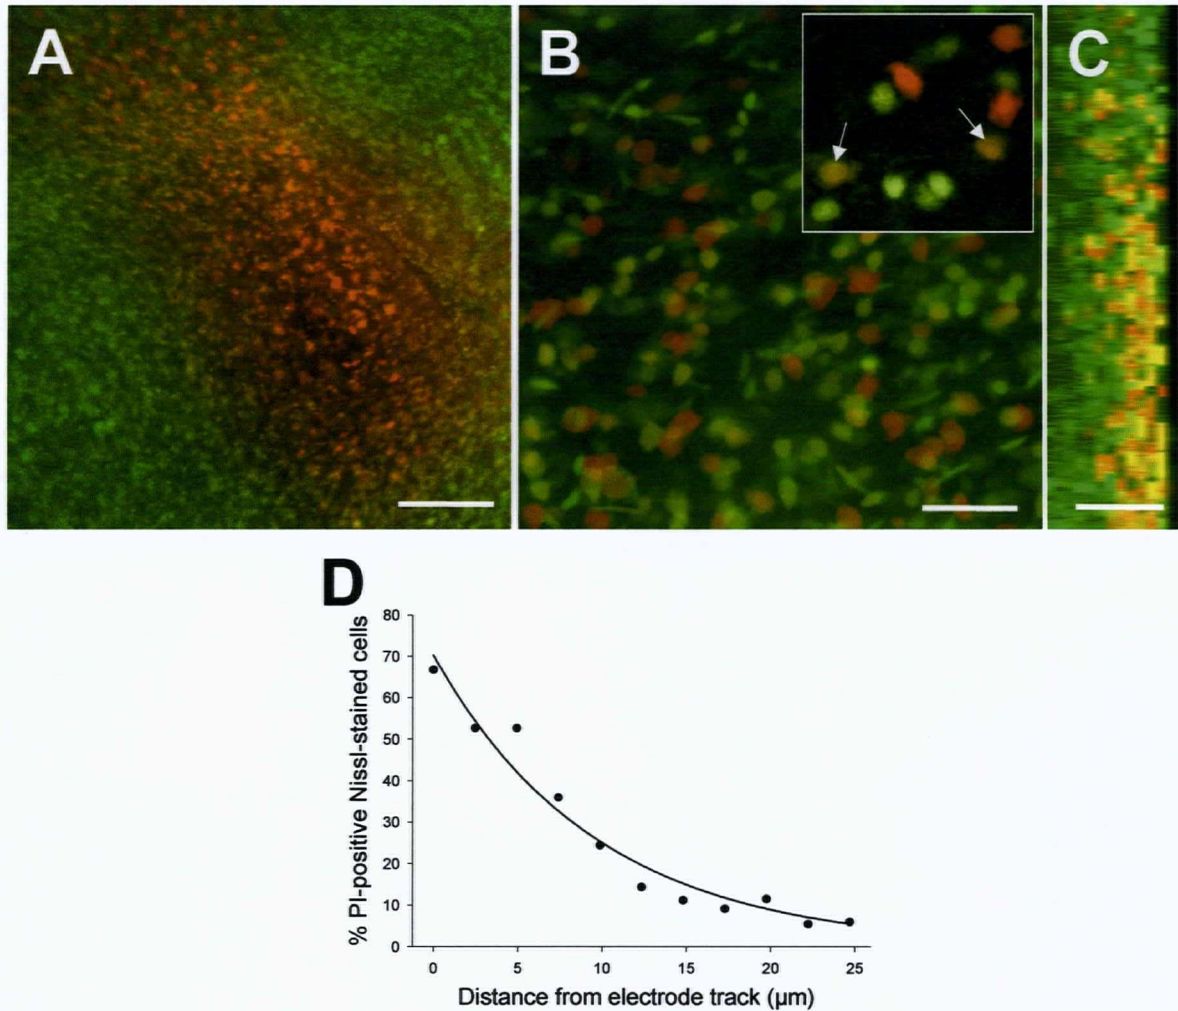


Figure 1.13 Estimation of damage around the polytrode.

(A) Low power photomicrograph illustrating that PI-positive (PI+ve) cells are restricted to the vicinity immediately surrounding the polytrode. Scalebar = 150μm. (B) Confocal reconstruction, 30μm thick, imaged directly in front of the polytrode track. (*Inset*) Green fluorescent cells are undamaged neurons (including some glia), yellow cells are predominantly PI+ve damaged neurons (*arrows*), and red cells are PI+ve non-neuronal cells. Scalebar = 50μm. (C) Rotated projection of the field of view in (B) shows the limited extent of PI+ve cells with compromised membranes. Scalebar = 25μm. (D) The percentage of PI+ve cells decreased exponentially (*fitted line, $r^2 = 0.96$*) with distance from the polytrode track.

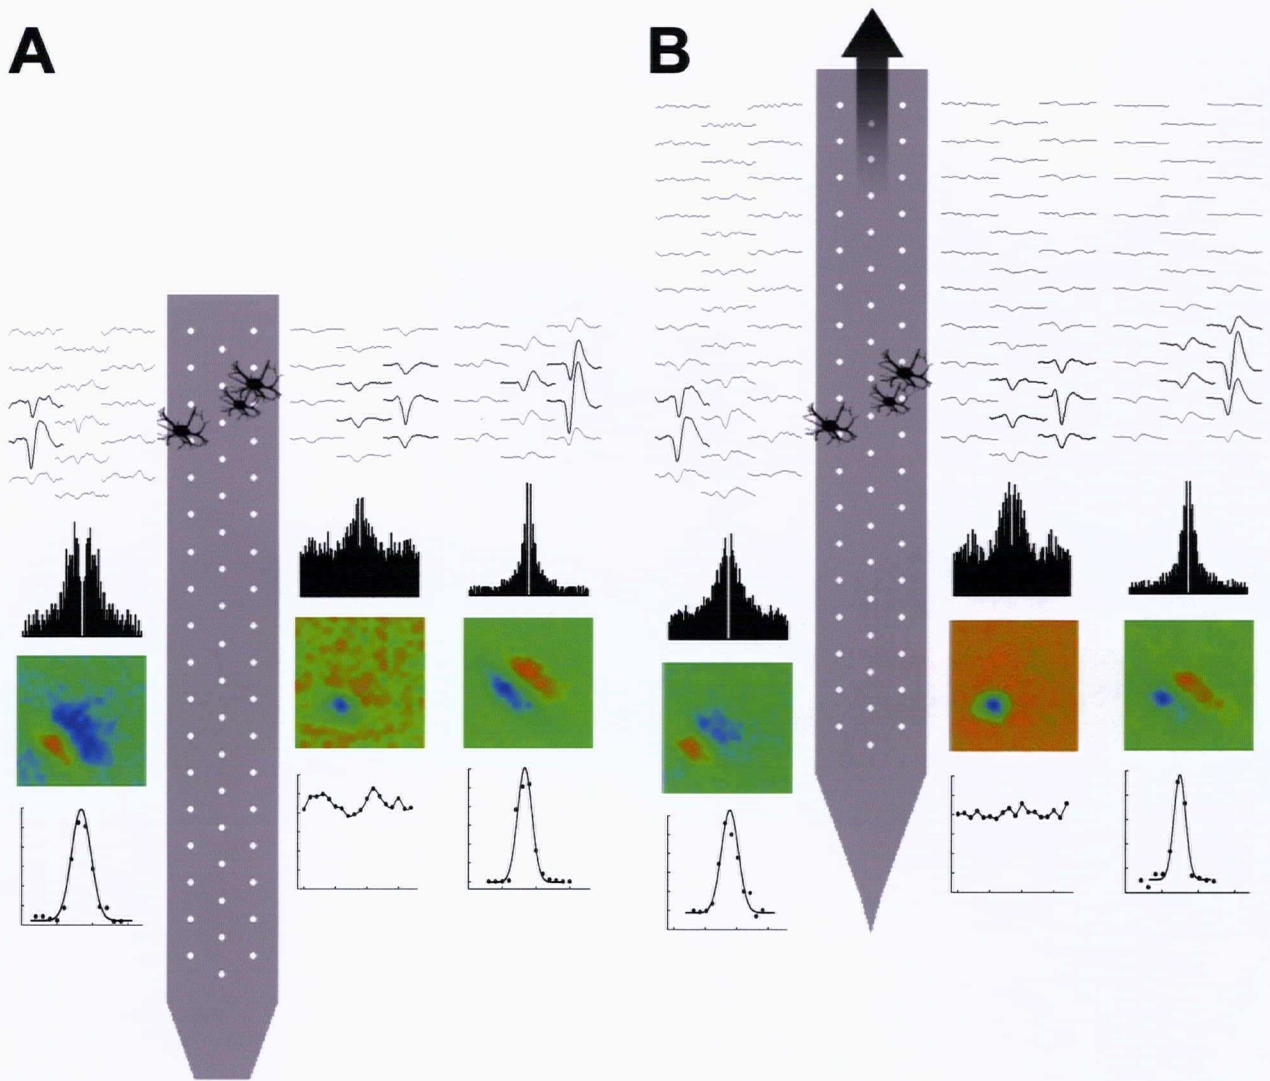


Figure 1.14 Neuron passage study.

An ensemble of active neurons was recorded (A) at the point of insertion, traversing cortical layers IV – VI, and (B) following retraction of the polytrode 400 μ m (layers II/III – V), four hours later. Averaged spike waveform epochs from 3 (of 24) visually responsive neurons are shown at both positions. Although the neurons were recorded on different sites, note the close similarity of the spike shapes and distributions. Estimated neuron locations were virtually identical. Autocorrelograms (± 50 ms, 1ms bin width), receptive field profiles ($8 \times 8^\circ$, 40ms post-stimulus), and orientation tuning curves (normalised to the peak firing rate) characteristic of the 3 neurons were the same at the two positions of the polytrode. Receptive fields were determined using binary m-sequence stimuli and reverse correlation (Jones and Palmer 1987); orientation tuning curves were derived from average responses ($n = 8$) to a drifting bar stimulus (0 - 360° , 20° increments).

1.5 Discussion

The desire to elucidate the function of biological neuronal networks has motivated the development of a variety of technologies able to record simultaneously from multiple neurons. Large scale recording of neuronal activity in the intact brain is considered by many to be a prerequisite for understanding the distributed coding mechanisms that underlie sensory-motor integration, perceptual abilities, learning, memory, and ultimately the neuronal basis of language and higher cognition (Churchland and Sejnowski 1992; Buzsaki 2004; Olshausen and Field 2005). Polytrodes are particularly well suited for this endeavour. No other currently available electrophysiological or imaging technique combines sub-millisecond temporal resolution with single-cell spatial resolution, and the capability to monitor the activity of hundreds of neurons spanning contiguous cortical areas (Campbell et al. 1991; Hoogerwerf and Wise 1994). The polytrodes described in this thesis are ideal for studies of columnar microcircuits because they enable exceptionally high-density recording of unit and field activity with minimal tissue damage. They offer demonstrably better single-unit isolation than single electrodes, stereotrodes, or tetrodes, and provide stable multiunit recordings for hours.

The finding that polytrodes with finer interconnects give robust recordings without increased noise or channel crosstalk augurs well for future polytrodes with even narrower conductors and spacing. The negligible crosstalk that was observed ($<0.5\%$) for $1.5\mu\text{m}$ feature sizes is in accord with predictions of an earlier theoretical study (Najafi et al. 1990) that concluded features could be scaled down to $1\mu\text{m}$ with less than 1% crosstalk. Current industrial scaling limits are in the submicron range, so even smaller high density polytrodes should be realisable in the near future. Another potential concern was the viability of neurons recorded by polytrodes with shank widths greater than $60\mu\text{m}$, a problem that has been reported by other users of similar devices (Csicsvari et al. 2003). However there was not any obvious deterioration in the number or quality of units recorded across the entire shank having a width of $\sim 200\mu\text{m}$ (Figure 1.7B). A narrower shank is nonetheless desirable for minimizing tissue displacement and neuronal damage (Claverol-Tinture and Nadasdy 2004). As stated earlier, the main factor determining the overall shank width, and in turn the maximum number of

recording sites on a single shank, is space for the interconnecting leads. Ultimately one-lead-per-site interconnects may become superfluous for polytrodes comprised of active transistor arrays with on-chip multiplexing. Field effect transistor-based polytrodes with thousands of sites have already been prototyped for *in vitro* applications (Fromherz 2003). However, until serious complications involving the durability and high intrinsic noise of these devices are resolved, passive polytrodes with 1 μ m or less feature sizes, narrower shank widths, and even more recording sites, will continue to provide state-of-the-art high-density multiunit recordings. Other process variations, such as multilevel metal for the interconnect leads, offer the prospect of reducing the shank width of high density passive arrays.

Two potential improvements to the current polytrode designs warrant mention. Firstly, penetration of the meninges could be aided by making the tip angle sharper (Najafi and Hetke 1990), or incorporating a silicon 'spine' on the back of the shank to make it more rigid. The latter can be achieved by withdrawing the polytrode from the final etch before a complete etch-stop is achieved (Najafi et al. 1990) without increasing the overall width of the polytrode. I have also attempted to soften the dura and pia by partially digesting it with collagenase (Zhu et al. 2002b), but this caused extensive spotted bleeding of the pial vasculature, which was deemed more detrimental than a small incision at the point of penetration. Secondly, the recordings reported here were made with unmodified iridium recording sites. Controlled electrodeposition of gold (using the apparatus described in appendix A) can increase site surface area without increasing site diameter, reducing impedances by an order of magnitude (T. Blanche, unpublished results). The lower input impedance of the sites in turn reduces the noise level. Optimizing the electrode-tissue interface by increasing the surface roughness and adhesive properties of the recording sites with synthetic polymers (Cui et al. 2003) is another avenue for enhancing recorded spike amplitudes. Either of these techniques could be adopted to improve the SNR of future recordings. If required, sites may also be modified for extracellular microstimulation following electrochemical 'activation' to increase their charge capacity (Weiland and Anderson 2000).

Despite the voluminous datasets generated (aggregate bandwidth ~2.8MB/s; ~10GB/hour for 54 channels including EEG), continuous acquisition has a number

of benefits. It obviates the need to set trigger thresholds online, which is time consuming and impractical for polytrodes with more than a few sites. It eliminates the possibility of missing or duplicating spike events due to inappropriate window discriminator settings. The standard tetrode approach of recording an epoch from all channels in response to a threshold-crossing is inappropriate for polytrodes extending over several millimetres, and the usual method of 'locking-out' the entire electrode array following a spike event makes detection of synchronous spikes impossible. Continuous acquisition does not solve the problem, but makes possible more sophisticated offline algorithms that can (see section 3.3.3). In any case, there is little bandwidth to be saved by making episodic recordings considering the large numbers of active neurons typically recorded with polytrodes. Furthermore, since no data is lost at acquisition, as new and improved methods of spike sorting are developed the possibility exists to return to the archived files and re-extract the spikes from the continuously recorded waveforms.

The use of stereotrodes and tetrodes drew attention to the value of spatially sampling individual neurons, exploiting differential spike amplitudes on different sites to improve single-unit discrimination. High-density polytrodes take this idea to its logical conclusion by recording from most of each neuron's field potential. The result is a further improvement in the reliability of single neuron isolation. Given the inherent 'myopia' of tetrodes, and the disparity between the number of neurons per virtual tetrode recorded with polytrodes and that usually cited for real wire-bundle tetrodes, the potential for mixed clusters in a highly active tetrode recording is probably considerably worse than that suggested by the case study presented here (Figure 1.9). The optimal sampling density and geometric configuration of recording sites needed to unambiguously resolve the activity of multiple neurons remains an open question. Each of the polytrode designs was, however, capable of recording individual neurons on multiple sites (Figure 1.8F), so the question of 'optimal' site spacing becomes more a question of specific application. Due consideration should also be given to the specific brain region and species under study. Rat hippocampus, for example, is only about $700\mu\text{m}$ thick, but the limited spatial extent of interneuron field potentials in the dentate gyrus requires a polytrode with a site spacing ideally less than $50\mu\text{m}$ (Freund and Buzsaki 1996). By way of contrast, only the two column and one of the three

column polytrodes are long enough (more than 1.2mm) to record simultaneously from all cortical layers of cat extrastriate visual cortex (Beaulieu and Colonnier 1985). In the proposed studies of the columnar circuits underlying receptive field dynamics (chapter 5) translaminar coverage was particularly important, but a sampling resolution higher than 50 μ m was not justified. Finally, it should be noted that higher overall neuron yields might be obtained with multiple tetrodes (Hoffman and McNaughton 2002; Jog et al. 2002) because they can be independently moved to foci of high activity, but not with the high density possible with single-shank polytrodes, and at the expense of knowing the exact anatomical location of recorded units.

There is arguably little value in estimating precise neuron locations (chapter 4) without knowing the location of the polytrode in the cortex with a comparable level of certainty. The online CSD analyses and unfixed fluorescent dye track reconstruction procedures described earlier are well suited for this task, enabling the depth and orientation of the polytrode to be determined with an accuracy of ~25 μ m (Figure 1.11, Figure 1.12). Imprecision caused by diffusion of the diI prior to imaging is of minor consequence because the shape of the polytrode is usually discernable and the coordinates of the recording sites are fixed with respect to the polytrode shank. Traditional methods, such as depositing iron from the electrode that is detected by the Prussian blue reaction (Green 1958; Brown and Tasaki 1961) are unsuitable for polytrodes as the platinum, iridium or gold recording sites are devoid of any iron. DC current lesioning (Hubel and Wiesel 1962) can localise the electrode with at best 100-150 μ m accuracy, even when differential shrinkage associated with tissue fixation is ignored. Dual electrolytic lesions with fixed multichannel silicon electrode arrays can, however, improve the accuracy to within 100 μ m (Townsend et al. 2002). Breaking the polytrode shank *in situ* (eg. Buzsaki and Kandel 1998) is another approach, but the probe is frequently displaced 200 μ m or more from its original recording location when cut or during subsequent histological processing (Bragin et al. 2000). Polytrodes that could otherwise be reused many times (Figure 1.10) are obviously rendered 'single-use' electrodes, a particularly undesirable aspect of this method. Electrode localisation using magnetic resonance imaging (MRI) has a similar margin of error of ~180 μ m (Fung et al. 1998; Jog et al. 2002), and given the high cost of MRI, its application for this purpose is not widely accessible. Compared with the alternative techniques

for electrode track reconstruction, the procedure described in this chapter is simpler, more accurate, and applicable to any single or multi-shank polytrode.

1.6 Next steps...

Having established effective recording and histological procedures, there are several important data processing stages between continuous acquisition of raw multichannel waveforms and the ultimate goal of studying the physiology of large neuronal ensembles. In fact, the majority of this thesis is dedicated to establishing a set of robust methodologies to take full advantage of the unique recording capabilities of polytrodes. The following chapter presents offline signal processing methods for enhancing waveform fidelity, in turn improving neuronal yield and single unit isolation. Chapter 3 describes spike detection and sorting algorithms designed specifically for polytrodes. Chapter 4 describes work aimed at precisely localising neuron location in 3D by modeling the high resolution field potentials measurements recordable with high density polytrodes. The prospect of differentiating interneurons from pyramidal cells is also explored. Chapter 5 introduces the concept of 'cortical micromapping', that is using neuron localisation and classification to bring together neural structure and function on the cellular scale. A new experimental paradigm is proposed for compiling a large multiunit database of unprecedented detail in order to address a variety of outstanding questions in visual neurophysiology, some of which are discussed in this concluding chapter.

chapter 2

Signal processing

2.1 Summary

Polytrodes have the potential to record simultaneously from hundreds of neurons. In practice, neuron yield is limited by spike waveform fidelity, which in turn affects the efficacy of spike detection and sorting algorithms. Judicious application of signal processing methods can help to fully realise the large-scale recording potential of polytrodes. Ideal bandlimited interpolation with sample-and-hold delay correction is shown to accurately reconstruct spike shapes, improve the reliability of threshold-based event detection, and facilitate accurate spike sorting by reducing waveform variability. A principal component analysis based denoising algorithm was explored but found to cause significant attenuation of small spikes while conferring only minor overall improvement in signal to noise. Rejection of transient noise artefacts was, however, achieved with a simple averaging procedure. Interpolation of high bandwidth data is computationally expensive. A cost-benefit analysis was made of interpolation rates ranging from 12.5kHz (Nyquist frequency, no interpolation) to 100kHz, taking into consideration the final application of the data. For most purposes data acquisition rates of 20–25kHz with offline interpolation to 50kHz was found to be ideal, with negligible gains above this rate. A practical benefit of upsampling is that storage requirements can be greatly reduced by sampling at or slightly above the Nyquist frequency, although the potential for aliasing at lower sample rates should be considered.

2.2 Introduction

Inadequate sampling or poor reconstruction of spike waveforms can result in errors of spike detection, distortion of spike shapes, and otherwise compromise the performance of spike sorting algorithms. Neurophysiologists routinely record spike activity at sample rates between 20 and 40kHz, yet according to sampling theorem (Nyquist 1928; Shannon 1949), for signals low-pass filtered at 5kHz, a 10kHz sample rate (*i.e.* the Nyquist frequency) should be adequate. Data sampled according to Nyquist criteria do not *appear* to provide an accurate representation of the original spike waveforms (Figure 2.1), hence the common practice of oversampling. However, the Shannon-Nyquist theorem (Shannon 1949) states explicitly that if a band limited signal $s(t)$ is sampled at a rate $2W$, such that its Fourier transform:

$$\mathfrak{F}[s(t)] = S(f) = 0 \text{ for } |f| > W$$

... then $s(t)$ is completely determined, and can be recovered from its samples, s_n :

$$s(t) = \sum_{n=-\infty}^{\infty} s_n \frac{\sin \pi(2Wt - n)}{\pi(2Wt - n)} \quad (2.1)$$

The Nyquist sampling rule is often misconstrued as implying that the original signal is recoverable from the raw samples without further processing. What is usually neglected is the correct application of equation 2.1 in the waveform reconstruction. A discrete implementation of this theorem using convolution was the basis for the bandlimited interpolation, or so called 'upsampling', of continuously acquired waveforms described in this chapter. Although this method of signal reconstruction is not new, I describe it here as it is rarely exploited by neurophysiologists. It deserves more attention because there are tangible benefits of upsampling with respect to reliability of spike detection, estimation of spike waveform parameters, and accuracy of spike sorting.

Equally important for successful isolation of single units in multiunit recordings is minimising contamination from noise. Radio-frequency interference, electrical line noise, ground loops, thermal (Johnson) noise, and static discharge all spoil the fidelity of spike recordings. Most of these exogenous noise sources were reduced or even eliminated in the present studies by the use of high input impedance headstage preamplifiers, the common mode rejection of the pseudo-

differential reference electrode, careful attention to proper shielding, and 'star grounding' configurations (Appendix B.1). Low frequency LFPs were removed by the high pass filter of the spike amplifier channels by design. Other sources are, however, inextricably part of extracellular multiunit recordings. Background neural noise, including stimulus related 'hash', arises from the spike discharges of distant neurons, which obscure the spikes of neurons adjacent to the polytrode, especially those with small amplitudes. The observation that neural noise is highly correlated and, compared with nearby spikes, spatially invariant (Rebrik et al. 1999) led to the development of linear array processing (Bierer and Anderson 1999) and principal component analysis (PCA) based algorithms (Musial et al. 2002) for selective removal of this noise. Both cited methods are mathematically equivalent (Musial et al. 2002), and the PCA cleaner was adapted here for use with the high-density contiguous recording sites of polytrodes.

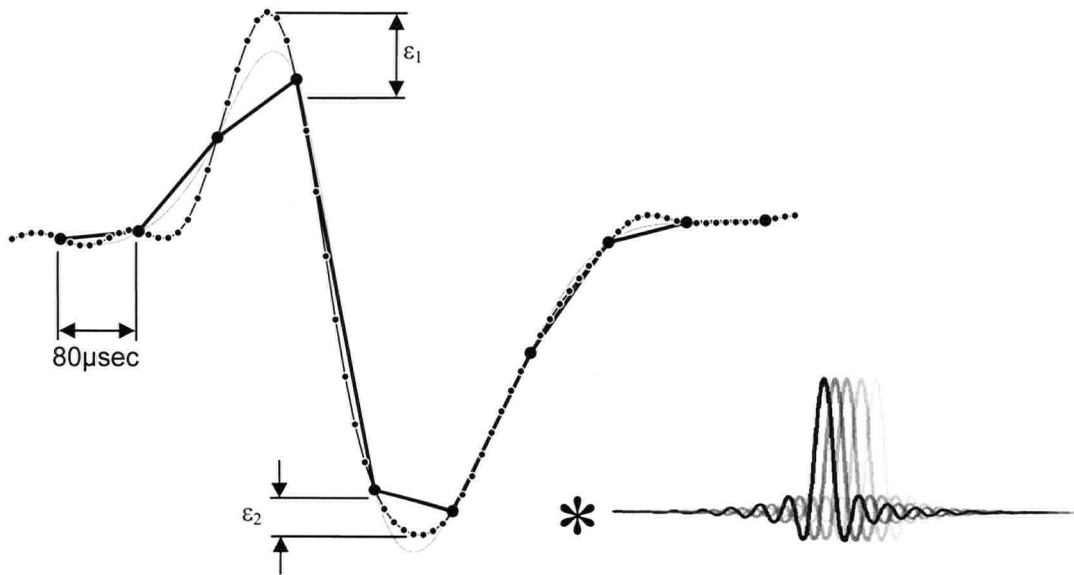


Figure 2.1 The case for upsampling.

A typical spike ($\sim 240\mu V_{pp}$ amplitude) sampled at the Nyquist rate (12.5kHz, *large circles*), linearly interpolated (*thick line*) provides a poor rendition of the original waveform sampled at 100kHz (*fine black line*) particularly at the spike onset, peak and valley. Spike amplitudes are consistently underestimated, and threshold-based event detectors set in the range ϵ_1 and ϵ_2 will miss a percentage of spike events. After interpolation to 100kHz (*small circles*) an accurate reconstruction of the underlying waveform is obtained. The cubic spline interpolated waveform (*fine grey line*) is shown for comparison. The inset shows the time-shifted windowed sinc kernels, one for each interpolant of the reconstructed waveform.

Noise transients caused by static discharge are more difficult to suppress. Due to the filtering properties of the amplifiers, these manifest as spike-like artefacts that typically appear on all electrode sites. Distinguishing these from real spike signals requires special consideration.

This chapter presents offline signal processing methods intended to improve waveform fidelity, spike detection and sorting, and ultimately neuron yield in multiunit polytrode recordings. Three techniques are described: bandlimited interpolation founded on the Shannon-Nyquist theorem; a PCA-based filtering algorithm for improving SNR; and a simple algorithm for rejecting transient noise artefacts. The effectiveness of each method was quantitatively assessed.

2.3 Methods

2.3.1 Bandlimited interpolation

The ideal method for reconstruction of analog signals sampled according to Nyquist criteria is, in the time domain, convolution with a sinc function:

$$s(t) = (s * h_s)(t) = \sum_{n=-\infty}^{+\infty} s(n)h_s(t - n) \quad (2.2)$$

where t = the time of the interpolant

$s(n)$ = the raw waveform samples

$h_s = \text{sinc}(f_s t) = \frac{\sin(\pi s)}{\pi s}$ = the sinc reconstruction filter

The sinc (meaning 'cardinal sine') function (Figure 2.1) is theoretically the perfect reconstruction filter since its Fourier transform is a box centered around DC of unitary width. It is, however, an infinite impulse response (IIR), as any signal that has a finite extent in the frequency domain must have an infinite extent in the time domain (and vice versa). The sinc keeps oscillating with ever decreasing amplitude, so for practical purposes it needs to be windowed. The rationale for windowing is to attenuate the sinc to zero at the extremes to produce a finite impulse response (FIR) filter kernel. By tapering the sinc function instead of simply truncating it, ripple in the transition band-pass (Gibbs' phenomenon) can be minimised. There are many choices of windows, and for interpolation the

choice is not critical. Commonly used symmetrical windows include the Bartlett, Hanning, Hamming, Blackman, Kaiser, and Lanczos windows. Each provide a subtly different trade-off between bandpass flatness, ringing and other reconstruction artefacts (Horowitz and Hill 1989; Smith 1997).

The sinc filter should be represented with sufficient resolution such that interpolating linearly between upsampled points does not introduce error greater than the quantization error of the analog to digital conversion (ADC). For 12-bit data, a filter kernel of length $n = 19$ with 9 'zero crossings' provides adequate precision (Smith 2004), yet is computationally feasible for the large signal bandwidth of polytrodes (1.35M.s^{-1} for 54 channels sampled at the Nyquist rate, $\sim 11\text{M.s}^{-1}$ upsampled to 100kHz per channel). Combining equation 2.2 with a Hamming window, w_h , gives the tapered FIR filter suitable for discrete time series convolution:

$$s(t) = \sum_{n=-9}^{+9} w_h s(n) h_s(t-n) \quad (2.3)$$

where $w_h = 0.54 - 0.46 \cos(\frac{\pi s}{n + \pi})$ (or other symmetric window function)

2.3.2 Sample and hold delay correction

Most acquisition cards have a single analog to digital converter and therefore cannot sample multiple channels simultaneously. As with the DT3010 cards described in this thesis, the usual solution is to combine sample-and-hold circuits with rapid signal multiplexing to obtain 'near-simultaneous' sampling across all channels. Resultant sample-and-hold delay (SHD) artefacts (*e.g.* Figure 2.3A) were corrected during the interpolation process by using appropriately phase-shifted sinc kernels, without additional computational overhead:

$$s_i(t) = \sum_{n=-9}^{+9} w_h s(n) h_{si}(t-n+t_i), i = 0 \dots N \quad (2.4)$$

where i is the ordinal position of the channel in the hold queue, and t_i is the appropriate phase offset, a multiple (or fraction) of the sampling period.

The phase-shifted windowed sinc kernels, h_{si} , one for each interpolant, for each SHD offset, were pre-calculated and stored in indexed arrays for efficient implementation of the convolution. The question of the optimal level of

oversampling was addressed empirically for interpolation factors ranging from one (Nyquist sampled, or no interpolation) to eight (giving an effective sample rate of 100kHz for data acquired at 12.5kHz). The test data comprised continuously sampled recordings from each of the polytrode designs.

2.3.3 Noise and artefact suppression

PCA is typically used to reduce the dimensionality of spike waveforms for spike sorting (Abeles and Goldstein 1977; Gray et al. 1995). It can also potentially improve the SNR of multichannel recordings by predicting the common mode noise present across electrode sites (Musial et al. 2002). PCs were calculated by first computing the normalised covariance, C , of the multichannel data, D , for each electrode site j of the n -channel polytrode, excluding the site being processed:

$$C(t, s_{\neq j}) = \text{cov } D(t, s_{\neq j}) = \frac{1}{n-1} \sum_{i=1, i \neq j}^n (s_i - \mu_i)(s_i - \mu_i)^T \quad (2.5)$$

The covariance matrix was then diagonalised by eigenvector rotations to find the largest eigenvectors (*i.e.* axes of principal variance), e , according to the largest eigenvalues, λ (Jolliffe and Morgan 1992). For each successive electrode site the algorithm computed the first three PCs derived from all other sites:

$$(C - \lambda_i I)e_i = 0, \quad i = 1, 2, 3 \quad (2.6)$$

where C and I are the covariance and identity matrices, respectively. Data from the site being processed were projected onto each of these PCs to give a set of normalised coefficients, n_k , representing the fraction of raw signal amplitude attributable to noise:

$$n_{k,j} = \frac{D(t, s_{\neq j})^T e_k}{e_k^T e_k}, \quad k = 1, 2, 3 \quad (2.7)$$

These components were then subtracted from the raw data of the site being filtered, and the entire process was repeated for every site in the polytrode array:

$$D_{\text{clean}}(t, s_j) = D(t, s_j) - \sum_{i=1}^3 n_{i,j} e_{i,j} \quad (2.8)$$

A second iteration of the algorithm was intended to correct for the influence of spike trains on other sites (Musial et al. 2002). Matlab (MathWorks Inc., MA) was used for the exploratory work described here.

In the original description by Musial and colleagues (2002) only the electrode site being filtered was excluded from the calculation of the noise PCs. This was appropriate for their simulated spike train data because spikes were seeded on only one channel (they further claim this was suitable for real data recorded with electrode sites spaced at least 100 μ m apart, however no empirical evidence was given to support this suggestion). To make the algorithm applicable to high-density polytrodes, where spikes from individual neurons appear on many adjacent sites, an 'inner lock-out radius' was incorporated whereby only sites outside the inner lockout were included in the PC estimations. Since the spatial extent of the dominant correlated noise sources had not been explicitly studied, a second 'outer-lockout' parameter was used to define the number of sites to include. Evaluation of the optimal combination of lockout parameters against overall SNR was made on short 3-10s data segments randomly selected from several polytrode recordings.

In addition to PCA-based filtering a simpler method was sought for suppressing cross-channel transient artefacts such as those caused by static discharge and electrical stimulation. Amongst those that were tested (*i.e.* covariance, cross-correlation, averaging and various derivative functions) a simple threshold factor, f , times the standard deviation, σ , of the n -channel arithmetic mean was chosen:

$$X[s(t)] = \frac{1}{n} \sum_{i=1}^n s_i(t) = 0 \text{ for } |X| > f \cdot \sigma \quad (2.9)$$

This is essentially a digital common-mode filter with a binary rejection criteria. For efficient implementation it was used to validate potential spike events only after they were detected (as described in chapter 3).

2.4 Results

2.4.1 Waveform reconstruction

Interpolation using windowed sinc convolution accurately reconstructed extracellular spike waveforms (Figure 2.1, Figure 2.2A, C). The representative spikes in these figures show that un-interpolated data, even when sampled according to Nyquist criteria, gave a poor rendition of the underlying spike shape. Accordingly, raw estimates of spike amplitude (Figure 2.2B, D) were consistently low, and spike to spike variability attributable to the sampling process was high. If spike-triggered averaging is used to generate spike templates and derive waveform metrics, as was the case in these examples, then amplitude variability is compounded by the temporal misalignment of the individual spike events.

Sinc interpolation increases the effective sampling rate, with concomitant improvement in waveform fidelity. There was a gradual improvement in spike shape with increased sampling rate, most noticeably for fast spikes and in regions of the waveform with high slew rates, including the peak and valley. Slower regions of the spike waveform, such as the repolarisation and after-hyperpolarisation phases, were in a sense already oversampled because they are comprised of lower frequencies. For example, the neuron in Figure 2.2B had an average estimated spike amplitude of $172\mu\text{Vpp}$ at 100kHz , but only $117\mu\text{Vpp}$ at 12.5kHz , 68% of the actual amplitude. Without adequate interpolation, average estimates of spike width were relatively accurate ($108\mu\text{s}$ at 12.5kHz vs. $96\mu\text{s}$ at 100kHz for the fast-spiking neuron), however trial-trial variability caused by temporal aliasing ($80\mu\text{s}$ at 12.5kHz) was unacceptably high. Spike width standard deviation was $44.8\mu\text{s}$ at 12.5kHz compared with only $5.3\mu\text{s}$ at 100kHz (*i.e.* it was proportional to roughly half the sampling period). The poor accuracy and precision of un-interpolated Nyquist-sampled data was inadequate for modelling neuronal field potentials and classification of cell type based on spike amplitudes and widths (chapter 4). For this reason all modeling was of spike waveform data upsampled to 100kHz .

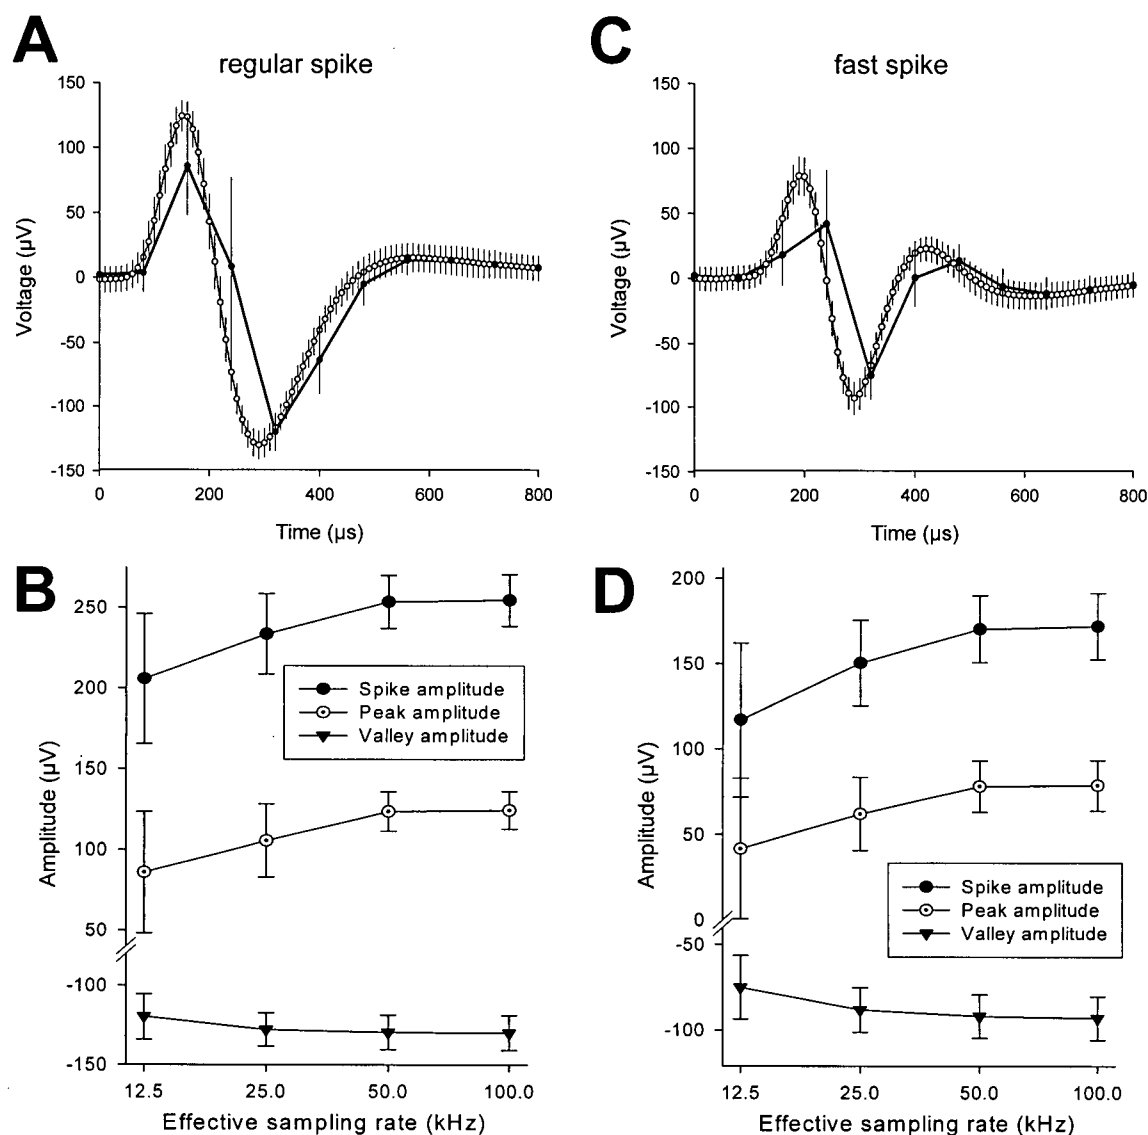


Figure 2.2 Spike waveform reconstruction.

Interpolation of Nyquist-sampled data faithfully reconstructs spike waveforms. (A) An averaged ($n = 313$) medium amplitude spike was crudely represented at 12.5kHz (—●—), whereas upsampling to 100kHz (—○—) accurately matched the underlying spike shape (grey line, acquired at 100kHz). (B) Waveform-derived spike parameters reflect the increased variance and underestimation of peak and valley amplitudes at low sampling rates. (C) A fast-spiking neuron ($n = 1640$) had even higher variability and distortion without adequate interpolation, (D) with consequent misrepresentation of spike parameters. Error bars are the mean \pm standard deviation.

While SHDs of a few microseconds are of no consequence for tetrodes, with polytrodes they accumulate to produce significant phase disparities in waveforms recorded across sites (Figure 2.3A, B). The SHD of a given site was proportional to

its channel index times the fastest 'burst-mode' sample interval of $1\mu\text{s}$ per channel (for the 32 channel DT3010 acquisition cards). The maximum phase *advance* caused by SHD was thus only $31\mu\text{s}$, however this was sufficient to corrupt estimation of the conduction velocity of BPAPs (Figure 2.3C), in addition to the instantaneous spike amplitude measurements needed for modelling spike field potentials. SHD artefacts were therefore routinely compensated for (equation 2.4) during the interpolation process (Figure 2.3, *lower panels*).

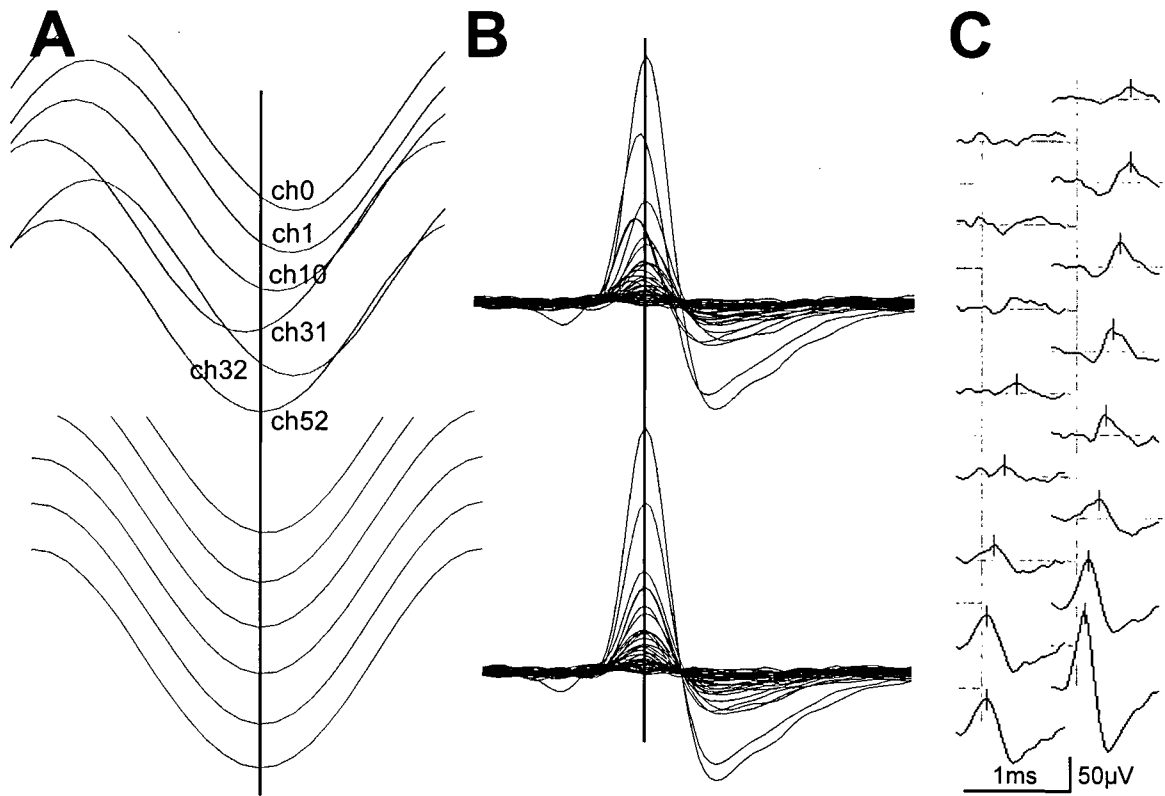


Figure 2.3 Sample & hold delay correction

(A) A 1kHz sine-wave signal sampled 'near simultaneously' on different channels of the two 32-channel data acquisition cards. Accumulation of SHDs produces artefactual phase misalignments (*upper panel*). The same waveforms interpolated with SHD correction (*lower panel*) are perfectly aligned (note that ch0/ch32, ch1/ch33, etc. are the corresponding pairs of channels on each DT3010 acquisition card, and are therefore subject to the same SHD). (B) A similar demonstration of the effect of SHDs on a multi-channel spike sampled across 18 recording sites. (C) SHDs are sufficient to mask precise measurements of biological phenomena such as the conduction velocity of BPAPs.

2.4.2 Reliability of event detection

Variability in the amplitude of spike peaks caused by digital sampling can result in missed spike events (threshold ranges ε_1 and ε_2 in Figure 2.1). The magnitude of this error was quantified by counting the number of spikes from a single neuron interpolated to different effective sampling rates over a range of trigger thresholds. In the first of two representative examples (Figure 2.4A) a large-amplitude neuron fired 303 action potentials, all of which were detected at a threshold level of +250 μ V or lower, irrespective of the sample rate. As the threshold was increased beyond the spike's peak amplitude, the number of triggers decreased, so that by +450 μ V no spikes were detected. The sigmoidal relationship (*i.e.* $y = a/(1 + e^{(x_0 - x)/b})$) between the number of detected events and the threshold level is characteristic of integration over a Gaussian distribution. Broader distributions produce shallower slopes, and *vice versa*, reflecting the combined variability of the spike amplitudes and the discrete, asynchronous nature of the digital sampling. Lateral shifts in the sigmoid result from changes in the mean spike amplitude at different sampling rates. Changes in the slope of the sigmoid (and hence the range over which the same spikes were detected) at different sampling rates indicate differences in spike amplitude variance attributable to sampling. For the large amplitude spike the shallower slope at 12.5kHz ($b = 23$) compared with that for 25kHz and 50kHz ($b = 16$) represents a small increase in sampling related spike amplitude variability. However, there was little benefit of higher sampling rates in terms of detection reliability, because lowering the threshold to +250 μ V gave perfect detection without an increase in spurious triggers from noise or other spike events. This was not the case in the second example (Figure 2.4B), where the broader spread of amplitudes at 12.5kHz ($b = 12$ vs. 8) meant that all spikes from this neuron ($n=1650$) could only be delineated from the background noise after interpolation. Even at the most favourable threshold (Figure 2.4B, *dotted line*), just 80% of the Nyquist-sampled spike events were detected, and lowering the threshold produced an exponential increase in the number of false positives triggered by other small amplitude neurons and noise. In all cases examined ($n = 15$), the proportionally largest improvement in detection reliability was from 12.5–25kHz, with no further gains beyond 50kHz.

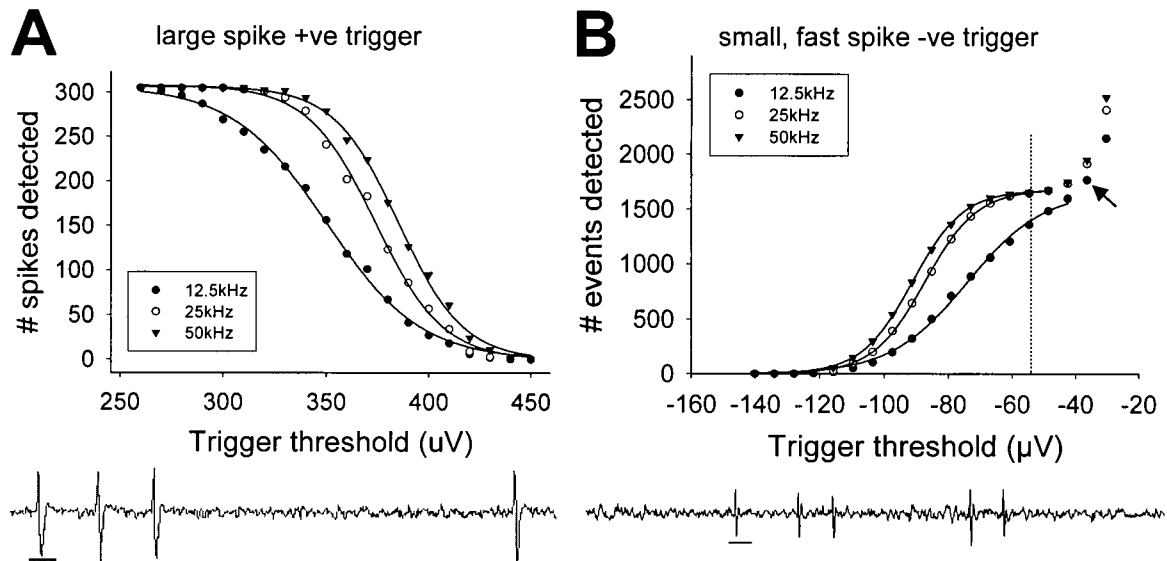


Figure 2.4 Interpolation improves event detection reliability.

Different trigger thresholds were required to detect the same spike events from a given neuron at different sampling rates. The fitted curves are sigmoidal functions (all $r^2 \geq 0.99$). (A) For large amplitude solitary spikes like that shown here (*lower panel*), upsampling gave no benefit in terms of detection reliability. All spikes from this neuron could be distinguished from other spikes and background noise at any sample rate simply by adjusting the threshold. (B) Fast spiking, lower amplitude spikes, however, were often only reliably detected at 25kHz or more. At the 'ideal' threshold for higher sampling rates (-55uV, *dotted line*) not all spikes were detected at 12.5kHz, but lowering the the absolute value of the threshold dramatically increased the number of false triggers (*arrow*). Note that in both examples the 100kHz plots were removed for clarity as they were identical to the 50kHz curves. Scalebars = 1ms.

2.4.3 Cluster variability

In order for spike clustering methods to effectively isolate single units, clusters from different neurons need to be clearly defined and have minimal overlap with the noise cluster along at least one feature dimension. One avoidable source of cluster scatter is sampling related variability (Figure 2.5). As before, the largest improvement was obtained upsampling from 12.5 to 25kHz. Two, four and eight-times oversampling each approximately halved the standard deviation of spike amplitude clusters compared to those generated from Nyquist sampled data (e.g. 28 to 15uV for the spikes shown in Figure 2.5A, 36uV to 21uV for Figure 2.5B). For spike sorting based on dimension-reduced waveform features, for

example PCA, band limited waveform reconstruction with precisely aligned spikes translated into a progressive reduction in cluster variability (Figure 2.5C). Finally, Figure 2.5D illustrates how the boundaries of different neuron clusters can be more easily delineated from each other and the noise cluster at sample rates above 12.5kHz. The cluster separation indices (Equation 1.2) were less than 2 for the overlapping raw clusters, and 6 or more for the 100kHz upsampled clusters.

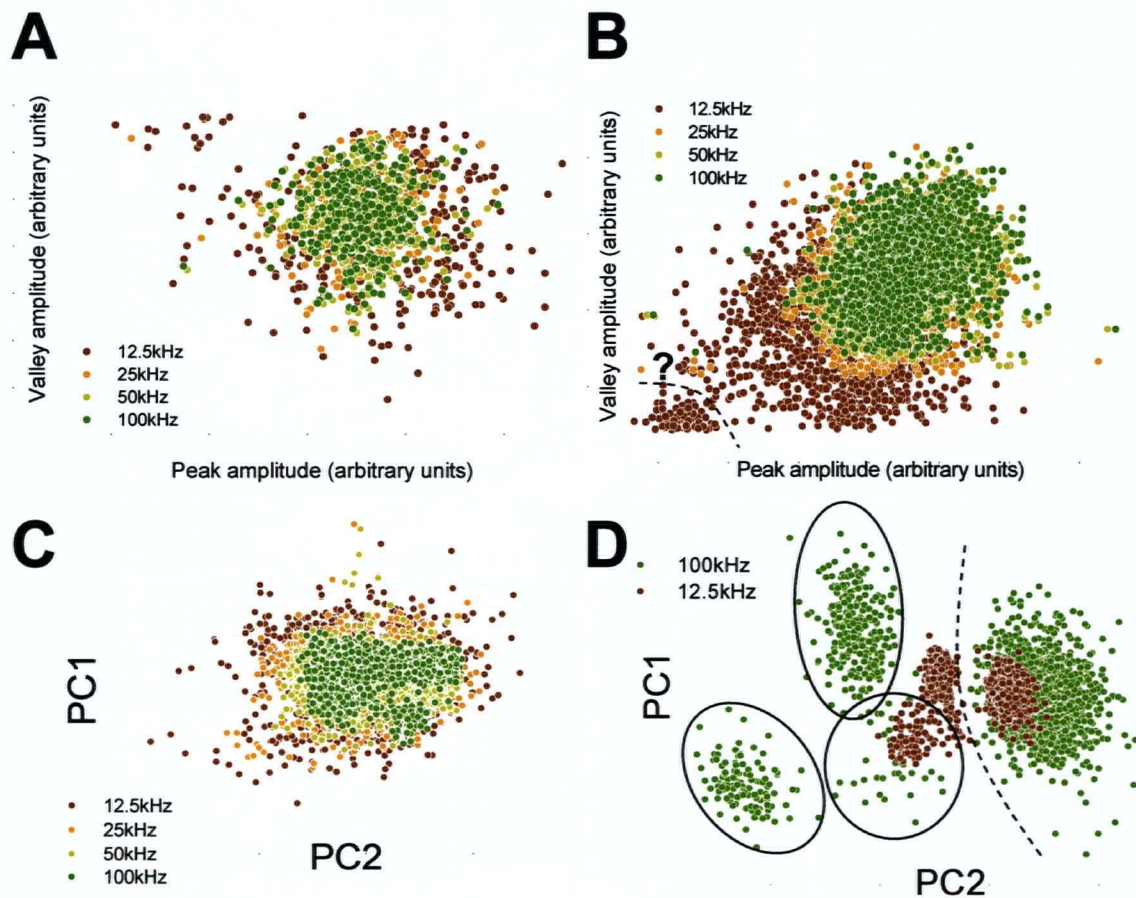


Figure 2.5 Interpolation reduces cluster variability.

Regardless of the features used for spike sorting, higher sampling rates yielded lower cluster variability. (A) Spike peak vs. valley amplitudes from a single neuron ($n=316$), mean-subtracted to show the cluster spreads at each of the four sampling rates tested. (B) Uncentred spike amplitude clusters from another neuron ($n=1253$). Underestimation of spike amplitudes produced a drift in cluster centres at low sample rates. High cluster variance, particularly at 12.5kHz, made separation of spikes from noise difficult (*dashed line*). (C) first vs. second PC scatter plots of the spike events in B. (D) PCA-based cluster plots of spikes upsampled to 100kHz produced three neuron clusters (*ellipses*), each distinct from each other and the noise cluster (*dashed line*). Clusters of the same spikes sampled at the Nyquist rate were confounded, and barely separable from noise.

2.4.4 Artefact rejection and denoising

The similarity of common-mode noise transients across sites was used to identify spurious triggers caused by these events (Figure 2.6). Large differences in the cross-site average amplitude of noise transients compared with real spike events meant a simple threshold (*i.e.* three times the standard deviation of the mean, equation 2.9) could distinguish one from the other without exception.

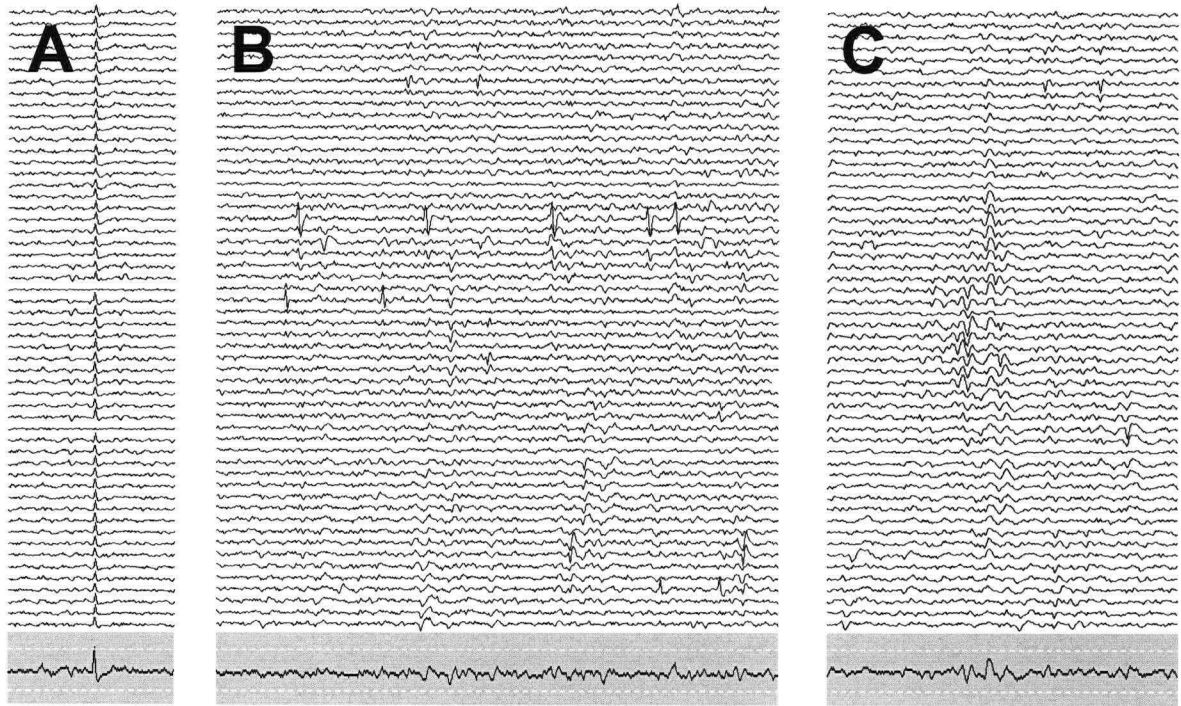


Figure 2.6 Robust artefact rejection.

(A) A typical artefact, in this example a voltage transient caused by a static discharge, was ubiquitous on ungrounded sites. The uniformity of the artefact was reflected in the mean (*shaded waveforms*, scaled 3-fold compared with raw waveforms). (B) Neuronal discharges, even large amplitude spikes that appear on many channels (C), were not strongly manifest in the average waveform. A simple bipolar threshold (*dashed lines*) could therefore be used to distinguish neuronal spikes from these sorts of artefacts.

More sophisticated forms of noise suppression were relatively unsuccessful. The PC-based filtering algorithm had a tendency to attenuate and distort real spike events, with only a marginal improvement in SNR (Figure 2.7). Generalised synchronous burst noise (Figure 2.7A) was broadly correlated across sites, with an average correlation coefficient of 0.43 (range 0.11–0.90). Background noise was

also correlated across sites, albeit to a lesser degree (average 0.16, range 0–0.42). In each case these correlations were almost entirely removed by the PC filter (0.09 and 0.02, respectively, Figure 2.7B). Nevertheless, only small improvements in overall SNR were obtained for spikes in background noise (for inner lockouts greater than 100 μ m). The largest increase in average SNR was 4.6% for an inner lockout of 300 μ m, however average spike amplitudes were attenuated by 6% (Figure 2.7C). Inclusion of the site being filtered or sites immediately around it in the PC calculation actually lowered the SNR because the spikes themselves were diminished by up to 80% (an average of 22% if every site was included). The modest gain in overall SNR might have been beneficial had the algorithm not predominantly attenuated the amplitude of low-amplitude spikes (Figure 2.7D). In the example portrayed, which was characteristic of the five recordings examined, small spikes below 75 μ Vpp were reduced on average by 6.1%, compared with 1.5% for spikes above 100 μ Vpp. Since bursting-related ripple and background noise were spatially uniform on the scale of the polytrodes (1–2mm), changing the radius of the outer lockout parameter did not affect these results.

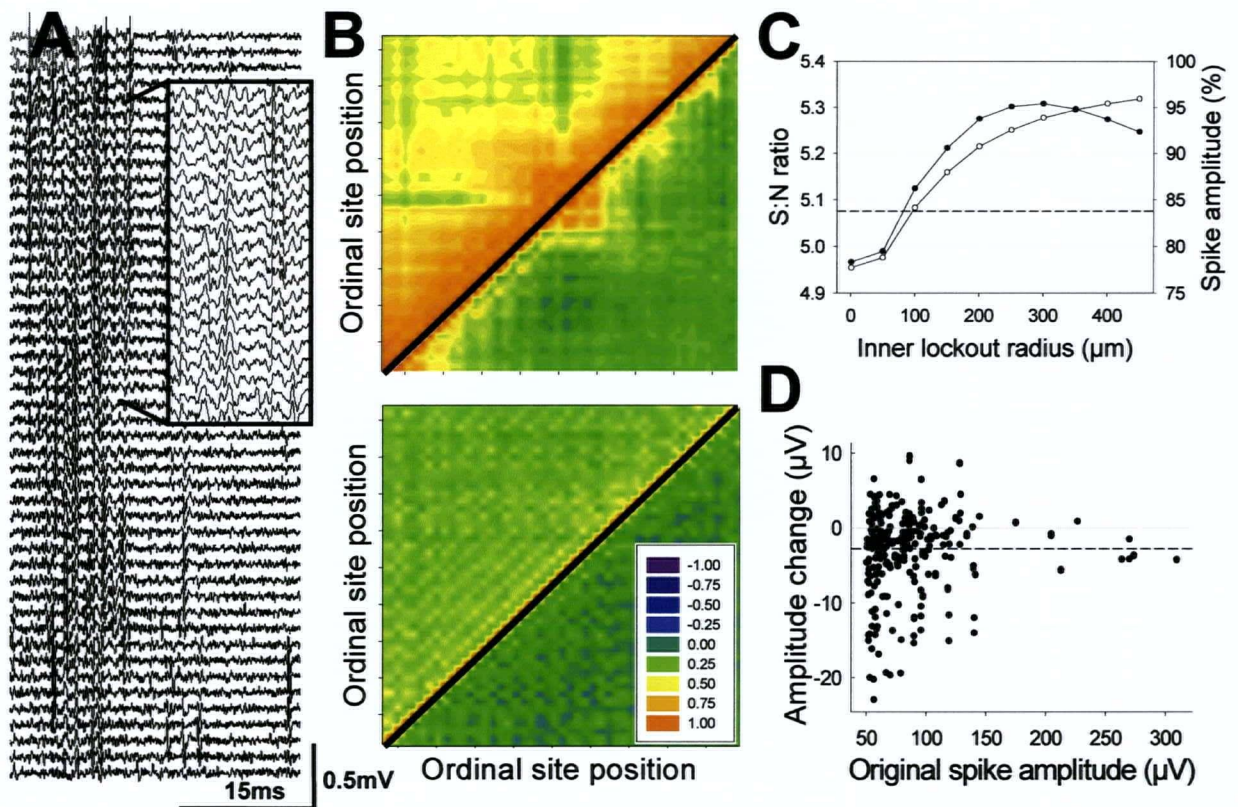


Figure 2.7 PCA-based denoising.

Correlated noise sources can be subtracted from raw spike waveforms to improve SNR. (A) Recording sites shared common noise, most notably during high-frequency synchronous bursts (*inset*). (B) Covariance matrices for two segments of polytrode data are shown. In each example the upper half of the matrix shows the degree of correlation in the raw data, whereas the lower half is the same data after denoising. Sites are ordered by their distance apart, and faulty sites were excluded. High-frequency ripple gave rise to large cross-channel correlations (*upper panel*), that were effectively removed by the denoising algorithm. Background noise showed weaker correlations (*lower panel*) that were also suppressed by denoising. (C) Overall SNR (—●—) and relative average spike amplitude (—○—) post-denoising depended on whether adjacent sites were included in the PC estimation (the *dashed line* shows the raw SNR). (D) On average denoising decreased spike amplitudes (*dashed line*), smaller spikes more so than large.

2.5 Discussion

The task of spike sorting (chapter 3) is challenging enough without the added variability caused by deficient waveform reconstruction. Linear interpolation of analog signals sampled at the Nyquist frequency cannot possibly capture the true shape of a spike waveform. Accordingly spike metrics that are needed for later

analyses such as spike sorting will be compromised. While the usual solution is to oversample, in this chapter I have demonstrated that windowed sinc interpolation can achieve the same end (as it should on theoretical grounds) without the burden of collecting and storing oversampled data. In addition I have shown that SHD correction removes systematic sampling delays that can have a deleterious effect on cross-channel spike waveform indices that depend on truly simultaneous cross-channel measurements.

That windowed sinc interpolation restores the original spike waveform does not imply that information over and above that contained in the raw data has been added. Rather, the information already contained in the data is being utilised properly according to sampling theorem. Upsampling 12.5kHz data to 100kHz will not produce waveforms with frequencies over 6kHz, since, if the data were sampled according to Nyquist criteria, these were already filtered out of the original data. Therefore, if only Fourier methods are used in the subsequent analysis of the raw data (*e.g.* McGill and Dorfman 1984; Kayser et al. 2003b) there would in fact be no advantage to interpolating in the time domain. However, most electrophysiological analyses such as spike detection, slope and amplitude estimation, peri-stimulus time histograms, burst analysis, reverse correlation and so forth are done in the time domain, and it is here that the benefits of correct interpolation will be evident.

Interpolation by convolution is computationally intensive. Both the quantity of data that must be managed and the time it takes to be processed is proportional to the rate of interpolation (Table 2.1). All significant improvements in accuracy and precision afforded by interpolation had been attained by 50kHz. Upsampling to 100kHz did not, for any practical purpose, justify the doubling in computation time or bandwidth. There is no single optimal rate of oversampling; it depends on how the data will ultimately be used. For analyses exclusively in the frequency domain, sample at the Nyquist frequency. If the goal is to generate tuning response curves from average spike counts, then for large amplitude spikes that are already well isolated it could be argued that undersampling is sufficient. Any missed spike triggers would be randomly distributed across trials, and the shape of the tuning curve would still be accurate. At the other extreme, attempting to infer intracellular state from multichannel extracellular waveform distributions (Henze et al. 2000), or studies of spike timing precision (Mainen and Sejnowski

1995; Oram et al. 1999), will benefit considerably from upsampling. Interpolation may also be applied to specific stages in the data processing stream. For example, the quality of PCA-based spike sorting relies on precise spike alignment (Lewicki 1998), however the realism of the spike waveform is unimportant so long as the shape is consistent. By upsampling to 100kHz to ensure reliable spike detection and peak alignment, followed by decimation to 10kHz (or less) to speed the calculation of the PCs, processing of large multiunit recordings from polytrodes can be made more efficient while improving the quality of sorting (data not shown). Likewise, waveform subtraction algorithms that attempt to decompose overlapping spikes (Atiya 1992; Lewicki 1994) perform poorly if the canonical spike templates are inaccurate. Misalignment of the spike template and the measured waveform can introduce spurious spike-like artefacts. To keep these residual artefacts below the RMS noise level, Lewicki (1994) estimated the required alignment precision to be in the range of 2-10 times the sampling rate, similar to that suggested here.

Table 2.1 Cost-benefit analysis of upsampling.

Interpolation factor	Effective sampling rate (kHz)	Spike detection	Spike amplitude	Spike shape	Spike width [†]	CPU load (54 channels per 100 ms)*	Bandwidth (KB/chan/s)
1x	12.5	+	+	+	+	≤ 2ms [†]	25
2x	25	+++	+++	++	+++	22.2±0.3	50
4x	50	++++	++++	+++	++++	44.6±0.5	100
8x	100	++++	++++	++++	++++	87.4±1.8	200

The benefits of spike waveform interpolation at different effective sample rates, against computational load and bandwidth requirements. Rating scale + to +++, relative to raw data sampled at the Nyquist rate (more pluses denote higher accuracy and precision).

* average ± the stdev. of 100 measurements. [†]overhead for sequential file streaming.

[‡]applies to both single spikes and the alignment precision of multiple spike epochs.

It is also possible to compute the interpolated sample points in the frequency domain using suitable phase rotations. But since the raw data is stored as time series, converting the data into the frequency domain, upsampling it, and converting it back to the time domain for further processing was slower than

direct convolution in the time domain, even when fast Fourier transforms (Press et al. 1994) were used.

Supplementary digital filters can be incorporated into the interpolation process, just as the Hamming window was convolved with the IIR filter. As convolution is associative, repetitious convolution of the time series for each filter pass can be avoided by convolving the actual sinc kernel with the filter or filters of interest. Kernels for a multitude of purposes, including high-pass, low-pass, and notch-filters, kernels for smoothing or removal of DC signals, differentiation, integration and so forth (Smith 1997), can therefore be applied without additional computational overhead.

Meijering (2002) gives a comprehensive review of the many available methods of interpolation, but only Fourier-based techniques such as bandlimited interpolation have a solid theoretical grounding. Alternative methods, such as cosine, Gaussian, polynomial and cubic spline interpolation are useful for display purposes but do not provide *bona fide* reconstructions (Schafer and Rabiner 1973). Cubic spline interpolation can, however, be used to good effect aligning spike waveforms by locating the true waveform peak and valley times (Wheeler and Smith 1987). It does so efficiently as only a few samples around the raw peak need be interpolated, and in turn PCA-based spike sorting can be made more effective, as already discussed. Aside from this specific purpose, cubic splines (based upon smooth, continuous first and second derivatives) are not useful for quantitatively accurate waveform reconstruction (Figure 2.1). Moreover, it is computationally prohibitive to process continuously acquired waveforms with cubic splines, making it unsuitable for improving spike detection reliability. In any case there is no reason for not using the theoretically correct Fourier methods of interpolation given that standard desktop computers are capable of processing more than 54-channels faster than real time, even to rates as high as 100kHz (Table 2.1).

2.5.1 Avoiding aliasing

There are few compelling reasons to heavily oversample analog signals during acquisition, as bandlimited interpolation is clearly an effective method of upsampling offline. However, before making the transition to lower acquisition rates the potential for signal aliasing must be considered, because the transition

bandpass of analog anti-alias filters can be fairly broad (Figure 2.8). One way to relate the relevant factors of filter roll-off, the signal (and noise) energy in the transition bandpass, and the level of acceptable distortion from aliasing given the overall dynamic range of the system, is the sampling ratio:

$$S.R. = \frac{s}{f_c} = 1 + \frac{f_r}{f_c} \quad (2.10)$$

where s is the sampling rate,

f_r is the frequency at the required attenuation, and

f_c is the -3dB corner frequency of the anti-alias filter.

Therefore, to maintain the precision of a 12 bit (72dB, or ~0.02%) acquisition system, with f_c set at 6kHz and 5 pole RC anti-alias filters (FA I-64 specifications state 100dB/decade roll-off), the desired f_r is 8.5kHz. This gives a SR of 2.4 and equates to a sampling rate of 14.5kHz, just above the Nyquist frequency based on the assumption of a 'brick-wall' filter with f_r equal to f_c (Figure 2.8). However, the SR is on one hand a conservative indicator, as it assumes a white noise spectrum (*i.e.* above and beyond the transition band), zero tolerance to aliased signals, and ignores other sources of noise that conspire to reduce the dynamic range, or effective number of bits (ENOB), of the system as a whole. A more realistic estimate might put the ENOB of the current system at 10 bits (60dB, or ~0.1% precision) given electronic noise, total harmonic distortion, and potential for crosstalk in the analog to digital circuitry during high speed signal multiplexing. At the same time, the measured filter decay was closer to 60dB/decade (Figure B.1A), somewhat less steep than the 100dB specified. Finally, the signal and noise power spectra were not white. Energy at the corner frequency, of background neural hash recorded *in vivo* (with all apparatus connected and switched on, but excluding spikes), was ~0.02% full-scale (Figure B.2A), matching the f_r . Regular spikes (RS) on average had an even lower spectral power at 6kHz, however FS neurons had significant energy well beyond 6kHz (Figure B.2B). Taken together, a more realistic SR for this acquisition system would be 3.4, avoiding the possibility of aliasing spikes from FS neurons by sampling at 20kHz. If anti-alias filters with f_c lowered to 5kHz and a steeper roll off were substituted into the amplifiers, or if up to 10 μ V_{pp} aliasing was deemed acceptable, then the adjusted SR drops back to ~2.4, a sampling rate of ~12kHz. Of course the opportunity to record from FS

neurons and distinguish them from RS neurons would be significantly compromised by filtering and sampling at this lower rate.

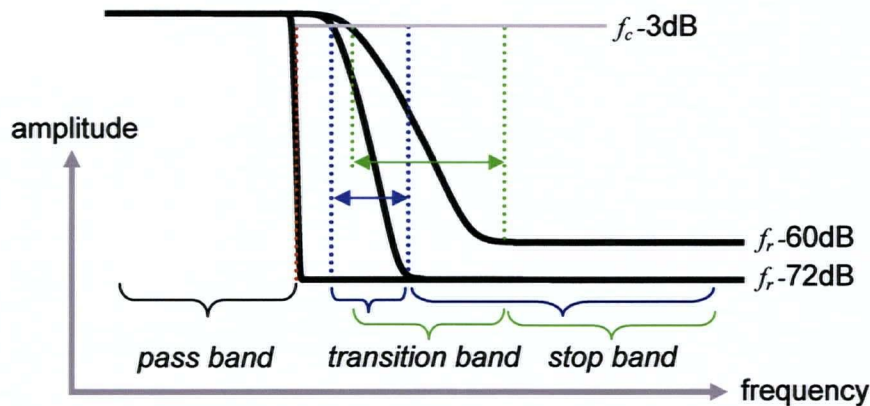


Figure 2.8 Transition bandwidth, sampling ratio and aliasing.

Calculation of an appropriate sampling ratio depends on the corner frequency f_c , the required level of attenuation f_r , and the slope of the anti-alias low-pass filter. The pass band is defined as the range of frequencies below f_c that, for an ideal filter, pass signals without changing their amplitude. The frequency range between f_c and f_r is called the transition band. Signals in this range are subject to aliasing if the simplest interpretation of the Nyquist sample rule is applied (*ie.* $2 \times f_c$). Frequencies above f_r are referred to as the stop band. For an ideal 'brick wall' filter (*red dotted line*) $f_r = f_c$ so the SR = $1 + 1/1 = 2$, in accord with Nyquist criteria. For narrow transition bands (*blue dotted lines*), the SR will be slightly higher, for broad transition bands (*green dotted lines*) the SR will be higher still.

From a practical standpoint, there are other reasons for acquiring and storing data at slightly more than twice the Nyquist frequency (as defined by f_c). Online monitoring of Nyquist-sampled spikes have poor definition, so moderate oversampling avoids the need to upsample online. Oversampling allows for considerably shorter sinc FIR filter kernels for an equivalent level of interpolation precision, because the number of zero crossings is determined by the width of the transition bandpass. A small percentage increase in the original sampling rate affords a larger percentage saving in computation time because the added bandwidth is a larger percentage of the filter transition bandwidth than it is of the original sampling rate (Smith 2004). For example, given a f_c of 6kHz, the transition band for a sampling rate of 20kHz is about 2.5kHz, while 25kHz provides a 5kHz transition band. A 25% increase in sampling rate halves the work per sample

during upsampling. High precision interpolation is thus possible with filter kernels containing fewer than 20 points.

2.5.2 Noise reduction

Noise is the scourge of multiunit recording. Extracellular spike amplitudes from the same neuron may fluctuate, particularly during bursting (Buzsaki et al. 1996; Fee et al. 1996b), and additional noise that adds to this intrinsic variability compounds the already difficult task of isolating single units or deriving accurate spike metrics. Even large amplitude spikes can be hard to distinguish from each other in the presence of large amplitude instrument noise or neural 'hash' (Figure 2.7A). Small amplitude spikes are most susceptible, often rendered undetectable or impossible to sort. It is therefore unfortunate that the PCA denoising algorithm did not only fail to improve the SNR of small spikes, but had a tendency to abolish them entirely. This negative bias may have arisen because small amplitude spikes are typically from more distant neurons, and the farther neurons are from the polytrode the more distributed their fields are across the polytrode (Figure 4.4). It appears that the inner lockout designed to allow for this was not a complete solution. Another likely failing of the algorithm was that the spiking probability of neuron ensembles are not independent, and weak but consistent correlations amongst the spike trains could produce the net average decrease in spike amplitudes observed. Perhaps the only obvious potential use of the PCA denoising algorithm, at least in its current form, is for extracting a subset of isolatable spikes from large amplitude brain ripple. This would be especially valuable in studies of the spike patterns associated with rapid image processing (Fabre-Thorpe et al. 1998), since the spikes of interest will frequently be embedded in stimulus onset-related burst activity.

For more generalised denoising, recently developed algorithms based on independent component analysis (ICA) show some promise (Iriarte et al. 2003), although they have not yet been tested in the context of multiunit recordings. Wavelet-based algorithms, like those used to analyse single-trial event-related potentials (Quian Quiroga and Garcia 2003) and suppress noise artefacts in microneurographic EMG recordings (Diedrich et al. 2003), may also be useful for denoising multiunit recordings (Zouridakis and Tam 1997; Letelier and Weber

2000) by virtue of the dissimilar power spectra of spikes and background neural noise (Figure B.2).

2.6 Conclusions

The original impetus for upsampling and denoising was to lessen the detrimental effects of avoidable sources of variability in order to obtain more accurate spike waveform measurements and, in turn, improve spike detection and sorting. In this regard, sampling at 10–20kHz (depending on the specifications of the system's amplifiers, filter settings, and signals of interest) followed by bandlimited interpolation to 50kHz is strongly advocated, whereas PCA-based denoising was demonstrated to be of limited value for polytrode recordings.

chapter 3

Spike detection and sorting

3.1 Summary

Spike sorting techniques developed for tetrodes are unsuitable for polytrodes. In this chapter, a new sorting algorithm is described that combines unsupervised template generation with multisite template matching (MTM). The strengths and weaknesses of various detection methods are also explored. Of five methods tested, the most effective for distinguishing spikes from noise used a criterion where the waveform must cross two dynamic thresholds of opposite polarity within $\pm 150\mu\text{s}$. On tests with real and simulated data this method worked as well or better than more sophisticated algorithms. Spikes with a signal to noise ratio of one could be reliably detected. A spatiotemporal 'window-discriminator' was used to assess each spike's location and extent to determine which channels to lock-out and for how long. This minimized the chance of missing near-synchronous spikes on adjacent polytrode sites, and enabled detection of synchronous spikes on more separated sites. For spike sorting, multichannel spike templates were automatically generated using an iterative binary split algorithm followed by extended k -means clustering in the entire 54 channel waveform space. Candidate templates were then truncated and fit across the continuously acquired time series to determine spike times for the whole recording. MTM was accurate for signal to noise ratios as low as one, and resilient to partial spike overlap. Minimal user input was required to confirm the validity of templates and set goodness-of-fit thresholds for each template. Unlike most existing methods, MTM has the advantage of being scalable and computationally feasible for large scale recordings comprising millions of spikes. Although originally conceived for polytrodes, these detection and sorting methods may be generally applicable to other multiunit or wire-bundle electrodes.

3.2 Introduction

The goal of spike sorting is simple: unambiguously identify the number of neurons in a multiunit recording, and assign spikes to identified neurons. Unfortunately the simplicity of the goal often belies the difficulty of the task.

Single-unit electrodes are predisposed to recording localised potential gradients because of the small surface area ($\leq 1\mu\text{m}^2$) of the exposed recording tip (Thompson and Patterson 1973; Humphrey 1979). By positioning the electrode tip sufficiently close to a single neuron such that the recorded signal is dominated by that neuron, a skilled electrophysiologist can isolate the neuron, or 'unit', from other nearby active neurons. Well isolated units do not require additional spike sorting. Multiunit electrodes, including the comparatively large surface area recording sites of polytrodes ($\geq 100\mu\text{m}^2$), are by design receptive to electric fields in a much larger volume of tissue (Buchwald et al. 1973; Kettenmann and Grantyn 1992; Henze et al. 2000). Multiunit recordings therefore comprise a complex mixture of spike waveforms originating from an unknown and potentially large number of neurons superimposed on background noise. Separating these signals requires a spike sorting algorithm, of which a myriad different methods have been proposed (reviewed by Schmidt 1984; more recently by Lewicki 1998).

There are a number of reasons why spike sorting is a non-trivial problem. Foremost amongst them is the similarity of extracellularly recorded spike amplitudes from nearby neurons, especially those equidistant from the recording electrode that have comparable morphologies. In such circumstances template matching (Millecchia and McIntyre 1978; Forster and Handwerker 1990; Bergman and DeLong 1992; Zhang et al. 2004) may be useful for discriminating neurons on the basis of differences in spike shape. Alternatively, stereotrodes (McNaughton et al. 1983) and tetrodes (Wilson and McNaughton 1993) utilise differential spike amplitudes across independent channels in order to identify individual neurons (Gray et al. 1995; Harris et al. 2000). This strategy can also be exploited by polytrodes with closely spaced sites [e.g. Figure 1.9; Csicsvari (2003)].

The spike field potentials of multiple neurons sum linearly in the extracellular space (Wehr et al. 1998). Juxtaposed neurons that fire together cause significant distortion of their respective spike waveforms, often making it impossible to disambiguate the sources, or worse giving rise to spurious correlations (Bar-Gad

et al. 2001). Yet for *in vivo* studies of local connectivity (e.g. Aertsen et al. 1989; Alonso and Martinez 1998; Snider et al. 1998) and spike timing dependent plasticity (Bi and Poo 1998), isolating spikes from neighbouring neurons is essential. A few methods (Atiya 1992; Lewicki 1994; Zouridakis and Tam 1997; Zhang et al. 2004) deal explicitly with the problem of spike overlap, but these algorithms involve iterative decompositions that are computationally prohibitive for single channel recordings, let alone polytrodes. Another class of algorithm exploits ICA to 'unmix' the contributions of all spikes, including those from distant neural activity (Comon 1994; Brown et al. 2001). ICA is more efficient than the other cited methods, but its application is limited to recordings where there are fewer active neurons (spike sources) than electrodes (however, see Takahashi et al. 2003).

A third major confound that affects spike sorting accuracy is the variability of the spike waveforms themselves. Spikes from one neuron are not identical each time the neuron fires, primarily due to variations in the inactivation state of voltage-gated sodium channels, which is most apparent during a burst (Figure 4.6B) or complex spike train (Fee et al. 1996b; Harris et al. 2000). The occurrence of a BPAP (Spruston et al. 1995; Buzsaki et al. 1996; Svoboda et al. 1999; Quirk et al. 2001) also dramatically alters the amplitude and distribution of the extracellular spike (Figure 1.8D). As with the spike overlap problem, few algorithms (Fee et al. 1996a) are able to compensate for this level of intrinsic variability.

Most algorithms dissect the problem of spike sorting in the following manner: spike detection; waveform or feature extraction; and either template matching or cluster analysis in order to obtain the spike times for each neuron.

Effective discrimination of spikes from background noise is the first step, irrespective of the method used. While spike sorting has been the subject of extensive study (Lewicki 1998), methods for improving the reliability of spike detection have received comparatively little attention (Obeid and Wolf 2003). As the first stage common to all spike sorting algorithms, detecting spurious spike events or missing valid spikes is as serious as their misclassification. Given the importance of this step, a number of different methods were evaluated here, including simple thresholding, time-windowed amplitude based filters, thresholding of non-linear energy transformed waveforms, and detection algorithms designed specifically for multisite electrodes.

In the second stage, a minimal set of spike waveform features is either manually selected (typically spike amplitude or width) or preferably is derived automatically using dimension reduction techniques such as PCA (Abeles and Goldstein 1977). Statistical clustering methods (Forgy 1965; Anderberg 1973; Massart and Kaufman 1983) are then used to assign spikes to neurons. If template matching methods are used, the usual practice is to fit the entire spike waveform of 'model' templates to classify the remaining spikes in the time series.

Multiunit recordings made with high-density polytrodes compound the practical difficulties of spike sorting. Most contemporary spike detection and sorting methods are not suitable. Manually defining the cluster boundaries of units recorded with contiguous polytrode arrays with more than a few sites is not practical because of the high dimensionality and sheer volume of data. Existing automated clustering algorithms are computationally intensive and do not scale well either. Therefore the desirable, if not essential, attributes of a polytrode spike sorter are that it be unsupervised, operating with minimal human intervention; scalable, suitable for hundreds of recordings sites and millions of spikes; accurate and sensitive, be able to sort and detect low amplitude spikes; and, ideally, it should be able to resolve synchronous or near-synchronous spikes.

Tetrode spike classification schemes typically rely on differential spike amplitudes, discarding information about spike shape, whereas template matching methods are usually only applied to single-channel multiunit data. The algorithm that is explored in this chapter is one that combines the strengths of both multichannel spatial clustering and template matching into an unsupervised algorithm for multisite template-based spike sorting. The issue of scalability is also addressed. Rather than attempting to cluster all the spike events in a given recording, the proposed method randomly samples a subset of spikes to build a set of representative templates using extended *k*-means clustering (Yang and Shamma 1988; Atiya 1992). After setting a goodness-of-fit threshold for each template, these are individually fit across the relevant channels to determine the spike times for each neuron across the entire file. By splitting the problem into two phases the computational load scales linearly instead of geometrically, making sorting of prolonged recordings with hundreds of neurons and millions of spikes more manageable.

3.3 Methods

The proposed spike sorting procedure is summarised in Figure 3.1. It comprises two main stages: i) a *learning phase*, involving threshold-based spike detection and unsupervised clustering of spike waveforms to generate multisite spike templates, and ii) a *classification phase*, whereby representative templates are fit to the original recording to obtain neuron-specific spike times. The table below contains a description of the parameters used throughout.

Table 3.1 Spike sorting parameters.

Symbol	Meaning
X	set of spike waveform samples (or spike-derived cluster features)
N	total number of samples comprising X
F	number of features (dimensionality of X)
k	cluster counter (the number of spike template classes)
K_i, N_i	sample indices comprising cluster, number of samples in class K_i
x_k	a sample assigned to cluster k
y^k	centroid of cluster k
d	Euclidian distance, similarity metric in the F -dimensional cluster space
D	maximum distortion of any cluster, determines cluster to split
v_0	1 st principal component of a cluster, determines axis for splitting
J, d_{min}	distance between any two centroids, the minimum allowable distance
ℓ_i, n_0	flags whether a cluster is locked, number of samples before a cluster is locked
N_{disc}	number of new samples discarded in current iteration, provides exit criteria
T_i	a final validated, truncated multisite template ready for fitting
N_c, N_s	number of channels in a template, number of samples for that channel (truncated)
$\varepsilon_{min}(t)$	local minima fit residual between template and original time series at time t

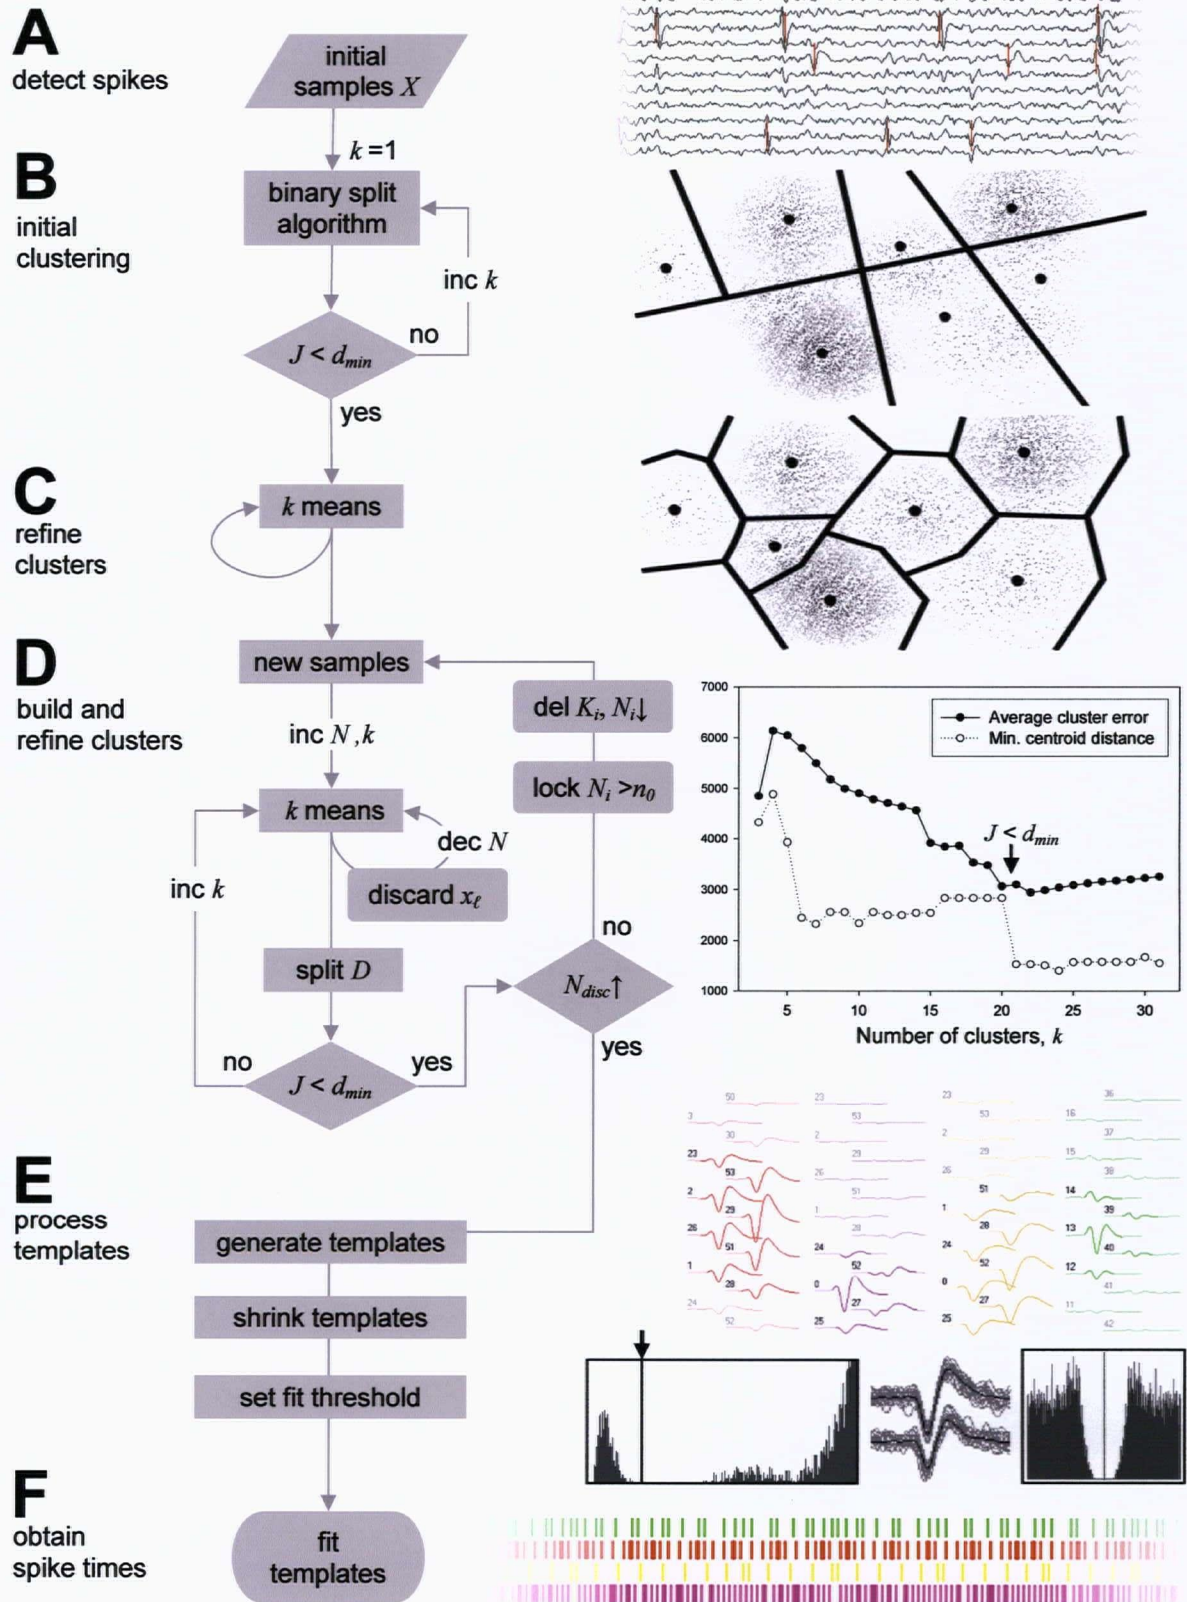


Figure 3.1 Spike sorting flow chart.

3.3.1 Spike detection

A prerequisite of any spike sorting algorithm, irrespective of the method used, is reliable differentiation of spike events from background noise. Even template matching methods, such as the one described in this chapter, require some form of event detection to extract spike samples for representative templates. Five different algorithms were evaluated, four of which are depicted (Figure 3.2).

The first, most common method was a simple bipolar amplitude threshold trigger, essentially a software window discriminator. An event was flagged whenever the voltage s on the i -th channel passed above or below a fixed threshold, f of σ_n , where σ_n was a measure of the background noise (typically four times the standard deviation, but see the following section):

$$f = 4\sigma_n \quad |s_i(t)| > f \quad (3.1)$$

Samples immediately following a trigger were ignored, or 'locked-out', to avoid re-triggering off the same spike (section 3.3.3). A similar but slightly more sophisticated form of spike amplitude thresholding, dubbed a 'multiphasic filter', took into account the bi- and triphasic shape of neuronal action potentials, specifically the fact opposite phases of a spike always follow each other in rapid succession (Figure 3.2B). Accordingly, spikes were triggered only when consecutive threshold crossings of opposing polarity occurred on a given channel s_i within a limited time window, δt :

$$\begin{array}{ll} s_i(t) > f & s_i(t) < -f \\ s_i(t + \delta t) < -f & s_i(t + \delta t) > f \end{array} \begin{array}{l} \text{+ve edge crossing;} \\ \text{-ve edge} \end{array} \quad (3.2)$$

In a sense this algorithm is a crude form of template matching, with a very loose definition of what constitutes a spike waveform.

The third method, the dynamic multiphasic filter, was a slight modification of the multiphasic filter, wherein the second threshold level was relative to the peak or valley of the initial phase trigger, and could precede or follow the initial trigger. The DC signal level was thereby ignored and candidate spikes were only required to have a minimum amplitude inflection occurring within a designated time window. This method is functionally analogous to a minimum slope or first derivative threshold detector convolved with a temporal smoothing window.

Neuronal spikes are characterised by transient high energy, high frequency waveforms compared with background noise. Non linear energy operators (NEO)

estimate the instantaneous frequency and amplitude of a signal (Kaiser 1990; Maragos et al. 1993). By accentuating spike amplitudes more than noise, such transforms have the potential to improve spike discrimination (Figure 3.2C). For continuous signals, the NEO ψ , has been defined by Kaiser (1990) as:

$$\psi[s_i(t)] = [s_i(t)]^2 - s_i(t) \frac{\partial^2 s_i}{\partial t^2} \quad (3.3)$$

For discretely sampled data, spikes were registered when:

$$\psi[s_i(t)] = s_i(t)^2 - s_i(t + \delta)s_i(t - \delta) > f^2, \quad 1 \leq \delta \leq 4 \quad (3.4)$$

where δ was usually equal to 1 (*ie.* 40 μ s), and a suitable threshold f^2 was set manually upon visual inspection of the transformed time series.

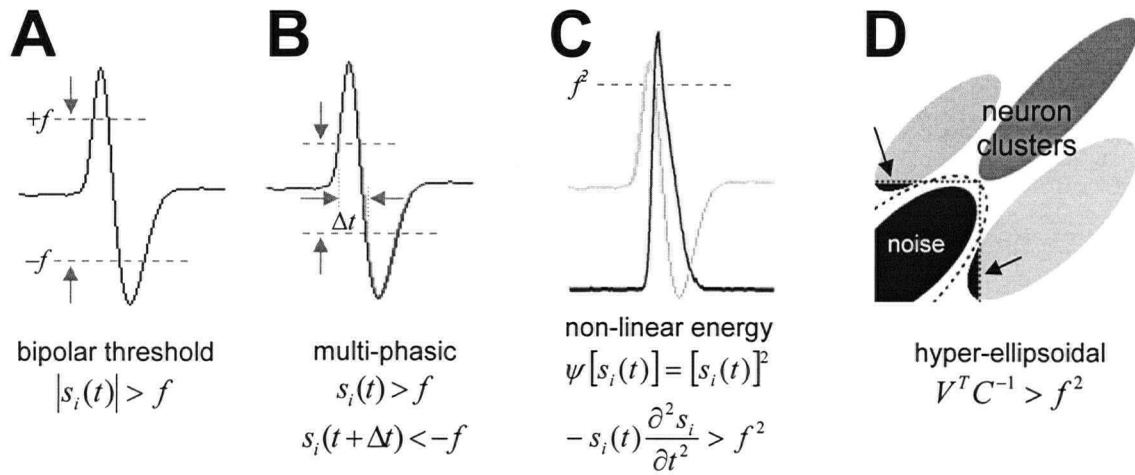


Figure 3.2 Spike detection algorithms.

Depiction of four (of the five) spike detection algorithms. (A) A bipolar threshold. Any waveform on any channel that deviates significantly from noise triggers a spike event. (B) Candidate spike events must cross two thresholds of opposite polarity within a specified time period, Δt . These relations are for positive going spikes; the converse applies for inverted spikes. Lower thresholds are possible with this hysteresis constraint. (C) The non-linear energy transformed waveform (*solid trace*) uses a single positive threshold. (D) A simple threshold on multiple channels corresponds to a hypercuboidal thresholding boundary (*dotted line*), which may not be ideal for delineating spike clusters from the noise cluster (*arrows indicate missed spike events*). The hyper-ellipsoidal algorithm used a multidimensional threshold applied to the signal covariance on nearby sites to partition spikes from the background noise (*dashed line*).

Each of the algorithms described thus far employed independent thresholds applied to individual channels. The hyper-ellipsoidal detection algorithm of

Rebrik and colleagues (1999) works on the assumption that both noise (see Figure 2.7) and spikes are correlated on adjacent recording sites. By computing the cross channel co-variance the hyper-ellipsoidal algorithm attempts to model the noise cluster. Thresholding in this space raises the possibility of more precisely partitioning the spikes from background noise (Figure 3.2D). To determine the hyperellipsoidal thresholding surface the cross-channel covariance matrix, C , was computed from random chunks of data (including spikes). Then, at every sample period, the adjacent multichannel waveforms (virtual tetrodes) were represented as a 4-dimensional vector, V . The spike detection criterion was:

$$\bar{V}^T C^{-1} > f^2 \quad (3.5)$$

As with the NEO method, the threshold factor f^2 was set arbitrarily by hand.

3.3.2 Automated thresholding

Spike detection methods require an appropriate detection threshold. Setting the threshold too high leads to missed spike events, missed templates, or misrepresentative templates. Setting it too low leads to detection of spurious events that cannot be classified, and can slow or even disable some sorting algorithms. An ideal threshold setting is one that detects all clusterable spikes while keeping the number of non-clusterable (noise) events to a minimum.

Various thresholding schemes were assessed, including fixed and dynamic thresholds, global (common to all channels) and independent (with each channel having its own threshold). The latter is justified if noise levels vary significantly from site to site. For practical purposes independent thresholds must be set automatically, especially in the case of high-channel polytrodes. Dynamic thresholds are potentially useful if the level of background noise changes significantly over the course of a recording. Two noise estimates were used to determine the thresholding factor f , regular standard deviation and another based on the median:

$$\sigma_n = median \left\{ \frac{|x|}{0.6745} \right\} \quad (3.6)$$

This measure is reported to be less affected by high firing rates than the standard deviation (Quiroga et al. 2004), providing a better estimate of the background noise irrespective of the level of spiking activity. Independent

thresholds were evaluated on individual channels, and the cross-channel average was used for automated global thresholds. Fixed thresholds were based on arbitrary 10s samples of multiunit data, while running statistics with variable window lengths (5ms to 50ms) were computed for the dynamic thresholds. Whether or not these more elaborate noise measures and thresholds were warranted is examined empirically in the following sections.

3.3.3 Partitioning of spike events

Samples immediately following a trigger must be ignored, or 'locked-out', to avoid multiple triggers off a single spike. Typically a fixed duration lockout of ~1ms applied to all channels has been used in the case of closely-spaced multichannel electrodes (*e.g.* tetrodes) where spikes appear on multiple channels. While this approach may be adequate for multiple independent single unit or tetrode recordings, it is unsuitable for large contiguous polytrode arrays. If following a spike event the entire polytrode is locked-out for a millisecond, then spikes on other channels will invariably be missed (Figure 3.3A), which is especially problematic for studies of neuronal synchrony or local circuit interactions. A partial solution to this problem is a 'spatiotemporal channel lockout' that takes into consideration knowledge about typical spike widths, and the spatial extent of extracellular spike potentials (Figure 1.8), to restrict both the size (number of channels) and duration (number of sample points) of the lockout.

Since the site configuration of polytrodes was known, following a trigger event the algorithm determined the spike's approximate location (the channel with the maximum amplitude, i_{max}) and field size (channels over-threshold within a user-defined radius of i_{max} , typically 150-175 μ m) to decide which channels to consider for lockout (Figure 3.3B). In the temporal domain, only a positive $\delta V/\delta t$ triggered a spike event for positive threshold crossings (vice versa for negative-going spikes), so it was sufficient to lock out only the first half-phase of the spike (Figure 3.3C). To manage the myriad possible spike shapes, bi-, and tri-phasic inflections were considered part of the same spike if successive triggers of opposing polarity occurred within 200 μ s. Every channel in the polytrode array maintained independent bipolar trigger lockouts ℓ_+ and ℓ_- according to these criteria, so a given channel within the spatially-defined domain of a spike event

was locked out for the minimum number of samples necessary to avoid a double trigger (or not at all). In this way the chance of missing simultaneous spikes on disparate electrode sites or near-simultaneous spikes on adjacent electrode sites was minimized. Suitable parameter values for the spatiotemporal lockout were determined empirically from a large representative sample of spike waveforms.

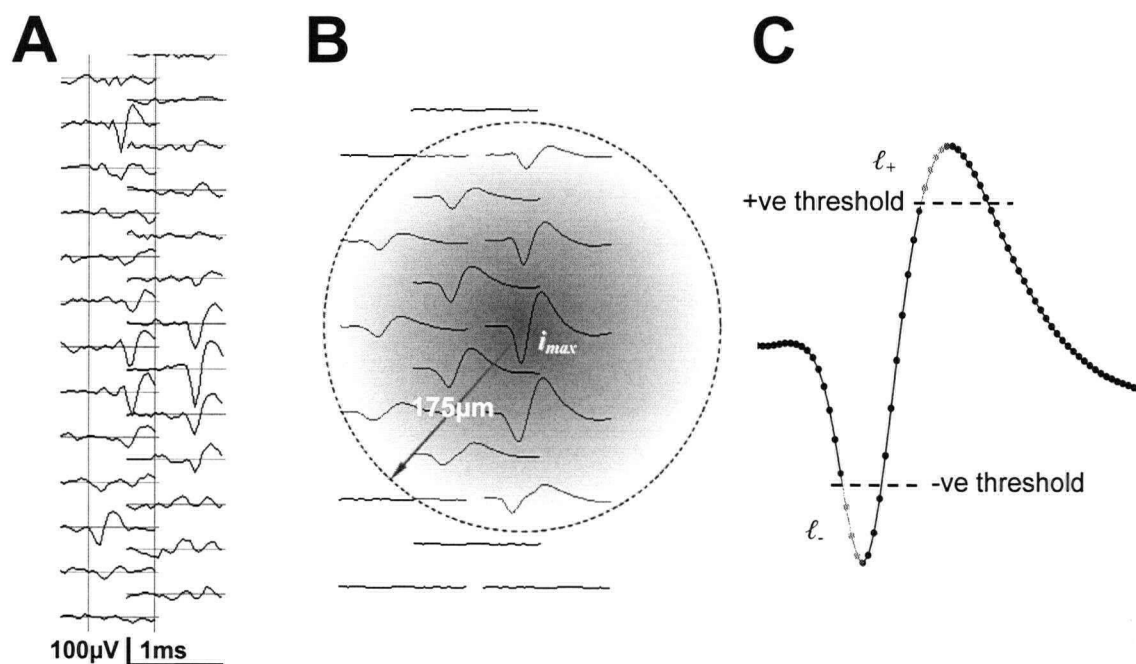


Figure 3.3 Spatiotemporal partitioning of spike events.

(A) Three spikes in a 1ms epoch from 26 (of 54) recording sites (54µmap2b) illustrates the problem of 'locking-out' the event trigger of the entire polytrode following a spike event. In this example, a global 1ms lock-out triggered by the lower spike event would inhibit detection of the subsequent two spike events. (B) A viable solution is to define a lock-out region encompassing only those sites (with signals over threshold) within 150µm of the largest spike amplitude, as shown here for a 54µmap1a recording. (C) The temporal duration of the lock-out (*greyed sample points*) can be reduced by considering the sign of the waveform derivative, and having independent bipolar trigger windows for each site.

3.3.4 Unsupervised template generation

The learning phase of the sorting algorithm began by detecting a couple of thousand spike events (Figure 3.1A), sampled at random from the whole recording. Spike waveform epochs of 1.0ms duration (250µs pre-trigger, 750µs

post-trigger) were extracted from all channels and stored in a large indexed array. These N initial samples $X = \{x_1..x_N\}$ constituted a single mixed 'supercluster' in a F -dimensional feature space, which for a 54 channel polytrode with 1ms epoch samples and 100 points per ms, $F = 54 \times 100 = 5,400$. Hereafter the terms 'cluster' and 'class' are used interchangeably to refer to the set of spike samples currently assigned to a given candidate template.

The binary split algorithm (Anderberg 1973; Sneath and Sokal 1973) is a simple but rapid form of divisive hierarchical clustering (Figure 3.1B). It was used to coarsely partition X into k initial clusters K . As the actual number of spike classes k was not known *a priori*, the algorithm split X into two clusters, and continued to split the resultant clusters in half until k intermediate clusters were created. At each iteration the cluster with the largest distortion D was split, until a valid cluster was split erroneously (as defined below). At this point the algorithm stopped, and the previous $k-1$ cluster assignments were retained.

The algorithm proceeded as follows: (i) the cluster counter k was set to 1. Initial spike samples $[x_1...x_N]$ were assigned to a single supercluster K and its centroid y_k was computed:

$$y_k = \frac{1}{N_k} \sum_{i=1}^{N_k} x_i, x_i \in K_k \quad (3.7)$$

Since the feature space was the entire multisite spike waveform, the centroid represents the emerging average spike template.

(ii) the cluster with the largest distortion D , a measure of the average distance of all points in a cluster to the corresponding cluster centroid, was calculated:

$$D = \arg \max \left\{ d_j = \frac{1}{N_j} \sum_{i=1}^{N_j} \|x_{ik} - y_j\|, j = 1...k \right\} \quad (3.8)$$

(iii) this cluster was then split into two sub-clusters along the hyperplane of greatest covariance by computing the principal eigenvector v_0 of the cluster:

$$\sigma_j = \begin{pmatrix} \sigma_{j,1} \\ \sigma_{j,2} \\ \vdots \\ \sigma_{j,f} \end{pmatrix} \arg \max \{ \sigma_{j,f} \} = v_0 \quad (3.9)$$

Those samples closest to $y_j + v_0$ were assigned to the existing subcluster K_k , and those samples closest to $y_j - v_0$ were assigned to a new subcluster K_{k+1} .

(iv) the centroids of the two new subclusters were calculated (equation 3.7).

(v) the cluster counter k was incremented and steps (ii) – (v) were repeated until a cluster was ‘oversplit’, as indicated by centroids that were too similar:

$$J = d(y_k, y_{k-1}) = \sum_{i=1}^f \|y_{ki} - y_{(k-1)i}\|, \text{ stop if } < d_{\min} \quad (3.10)$$

k -means (Forgy 1965; Lloyd 1982), a ‘nearest neighbour’ clustering algorithm, was then used to refine the cluster boundaries (Figure 3.1C), as follows:

(i) the k -means algorithm was initialised with the spike samples X , the clusters K , and set of k cluster centroids $[y_1 \dots y_k]$ defined by the binary split algorithm.

(ii) each sample $[x_1 \dots x_N]$ was assigned to its nearest cluster centroid $[y_1 \dots y_k]$ according to a Euclidean distance metric:

$$d(x_n, y_k) = \sum_{i=1}^f \|x_{ni} - y_{ki}\| \quad (3.11)$$

where a sample x_n belonged to cluster K_i if y_i was its closest centroid:

$$K_i = \{x_n \mid d(x_n, y_i) \leq d(x_n, y_j), j = 1 \dots k\} \quad (3.12)$$

(iii) the centroids of all clusters were recomputed (equation 3.7).

(iv) if any samples changed assignment, the algorithm returned to step (ii).

The standard k -means algorithm described above was extended to include provision for splitting mixed clusters (Ball and Hall 1965; Atiya 1992), whereby k -means was run repeatedly with additional classes (increasing k), at each iteration computing:

(v) J , the minimum distance between all pairs of centroids:

$$J = \arg \min_{i,j,i \neq j} \{d(y_i, y_j)\} \quad i, j = 1 \dots k \quad (3.13)$$

(vi) if $J < d_{\min}$ then, as with the binary split algorithm, $k-1$ became the final partitioning of the clusters (for the current X), and the algorithm exited.

(vii) empty classes were deleted. Assignment indices for the current clusters were saved. The cluster with the largest distortion (equation 3.8) was split, and samples were allocated to the resultant sub-clusters (equation 3.9). The cluster counter k was incremented by one and steps (ii) – (vii) were repeated until the exit criteria were met (equation 3.10).

An outer loop (Figure 3.1D) continued to extract new spike samples until no additional unique samples were found. At each iteration classes with sufficient samples were locked and excluded from further clustering. Outlier, or ‘dormant’ classes (defined below) were deleted. In this way spike templates representative

of the entire recording were built up in a piecewise fashion, without overloading the k -means algorithm. Hence, after converging at step (vi) for the first time,

(viii) new spike samples were added to a new supercluster K_{k+1} , and N and k were updated accordingly.

(ix) extended k -means, steps (ii) – (vii), was used to allocate the new samples to existing (or novel) classes. Those assigned to locked (completed) classes were discarded. A tally was kept of the samples discarded in this iteration, N_{disc} .

(x) any class with sufficient samples was locked:

$$\text{for } N_i > n_0, \ell_i = \text{true}, \quad i = 1 \dots k \quad (3.14)$$

(xi) dormant classes, those with 3 or fewer samples that had not acquired new samples since the previous iteration, were typically noise events and thus deleted.

(xii) cluster indices were saved. Steps (viii) through (xii) were repeated until the relative majority of new spike samples were discarded:

$$\Delta N_{disc} = \frac{N_{disc(prev)} - N_{disc(curr)}}{N_{disc(prev)}} < 10^{-4} \quad (3.15)$$

For the algorithm to run efficiently, implementation details were important. For example, only the cluster indices, not the sample waveforms, were moved in memory during cluster reassignment. Efficiency issues and the choice of robust exit criteria are considered further in the discussion.

3.3.5 Spike classification: multisite template matching

Having obtained a representative set of multichannel spike templates, these were displayed for user validation (Figure 3.1E). Clusters could be manually split, combined, or deleted if desired, with enlarged spike waveform plots and inter-spike interval (ISI) histograms to aid verification.

Each template was truncated to include only those channels with amplitudes above $15\mu V_{pp}$, and/or those channels within a pre-defined radius of i_{max} . The preferable method of truncation was subjective, as depending on the shape of juxtaposed templates flat channels could have as much discriminatory value as large-amplitude channels, however being overly inclusive ‘diluted’ the signal energy of small, localised spikes, rendering the template useless for fitting.

The set of truncated templates T was individually fit across a short ($< 1\text{min}$) arbitrary segment of the original recording in order to build a ‘goodness-of-fit’

histogram. The main difference between the fitting undertaken here and regular template matching was the multichannel composition of the templates. Each template was 'slid' across the relevant channels at the resolution of the interpolated data (10 μ s intervals). At each interval the least squares fit error ε was calculated. To economise on memory usage only the local minima of the residuals, ε_{\min} , and the corresponding timestamps were stored:

$$\varepsilon_{\min}(t) = \sum_{i=1}^{N_c} \sum_{j=1}^{N_s} (T_i(j) - s_i(t+j))^2 \ni \varepsilon_{t-1} < \varepsilon_t < \varepsilon_{t+1} \quad (3.16)$$

where N_c was the number of channels in the template, and N_s was the number of samples for that channel. Goodness of fit histograms, one for each template, were populated with the fit residuals. These histograms usually comprised two or more distributions (Figure 3.1E), one corresponding to valid matches (to the left of the histogram, with the lowest residuals), and other peaks representing invalid matches (the fit residuals of the current template to similar spikes and noise). A suitable goodness-of-fit threshold for the whole recording could now be set for each template, either manually or by fitting multiple Gaussian functions centred on the predominant distributions. An advantage of sorting spikes in this manner was that a statistically defined detection reliability could be set, which was particularly useful when there was not a clear separation between the valid-match distribution and that of the non-match distributions. In such circumstances a compromise could be established. For example, setting the fit threshold to be two standard deviations above the mean of the valid fit distribution would account for 95% of the spikes in the file being detected, while the overlap with the tails of the other distributions would indicate the incidence of false positive matches.

Finally, the validated, truncated templates were fit sequentially across the entire time series to obtain the sorted spike times for each neuron in the recording (Figure 3.1F). Spike times, along with their associated spike template, were saved to file for subsequent physiological analyses.

3.3.6 Evaluating spike detection and sorting reliability

The literature is awash with different spike sorting algorithms, each purporting to be the most effective for isolating spikes in multiunit recordings. In many cases, however, algorithms have been tested with unrealistic simulated data,

or on real data under ideal circumstances (high SNR, stationarity of spike shapes, and no overlapping spikes). It is therefore difficult to judge the validity of these claims. In the absence of independent validation from intracellular recording (Wehr et al. 1998; Harris et al. 2000), real and simulated data (upsampled to 100kHz) were used to assess the performance of the various spike detection and sorting algorithms presented here. Simulated data were compiled from 1s segments of *in vivo* noise (average $\sigma = 7\mu\text{V}$), replete with distant neural hash but devoid of any local, clusterable spikes. These noise fragments were concatenated into longer 8-10s segments and added to the background noise at pseudorandom, non-overlapping intervals with twelve varieties of multichannel spike templates representative of the spike shapes and distributions observed in cat visual cortex (Figure 3.4A, B). Using actual rather than simulated noise meant the inherent variability on different recording sites, structured or correlated noise across the polytrode, non-stationarities, and faithful noise spectra (Figure B.2) were retained.

Since the exact number of spike events was specified in the simulated spike trains, an explicit comparison could be made of the different detection algorithms in terms of their accuracy and resistance to noise by tabulating the number of correctly detected and false positive spikes. Spike amplitudes ranging from 49–98 μV_{pp} , corresponding to a SNR of ~1 to 2:1, were used for these tests. A second simulated file, seeded with unmodified spikes, was used to evaluate the template generation stage of the spike sorting algorithm. In a third simulation, the spike sorter was provided with spike templates ranging from 35–84 μV_{pp} to assess the reliability of MTM under conditions of low SNR. Finally, the sorting algorithm's tolerance to spike overlap (Figure 3.4C) was gauged by systematically varying the spatiotemporal overlap of various multisite spike templates having different shapes and relative amplitudes.

All tests of the spike sorting algorithm were 'blind' in that the simulated files were generated by one person, and the number, type and variety of spikes were unknown to the person doing the sorting (*i.e.* choosing the trigger thresholds, validating the templates, setting of fit thresholds). Real recordings, of course, provide the most realistic test of sorting performance, as they contain additional sources of variability not found in the simulated files, such as non-stationarities of spike amplitude and shape (Fee et al. 1996b; Harris et al. 2000; Quirk et al. 2001).

However, as mentioned previously, such tests are useless for quantitative validation as they lack the 'gold-standard' reference of the synthetic recordings.



Figure 3.4 Real and simulated spike trains.

(A) Eight channels of actual recording, compared with (B) the simulated test files, seeded with a realistic variety of spikes of different polarity, amplitude, shapes, cross-channel spreads, and firing patterns, embedded in real background noise. Note the variability of spikes from the same template due to the noise. (C) Spikes from two neurons recorded on the same site illustrate the problem of spike overlap (*arrows*).

3.4 Results

3.4.1 Spike detection

Noise varied considerably across channels of the same recording. This was partly due to differences in amplifier channel noise, but largely because of variations in background neural noise in a given brain region (Buchwald and Grover 1970). In the example shown in Figure 3.5A, the average noise level was $10.8\mu V_{rms}$ on the noisier site, nearly two-thirds greater than that of the quiescent site ($6.7\mu V_{rms}$). Across all sites the noise ranged from $5\text{--}15\mu V_{rms}$, a three-fold variation.

Within a given site the variation over time was less, up to 1.5-fold. These results justify the use of independent, dynamic noise estimates for determining automated spike detection thresholds, especially as multiples of the noise estimate are generally used. For instance, a suitable threshold based on the least noisy sites might be $\pm 20\mu\text{V}$, whereas the suggested threshold for the noisiest sites would be $\pm 60\mu\text{V}$. A global threshold fixed at either of these levels (or in the middle) would be ideal for neither site.

Both the median-based and regular standard deviation noise estimates gave comparable noise measures in the absence of spikes (Figure 3.5A). However, just a single spike increased the standard deviation noise estimate by a factor of two (Figure 3.5B), and multiple spikes produced a sustained three-fold overestimate of the noise. The median-based noise estimate was much less susceptible to the influence of spikes, even during rapid spike bursts (Figure 3.5B). Using a broader time window or replacing the dynamic estimate with a fixed estimate based on an arbitrary sample period did not rectify the problem. Measurements of artificially seeded spike files (Figure 3.5D) showed that for an average aggregate spike rate of 100Hz, the standard deviation overestimated the noise level by a factor of two, while the median estimate remained accurate to within 10% of the true value. In circumstances where it was desirable to increase the noise estimate, for example during periods of high-frequency neural 'hash', the median noise estimate responded appropriately (Figure 3.5C), so there was no obvious drawback of using this alternative measure.

A large sample of neurons ($n=255$, from 3 cats, 8 penetrations, all polytrode designs) was used to establish suitable parameter values for the spatiotemporal partitioning algorithm. For the lowest practical threshold setting of $\pm 20\mu\text{V}$, the average number of locked sites per template was 4.2 ± 2.8 (\pm standard deviation, range 1-13). This corresponded to an average lockout radius of $47 \pm 22\mu\text{m}$ (range 0-156 μm), depending on both the neuron and the polytrode site configuration. Of those sites that were locked the average temporal lockout (from threshold crossing to peak or valley inflexion), was $57 \pm 22\mu\text{s}$ (range 10-160 μs). Therefore the duration of the spatiotemporal lockout for most spikes was at least an order of magnitude shorter than that usually used for window discriminators. These analyses provided a solid basis for limiting the upper bound of the spatial lockout radius to 175 μm , and the width of the multiphasic filter (methods 2 & 3) window to $\pm 150\mu\text{s}$.

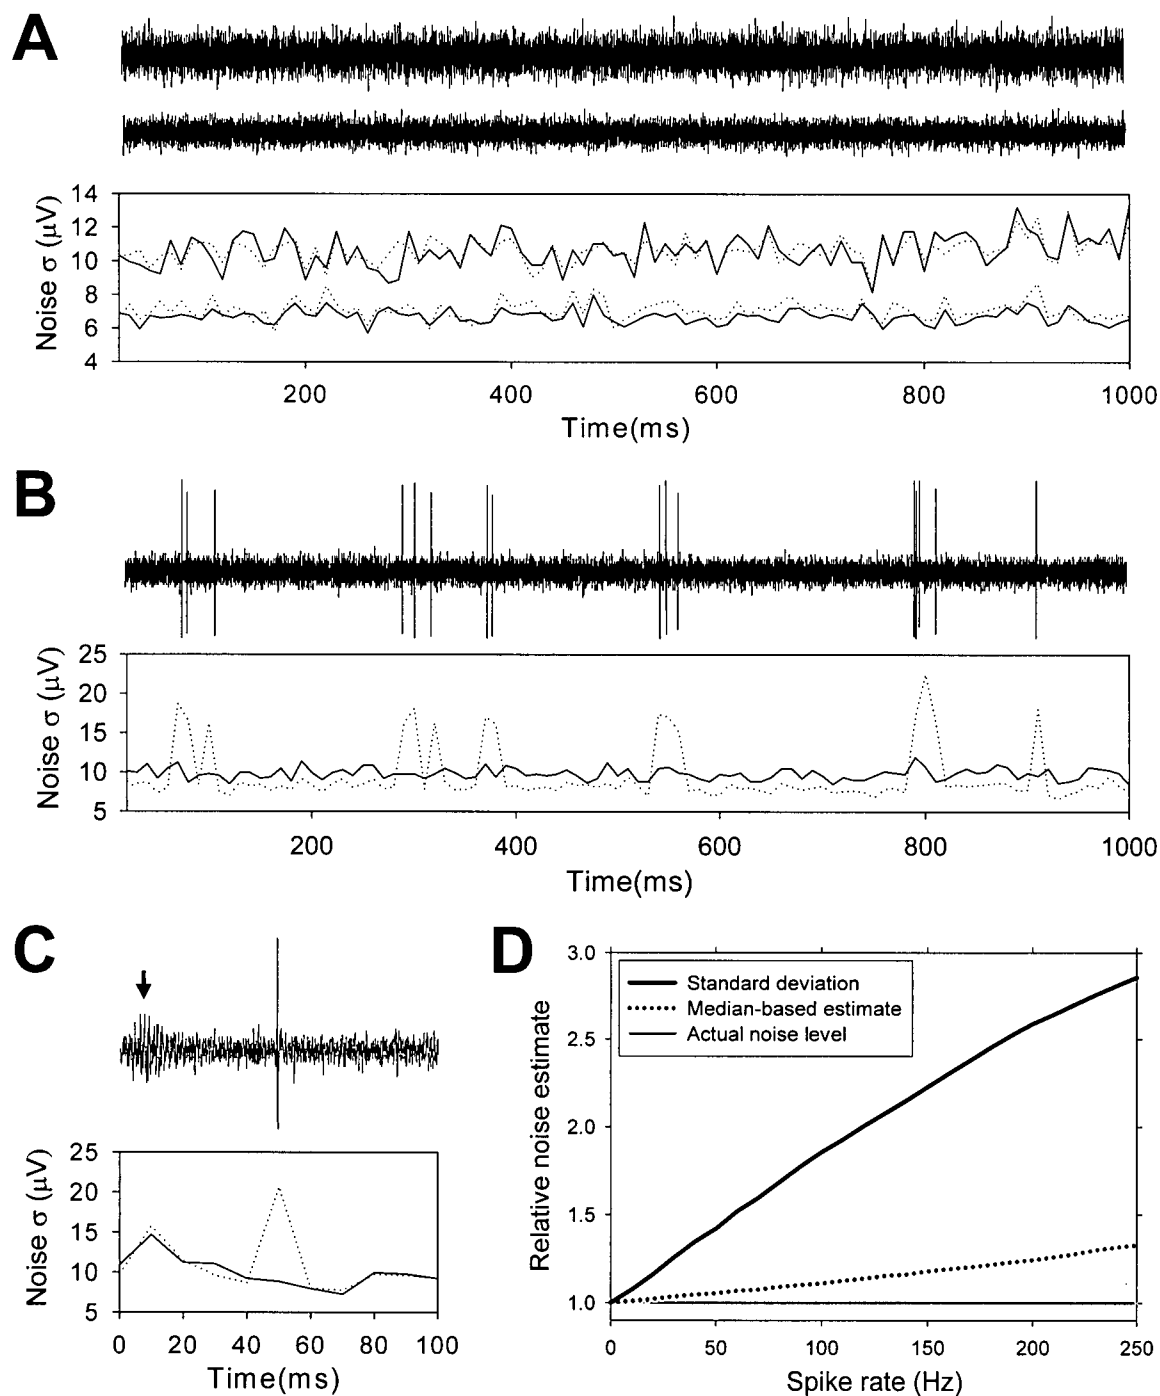


Figure 3.5 Noise estimates for automated thresholding.

(A) Background noise varies across channels and over time. Two channels from the same recording are shown with their respective standard deviation (*solid traces*) and median (*dotted traces*) noise estimates. Note the close correspondence between the two measures. (B) Spikes produce an artefactual elevation in the standard deviation noise estimate. (C) Both measures respond appropriately to the noise envelope associated with high levels of neural hash (*arrow*). (D) The median estimate is accurate even for high spike frequencies. Estimates were calculated on 10ms non-overlapping time windows.

The value of the spatiotemporal lockout was exemplified in highly active recordings (Figure 3.6). As MTM (see following section) gave the most reliable spike detection in high activity regimens, even for partially overlapping spikes, it serves here as a reference for the other detection methods. This 5s multiunit recording segment had an average aggregate spike rate of $\sim 500\text{Hz}$ across the whole polytrode. In such a recording there were inevitably many coincident spikes, resulting in nearly one third of the spike events (those identified by template matching) being missed by the simple window discriminator with a global, fixed duration post-trigger lockout of 1ms (Figure 3.6B). Reducing the lockout to $300\mu\text{s}$ meant 87.8% of the spikes could now be detected. The spatiotemporal lockout performed most reliably, detecting 98.9% of all spikes without false positives or double triggers. Only 28 of the 2522 identified spikes were missed because they were eclipsed by other coincident or near-coincident juxtaposed spikes (Figure 3.6A). Although this recording segment was deliberately selected because of its high level of activity, it is precisely these periods where synchronous oscillatory phenomena (Gray and Singer 1989) and monosynaptic interactions of interest are likely to be observed.

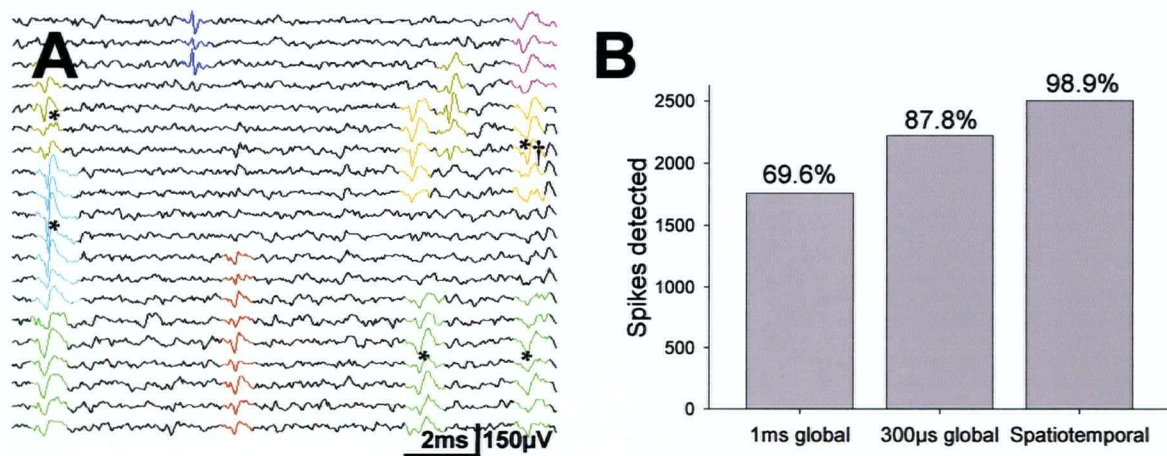


Figure 3.6 Spatiotemporal spike partitioning.

(A) The difficulty of detecting coincident and near-coincident spikes is illustrated in this short recording extract. A subset of channels contained 8 units identified by the multichannel template sorter (*colour coded*). Spikes missed by the 1ms and $300\mu\text{s}$ global lockouts (*) were detected by the spatiotemporal algorithm, with one exception (†). (B) The number of correctly detected spikes for fixed duration global and spatiotemporal lockouts in a highly active 5s multiunit recording segment. Percentages are relative to the number of spikes detected by the MTM algorithm over the same period.

'Accuracy' is defined here as the combined number of missed and false +ve triggers expressed as a percentage of the total number of seeded spikes. Spikes with amplitudes over $90\mu V_{pp}$ (14 times the noise standard deviation, SNR $\sim 2:1$) could be detected with 100% accuracy with any of the spike detection algorithms. Differences in the accuracy of the various detection methods were most apparent for spikes between $50\mu V_{pp}$ and $90\mu V_{pp}$, especially for very small spikes around $56\mu V_{pp}$, or 8 times σ_{noise} (Figure 3.7A). When comparing the performance of the different algorithms, what matters is not the absolute threshold at which 100% of the spikes are detected, but rather the highest level of accuracy attainable at any threshold. Only the dynamic multiphasic trigger (method # 3) was capable of detecting spikes of this amplitude with better than 95% reliability and less than 5% false triggers, albeit over a narrow $2\mu V$ threshold range. This suggests that setting a dynamic, automated threshold to $22\pm 1\mu V$, or 3 times σ_{noise} , would be ideal for detecting small spikes with minimal contamination from noise. The other methods were less accurate at this low SNR. At 95% correct detection, the bipolar, multiphasic, NEO and hyperellipsoidal algorithms gave 11%, 90%, 9%, and 25% false positives, respectively. Conversely, to ensure no more than 5% false positives were detected on average, the detection reliability for valid spikes dropped to 90%, 75%, 89%, and 83%.

For spikes over $65\mu V_{pp}$ all algorithms, including the bipolar threshold, had a threshold that gave at least 95% detection reliability with no more than 5% false positives. What differed was the range of thresholds over which this performance criterion was met (Figure 3.7B). This range is relevant because it is an important practical matter whether or not the optimal threshold setting can be determined easily or automatically (*e.g.* for $65\mu V_{pp}$ spikes the threshold for the hyperellipsoidal and multiphasic algorithms must be within a microvolt of the optimal setting to achieve the specified accuracy). As the amplitude of the spikes increased, the threshold tolerance relaxed. By $90\mu V_{pp}$ the performance of the bipolar, dynamic multiphasic, and hyperellipsoidal detection algorithms were virtually the same. The two exceptions were the multiphasic trigger, that was slow to converge with the other methods because it consistently missed low amplitude monophasic (unipolar) spikes, and the NEO transform, which gave a useful detection range that was superior to all other methods for spikes above $70\mu V_{pp}$.

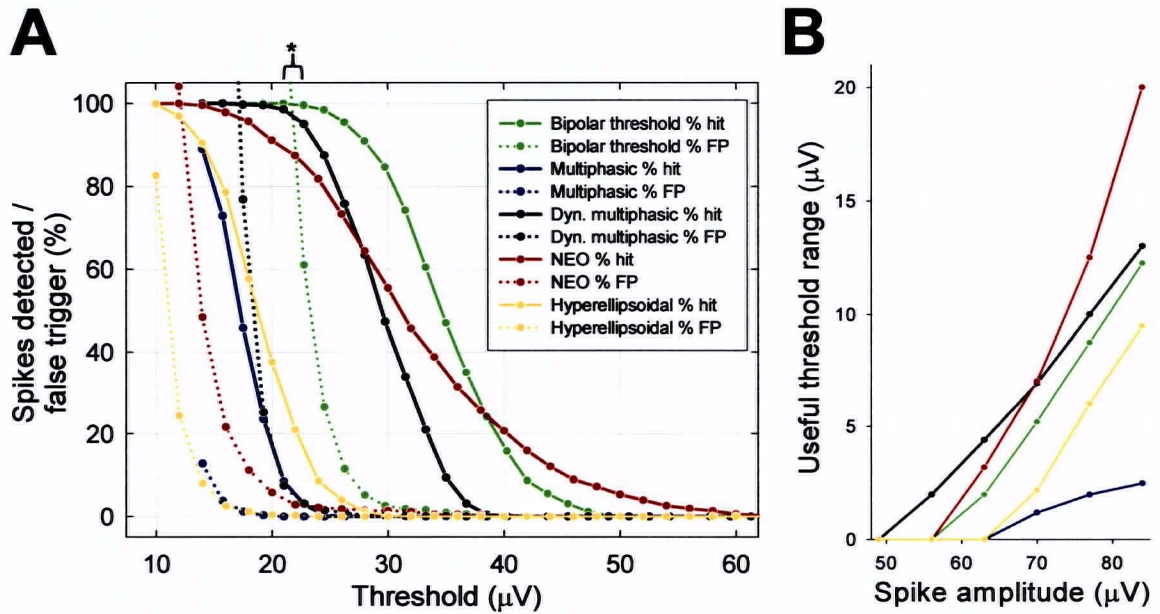


Figure 3.7 Low amplitude spike detection.

(A) Performance of the five different detection methods for small spikes (56μV_{pp}). At a SNR of just 1.1:1, only the dynamic multiphasic algorithm (*black traces*) was capable of detecting spikes with over 95% reliability (correct hits) and less than 5% false positives (FP) over the narrow threshold range shown (*). (B) Valid detection ranges, as defined in (A), for the different methods at different spike amplitudes. Note that the x-axis in (A) and the y-axis in (B) for the hyperellipsoidal and NEO plots are not in units of μV, but were adjusted to allow comparison with the amplitude based methods.

The relative performance of the five detection schemes at different SNRs is summarised in Table 3.2, ranked in order from least to most effective according to accuracy (for optimal thresholds). The detection reliability of MTM is included for comparison (see Figure 3.9). As above, the dynamic multiphasic filter was the most effective at all SNRs, only slightly less accurate than MTM for the lowest SNR, and second only to the simple bipolar threshold in terms of computational simplicity. These analyses also indicated optimal threshold settings, relative to the noise level, for attaining the highest possible accuracy regardless of the SNR. For the amplitude-based detection methods – the bipolar, multiphasic, and dynamic multiphasic – the ideal thresholds were 4, 2.5, and 3.5 times the noise estimate (σ_n), respectively. No consistent optima were found for the hyperellipsoidal or NEO transforms. Rather, the optimal threshold varied depending on the spike, its distribution, and the SNR, making automation of these methods difficult.

Table 3.2 Comparative performance of different spike detection schemes.

Detection method	SNR 1.0 $P_m \mid P_f$	1.1 $P_m \mid P_f$	1.3 $P_m \mid P_f$	1.7 $P_m \mid P_f$	2.0 $P_m \mid P_f$	Computational complexity *
hyperellipsoidal	.243 .155	.155 .056	.067 .050	.002 .012	.000 .000	35M, 33A, 10S, 1C
multiphasic	.117 .230	.095 .081	.037 .076	.002 .005	.000 .000	$2 \leq C \leq 15^\dagger$
non-linear energy	.162 .120	.088 .059	.036 .027	.001 .002	.000 .000	2M, 1S, 1C
bipolar threshold	.125 .165	.090 .053	.043 .025	.002 .002	.000 .000	2C
dyn. multiphasic	.093 .078	.050 .032	.026 .015	.001 .002	.000 .000	$2 \leq C \leq 15^\dagger$
template matching	.065 .070	.032 .035	.015 .000	.000 .000	.000 .000	$N_s M, (N_s - 1)A, 1C$

The probability of false negatives (P_m), and false positives (P_f) is indicated for the 'optimal' threshold at each SNR. The optimal threshold was determined by maximising the accuracy (*ie.* $\text{argmax} [P_m - P_f]$). * C, M, S, and A denote the number of comparison, multiplication, subtraction, and addition operations per sample period, per site, respectively. † the actual number of comparisons depends on the spike width and sample rate. N_s is the multisite template length. The complexity shown for the hyperellipsoidal method is for $N_c = 4$, as described by Rebrik et al. (1999).

3.4.2 Spike sorting

When applied to real recordings, unsupervised extended k -means clustering of spike waveforms produced distinct, pure multisite spike templates (Figure 3.8). The few hundred neurons that were studied in this thesis were all obtained using this algorithm (for other examples, see Figures 1.9, 1.14, 3.9, and 4.7).

Simulations with synthetic data provided valuable information about the performance limits of MTM, in addition to indicating areas for improving future implementations. The results of the blind template generation were as follows: 38 of the 45 unique spike templates (84%) were identified correctly. Of the spike classes that were missed, three had amplitudes smaller than the trigger thresholds specified by the user, two were rejected by the sorting algorithm as outliers ($N_i < 3$, cluster size not increasing), and two were never detected because only ~50% of the file was sampled, and these spikes were only seeded once each. These results highlight the importance of effective event detection, and making sure the whole

recording is adequately sampled when making candidate templates. Of the templates that were correctly identified, and subsequently fit across the entire file, 27 (73%) identified spike times with perfect accuracy, 7 templates (19%) missed fewer than 1% of the seeded spikes, two missed less than 5% of the spikes (1 of 43, and 1 of 56 seeded), and the template with the poorest sorting accuracy still managed to classify over 90% of its spike events correctly. The latter template missed 8% of the spikes because it was barely distinguishable from noise (SNR 0.9), and an overly conservative fit threshold was chosen (manually). Overall, for those spike classes that had templates, less than 0.1% of the total number of spike events in test file were missed. No false positive events, due to either misclassification or noise, were identified with any of the templates.

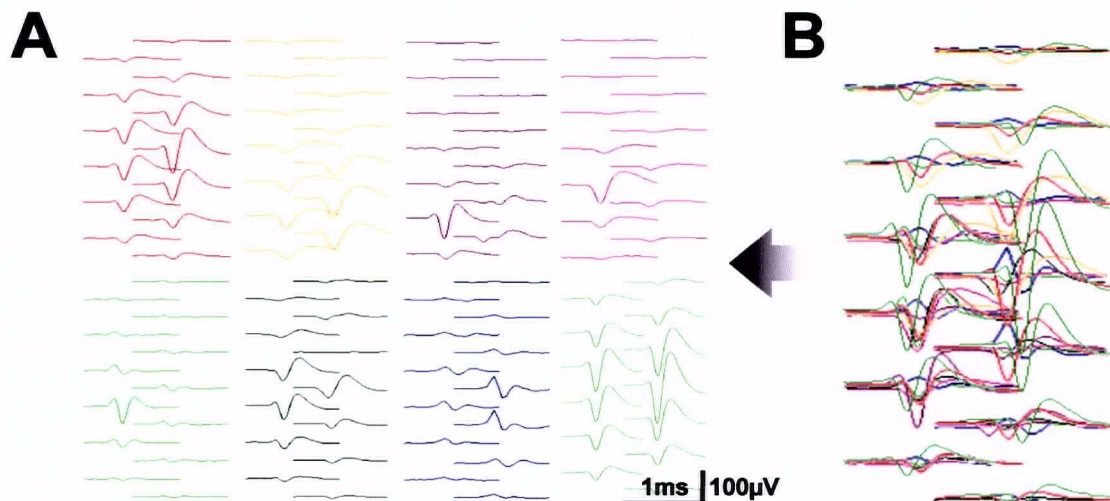


Figure 3.8 Unsupervised template generation.

(A) Examples, in a local group of channels, of multisite spike templates automatically extracted by the waveform clustering algorithm (data from a 54μmap2a polytrode). (B) The superimposed templates give an impression of the unmixing problem the clustering algorithm must solve.

Analysis of sorting performance using attenuated templates (Figure 3.9A) demonstrated that MTM was capable of perfect accuracy (zero missed spikes, zero false positives) for spikes with amplitudes over $60\mu V_{pp}$, a SNR of just 1.2. The data from four very different spike templates (two broad, two focal, biphasic and triphasic waveforms) were pooled because they gave very similar results, however one of the broad triphasic templates could detect spikes with 95% reliability and

no false positives at $42\mu\text{V}_{pp}$, which was considered impressive for spikes embedded in noise of equal or greater amplitude. Together (fitted curves in Figure 3.9A) there was an inverted symmetry between the exponential drop in the number of correct spike detections and the corresponding rise in false positive triggers at low SNR. The lowest spike amplitudes that could be consistently classified with 95% accuracy were around $50\mu\text{V}_{pp}$. If this result is compared with the best spike detection algorithm (dynamic multiphasic), extrapolating the black line in Figure 3.7B to the 'optimal' trigger threshold, it is clear that MTM performs as well or better than the best threshold-based method of spike detection (see also Table 3.2). From a practical standpoint, the dynamic multiphasic spike detector is well-suited for building candidate spike templates that can subsequently be used to detect and sort spikes with even higher accuracy using template matching.

MTM was in some circumstances surprisingly resistant to sorting errors caused by overlapping spikes with distorted spike waveforms. Large spikes were unaffected by smaller overlapping spikes even when completely superimposed (Figure 3.9B). The lesser amplitude spike was rendered undetectable once more than about 10% of its waveform (on either side) was contaminated by the large spike. For spikes of similar amplitude (Figure 3.9C), the template with appreciable signal on only two sites fared worse than the broader spike with signal on five sites. With the latter spike, detection accuracy was maintained above 90% except when the peaks coincided ($\pm 50\mu\text{s}$), presumably because two of the sites were only weakly misshapen by the residual signal of the other spike. These results were for spatially overlapping spikes (a common i_{max}). Displacing the two templates by just one channel ($50\mu\text{m}$ intersite spacing) meant the detection and classification accuracy of either spike was far less severely compromised. As in the previous example the broader spike was least affected (Figure 3.9D), with near perfect accuracy regardless of the temporal superimposition. The more focal spike still maintained a high level of accuracy (above 90%) providing less than ~40% of its waveform was contaminated by the other spike. Note that in each of these examples the drop in accuracy was due not to false positives (noise events) or sorting misclassification, but rather missed spikes (*ie.* loss of valid spike detection). Excessively distorted spikes have a fit residual well beyond that of spikes to noise, so a single goodness-of-fit threshold that can accommodate such spikes unavoidably includes an inordinate number of noise events. The results

presented here were representative of five other neuron pairs with different combinations of spike shapes, relative sizes, and spatiotemporal overlap.

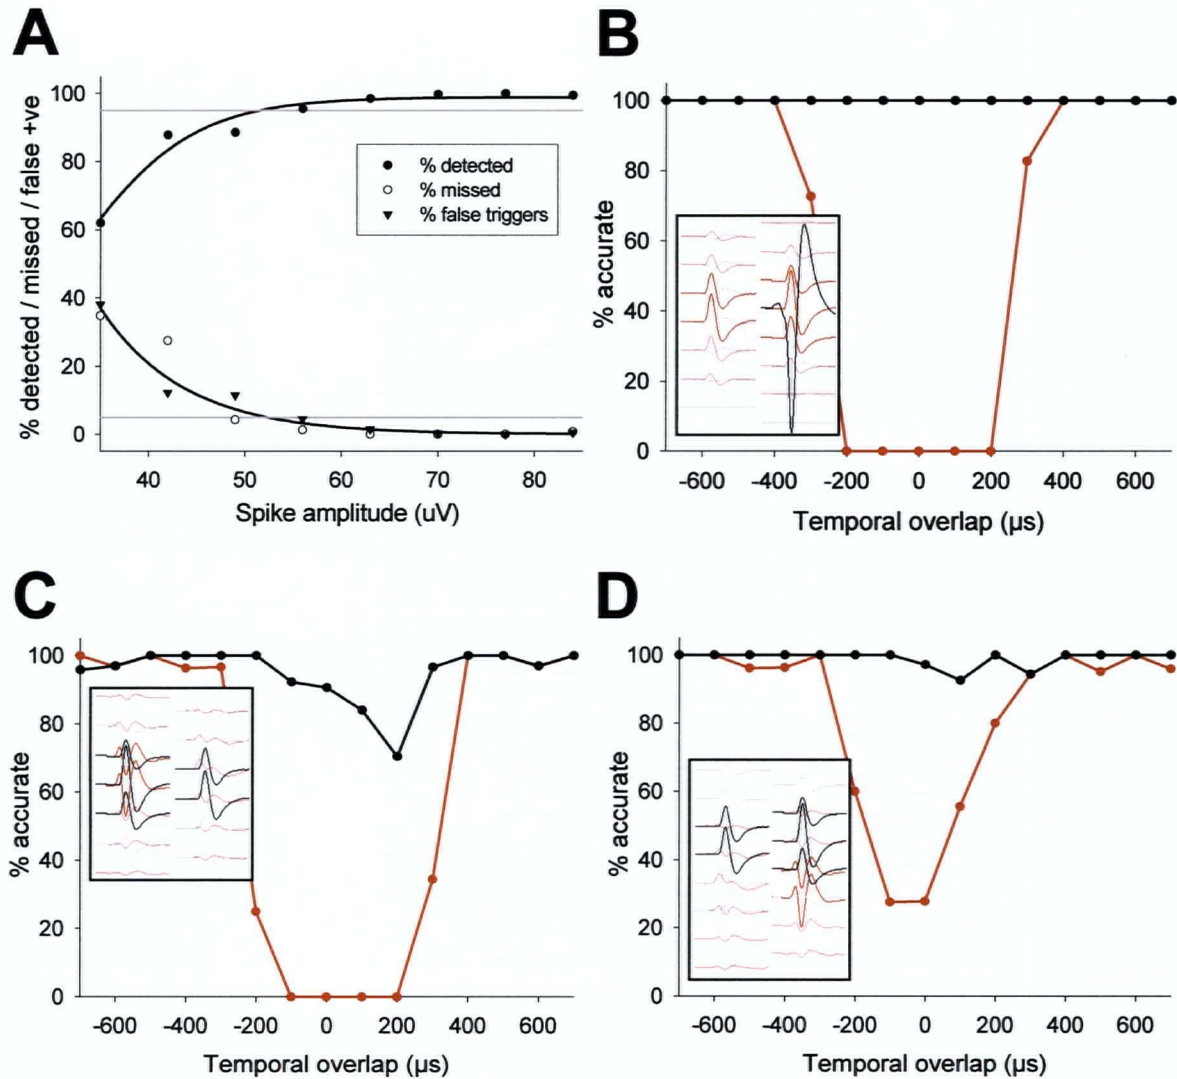


Figure 3.9 Sorting of low SNR and overlapping spikes.

(A) Multisite template matching was perfect for spike amplitudes above $65\mu V_{pp}$ (SNR 1.2). The number of missed spikes and false positive events increase rapidly for SNR below ~ 1 . The *grey lines* indicate the 5th and 95th percentiles. Fitted curves are exponentials ($r^2=0.94$). Sorting performance was also assessed as a function of spatiotemporal spike overlap. (B) Only the small spike was affected by the large spike. In each example the *insets* show the spatial relationship of the two spikes under study. Percent accurate refers to the combined percentage of missed and false +ve spikes. (C) Two spikes of similar amplitude and spread, with i_{max} overlapping and (D) i_{max} offset by one channel. The spike distributed on more sites was less susceptible than the smaller, more focal spike.

3.5 Discussion

Although multiunit spike sorting has been studied extensively, currently available algorithms are not directly applicable to polytrodes either because they do not scale well, making them computationally intractable, or because they rely heavily on user input, which is not practical for large contiguous electrode arrays. The MTM implementation described here was adapted from existing clustering and template matching procedures in order to make them suitable for the large numbers of neurons and spikes routinely acquired with polytrodes.

The simple method for partitioning spike events was capable of detecting 98.9% of all spikes, even during periods of high activity with frequent synchronous events. Median-based, dynamic estimation of background noise provided a basis for automating threshold settings for optimal spike detection, avoiding the arduous task of manually adjusting appropriate thresholds for each channel. When combined with spike detection methods based on measurements of a large library of spike waveforms, very low amplitude spikes (~ 0 dB) could be detected with minimal contamination from noise events. Finally, MTM was effective at sorting high-bandwidth polytrode data with a high degree of accuracy, even for spikes with a SNR less than one and partial spike overlap.

3.5.1 Detection and sorting of low SNR spikes

Poor spike detection can disable clustering algorithms if an excessive number of noise events are included among the valid spike samples. The template generation algorithm described here was no exception, slowing dramatically in the face of growing numbers of outlier classes. As the first stage of all sorting procedures, low SNR spike detection is essential to obtain a representative sample of the neural population, including small amplitude fast spiking cells. Even if spikes are ultimately detected by other means, as is the case with MTM, efficient generation of candidate templates still requires accurate spike detection.

A superficial comparison of the five detection methods (including MTM) suggests that the differences in performance are subtle, but because the incidence of false positives rises precipitously at low detection thresholds, apparently small improvements should not be disregarded. The dynamic multiphasic method

ranked the best, able to reliably detect spikes with unitary SNR, or 15–25% smaller amplitudes than the other methods (Figure 3.7B). It is gratifying that this algorithm is also one of the least sophisticated, fastest methods to compute (Table 3.2), which is an important consideration if any of these methods were to be used for real-time event detection or implemented in VLSI hardware for brain machine interfaces (Obeid and Wolf 2003). The next best detection method was the NEO transform, followed closely by the bipolar threshold, the hyperellipsoidal, and finally the (non-dynamic) multiphasic filter. The hyperellipsoidal algorithm can be disregarded on the basis of weak performance and computational complexity. The multiphasic trigger with fixed counterphase thresholds should also be avoided as it consistently missed monophasic spikes (for examples of such spikes, see Figures 1.8, 1.9, 1.14, 3.8, 4.7). The NEO performed less well than the dynamic multiphasic at low spike amplitudes, but surprisingly out-performed all methods for spike amplitudes above $70\mu V_{pp}$. This finding is consistent with an earlier study claiming superior spike detection with NEO (Kim and Kim 2000), however the authors only provided evidence for spikes with SNR above 1.3:1. Had they compared the detection performance of NEO with other methods, even a simple threshold, at lower SNRs (Figure 3.7, Table 3.2), they presumably would not have made such a sweeping assertion. Nevertheless, it may be possible to improve NEO's performance at lower spike amplitudes by shifting the bandpass of the transform to emphasise the spike frequency energy, and de-emphasise the noise-dominant energy spectrum [*ie.* by changing δ in equation 3.4, and adding a variable-width Bartlett smoothing window (Mukhopadhyay and Ray 1998)].

The performance of the dynamic multiphasic trigger compares favourably with other spike detection schemes, including more sophisticated methods such as the wavelet transform (WT). The discrete WT can be viewed as a filter optimised for extracting spikes by preserving overall waveform shape but filtering out most of the background noise. WTs are purported to be very useful for spike detection (and sorting) (Zouridakis and Tam 1997; Hulata et al. 2000; Letelier and Weber 2000; Hulata et al. 2002; Laubach 2004; Quiroga et al. 2004), however thus far it has only been tested on high SNR data ($\geq 2:1$). The WT is commonly used to construct basis sets for decomposition methods dealing with overlapping spikes, and their application for this purpose will be discussed in the following section. Yang and Shamma (1988) carried out a detailed comparison of several spike detection

methods, including a multi-threshold trigger akin to the multiphasic trigger (but without the temporal window), matched filters (spike templates comprising three points to define waveform shape), and simple thresholding. For a low SNR (1.3:1) simulated spike embedded in white noise, the methods were ranked in the same order reported here, however the matched filter performed less well than MTM in absolute terms (6% vs. zero false positives and 1.4% vs. 1.5% missed spikes), highlighting the value of using the whole spike waveform in template matching.

Before leaving the discussion on spike detection it should be mentioned that, for sake of simplicity, fixed thresholds (with $3.5\mu\text{V}$ test intervals) were used in all simulations. Using dynamic, channel-independent thresholds would have made interpretation of the comparative performance of the detection methods more difficult. In any case the noise was more or less stationary ($7.0 \pm 0.5\mu\text{V}$) and similar across channels (mean range $6.7 - 7.3\mu\text{V}$), so the outcome would likely have been similar. Incorporating dynamic multichannel thresholds (based on the median noise estimate) would, if anything, be expected to improve overall detection performance by optimising the threshold setting at every moment in the time series. Dynamic thresholds compensate for lateral shifts in the accuracy curves as the SNR changes (Figure 3.7A), but should not change the slopes or 95/5 percentiles (differential hit/false positive percentages) unless the spectral content changes significantly with noise amplitude. It is therefore reasonable to conclude that the ranking of the various methods would be the same for dynamic thresholds applied to real data with more variable noise levels.

Comparing the sensitivity and specificity of MTM with existing spike sorting methods is problematic, primarily because important details such as the SNR of the spikes, the similarity of the spike shapes, and the incidence of spike overlap, are often lacking in published works (e.g. Fee et al. 1996a; Hulata et al. 2000). Even when such details are provided it is often impossible to disambiguate the relative contribution of the spike detection stage from the feature extraction or spike sorting stages in terms of the overall performance of the algorithm. Few papers in the literature report the comparative performance of their technique in the context of more established methods for the same test dataset. Two notable exceptions were the study by Wheeler and Heetderks (1982) who found that optimal filtering methods or feature-based clustering using PCA did not classify as accurately as template matching. Atiya's (1992) single channel template

matching algorithm showed accuracy curves indistinguishable from those obtained here for MTM (Figure 3.9A), but only for tests of two dissimilar, non-overlapping spike classes. The accuracy dropped significantly (by about 50%) for three similar neuron classes, which demonstrates the advantage of fitting multisite over single channel templates. In general it suffices to say that none of the published spike sorting methods – including single channel template matching (e.g. Yang and Shamma 1988; Bergman and DeLong 1992), single and multichannel feature-based clustering (e.g. Gray et al. 1995; Fee et al. 1996a; Harris et al. 2000), classification with artificial neural networks (e.g. Kim and Kim 2000), and wavelet-based methods (e.g. Letelier and Weber 2000) – claim to be able to sort spikes reliably for SNRs less than one. MTM therefore appears to be as effective, or more effective, than both conventional and more recent methods.

Lowering the reliable detection and sorting limits from $\sim 65\mu\text{V}$ to $\sim 50\mu\text{V}_{\text{pp}}$ affords a non-linear increase in the volume of potentially recordable neurons, which is further enhanced by the slow radial dependence of signal decay from neurons more than $\sim 70\mu\text{m}$ from the polytrode (*ie.* around this amplitude range, see chapter 4). This fact is reflected in the distribution of spike amplitudes, skewed towards low amplitude spikes (Figure 1.8F). The actual distribution is presumably even more skewed had the improved methods described in this chapter been used to construct the histogram. Estimates of the number of cells potentially recordable with tetrodes are similar. On purely geometric grounds (ignoring the shape of the field decay) a roughly 10-fold increase in the number of recordable cells is possible if spike detection and sorting methods can be made reliable in the range $45\text{--}65\mu\text{V}_{\text{pp}}$ instead of $\geq 65\mu\text{V}$ (Henze et al. 2000; Buzsaki 2004). Of course, detection of low SNR spikes is of little value unaccompanied by a spike sorting method of comparable sensitivity. In this regard the detection and sorting capabilities of the dynamic multiphasic spike detector and MTM are well suited.

3.5.2 Sorting of overlapping spikes

Previous studies looking at neuron interactions on short time scales ($< 2\text{ms}$) with multiple tetrodes or multi-shank silicon electrodes (e.g. Bartho et al. 2004) have by necessity restricted their analyses to units recorded on independent tetrodes or shanks. Two methods presented in this chapter, the spatiotemporal

spike partitioner and MTM, enable these sorts of analyses to be done on adjacent units recorded on the same polytrode shank, except when there is almost complete spatiotemporal spike overlap. The study of local (intracolumnar) cortical circuits will therefore be greatly facilitated by these methods.

Assessing the validity of accuracy claims for algorithms specifically designed to sort overlapping spikes is, as stated earlier, hampered by inadequate details of the test datasets. While the incidence of spike overlaps was usually provided, the degree of spike overlap (*ie.* the percentage of waveform superimposition) was *never* specified. Zhang et al. (2004) used iterative template matching of single channel data wherein spike events that did not match any existing template were fit against all possible combinations of phase-shifted templates until an acceptably low residual was found (according to a χ^2 criterion). Using simulated data comprising four relatively similar spike shapes with ~10% ($n=11$) of the seed spikes overlapping, they reported that only 2.4% of the total spike events were misclassified, however two-thirds of these misclassifications were of overlapping spikes. Unfortunately, since little information was provided about the degree of temporal overlap, the main claim of the paper is not open to further scrutiny. For the three (correctly classified) overlapping examples that were shown, the overlap appears to be around 50% in one, minimal in another, and perfectly superimposed in the third, however in the latter example both spikes were of the same polarity and their combined waveforms were unlike any of the base templates. Without a larger sample size and documentation of overlap statistics, it is impossible to determine whether these results were typical or exceptional. The same can be said of the other papers in the literature, where careful scrutiny of the (presumably) representative examples never showed overlap greater than ~50% of the spike waveforms (Zouridakis and Tam 1997; Hulata et al. 2000; Letelier and Weber 2000; Hulata et al. 2002; Takahashi et al. 2003; Laubach 2004; Quiroga et al. 2004).

While ICA- and WT-based decomposition methods probably do assist in the detection and sorting of overlapping spikes under certain circumstances, the contention here is whether they are more effective than MTM, especially given the computational burden of these algorithms. The overall impression is that MTM, by virtue of the multisite templates, performs as well as sorting algorithms that deal explicitly with spike overlap, except for occurrences of near complete spatiotemporal spike overlap (Figure 3.9C). In fairness, all the cited methods were

applied to single channel spike data, with one exception (Takahashi et al. 2003), and might perform better on multisite data. It would therefore be informative to use iterative template matching as described by Zhang et al. (2004) or Atiya (1992) as an adjunct method to MTM. In any case it should only be used for neurons whose template channels are shared (common i_{max}), and where the time series indicates frequent spike overlaps to justify the computational investment.

3.5.3 Refining template generation

The classification of spikes with MTM hinges on the prior availability of clean, representative templates for each of the active neurons in a recording ensemble. Template generation for polytrodes would preferably be entirely unsupervised, but in practice a few candidate templates usually require manual splitting, recombining, or removal of outliers. When this happens, the algorithm has usually confounded (*ie.* left unsplit) small amplitude, focal neurons, at the same time oversplitting large amplitude, homogenous clusters ostensibly from a single neuron with high intrinsic spike amplitude variability (*eg.* burst firing cells). Two interrelated aspects of the algorithm – robust selection of the cluster for splitting and when to stop splitting – are key to reducing these sorts of errors, in turn minimising post hoc user input.

The current implementation selects the cluster to split based on its variance, or maximum distortion (equation 3.8). Alternatively, the cluster with the maximum number of elements could be selected, but this implicitly assumes a uniform firing rate across the neural population, which is clearly not the case (Figure 3.10B). The trouble with the distortion measure, in spite of normalising for the number of cluster elements, is that it will select a broad, unimodal cluster in preference to two small, closely spaced heterogeneous clusters each having low variance. More specifically, variance measures are prone to such failures because they ignore the shape of the cluster distributions. k -means suffers from the same limitation in that the cluster assignment rule depends solely on a distance metric, irrespective of cluster variance (equation 3.12). In both cases, a useful approach is to model the cluster distributions as mixtures of Gaussians (*e.g.* Lewicki 1994). Given such a model, cluster boundaries may be delineated more precisely, and standard parametric statistical tests for unimodality and normality can be used to help

evaluate whether a given cluster requires dividing into one or more subclusters. The drawback is about half of all neurons deviate significantly from this assumption of isotropic, Gaussian variability, therefore mixture of Gaussian models are likely to overfit the data, splitting single unit clusters into multiple clusters (Fee et al. 1996b). A more sophisticated approach is to model the non-stationarity of the background noise and the intrinsic waveform variability, the latter predicted by the inter-spike interval (Fee et al. 1996a). Taking into account the shape of the clusters is more computationally intensive, and requires *a priori* knowledge of typical distributions of spike clusters, but doing so should result in more reliable partitioning and identification of mixed clusters.

Closely related to the best choice of which cluster to split are robust criteria for when to stop splitting. This is critical in divisive clustering algorithms such as this one, since they have no provision for combining clusters. Combining over-split clusters is almost as laborious as having to manually define templates from the raw data. The present criterion (equation 3.13) uses an arbitrarily defined threshold (d_{min}) representing the minimal 'acceptable' similarity between two spike templates. In practice this works fairly well, but deciding on a fixed, universal threshold is untenable because suitable thresholds will depend on the level of noise and variability of the spikes. The convergence profile of various cluster measures for different datasets (Figure 3.10A) suggests that the relative drop in d_{min} combined with a minimum average cluster distortion may be a better criteria. Furthermore, splitting all clusters and testing for the *absence* of critical drops in d_{min} instead of just the cluster with the maximum distortion should reduce the chance of prematurely exiting the algorithm before all heterogeneous clusters are unmixed. Finally, since none of the existing cluster measures use information about the shape of the underlying cluster distributions, another possible improvement would be to qualify any splits by assessing the normality of the two sub-cluster distributions.

In highly active recordings (eg. Figure 3.6A) small, focal spikes represent a small fraction of the overall signal energy compared with the entropy across the polytrode. In these circumstances multiple small spikes frequently associate into a large noisy supercluster that must be parsed by hand. One possible solution may be to include only a subset of the relevant channels in the feature space of a given template. Like the spatiotemporal trigger lockout, a mask centred on i_{max} could be

used to null all channels outside a set radius, which should not only aid in the identification and extraction of small spikes, but also lower the incidence of spurious transient templates comprised of disparate synchronous spikes.

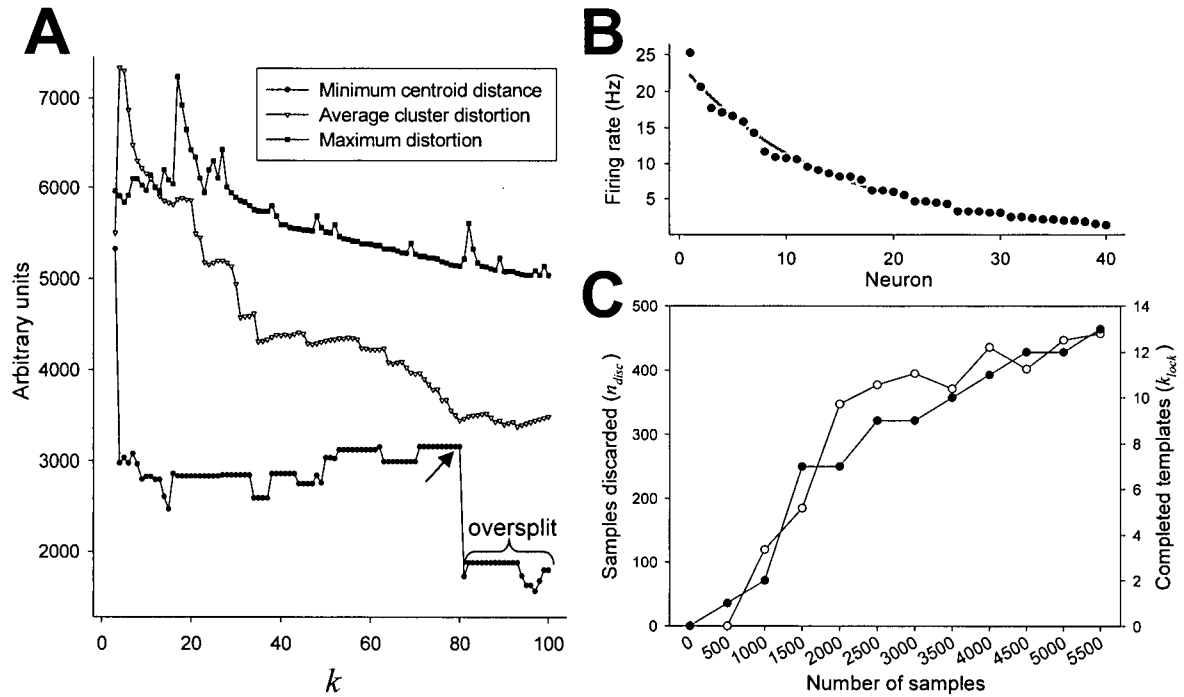


Figure 3.10 Automated exit criteria and sampling issues.

(A) The relative change in the minimum distance between any two cluster centroids (\bullet) combined with the average cluster distortion (∇) provides a robust criteria for stopping cluster splitting and determination of k (arrow). (B) The majority of spikes come from a small number of highly active neurons. An exponential function ($r^2 = 0.98$) describes the spike frequency distribution of a typical recording ensemble ($n = 40$). (C) A large number of spikes must therefore be sampled in order to assemble representative templates for infrequently firing neurons. The current algorithm deals with redundant samples by 'locking' completed templates containing sufficient samples (k_{lock} — \square), and discarding new samples belonging to these templates (k_{disc} — \bullet). When the majority of new samples were discarded, the entire template generation procedure was stopped.

An important corollary of using finite, random sampling of spike events to build the templates is that the sampling must be sufficiently exhaustive so that no spike class is overlooked. The firing distribution of a typical neural population is dominated by the activities of a few neurons, making adequate sampling even more difficult (Figure 3.10B, C). Short of visually inspecting the entire time series for missing spike classes, the statistics of the noise (*ie.* the standard deviation,

probability density function, and power spectrum) should match the raw data with spikes removed. Zhang (2004) gives an example of this method of validation. The current algorithm could also be extended to include a 'second-stage' sampling of the raw data with all first-stage spikes removed to further reduce the risk of missing infrequently firing neurons.

3.5.4 Improving sorting efficiency

The template generation stage of the sorting algorithm is very similar to that described by Atiya (1992) for single channel recordings, with some notable differences. Atiya's algorithm used iterative k -means to arrive at the number of likely neuron classes, increasing k and re-randomising the centroids with each iteration until the exit criteria were met. In order to improve the sorting efficiency for high dimensional multichannel spike templates, this highly inefficient randomisation step was replaced with divisive clustering (steps i to v) to obtain initial centroids and a conservative estimate of k . Since k derived from divisive pre-clustering was close to but consistently less than the definitive k arrived at following k -means, this modification did not increase the chance of overestimating the number of initial spike classes. This was critical as the algorithm has no mechanism for combining over-split clusters.

The fitting time for MTM is proportional to the number of templates, the number of sites and samples per truncated template, the sampling rate, and of course the recording duration. On an Athlon 1800+ processor running assembler code, the fitting of ~100 neurons to interpolated (100kHz) data currently takes ~20 times the duration of the original recording. It is not obvious that an effective sample rate of 100kHz is warranted, and so decimating the templates and fitting to the raw time series may offer a four or five-fold decrease in fitting time without compromising the sorting performance. The similar SNR limits of the extended multiphasic filter using a far cruder 'template' suggests such decimation will not adversely affect the sorting reliability. Selecting four or five canonical waveform points (Yang and Shamma 1988) per channel may also be just as effective as fitting the entire waveform. On the other hand, longer epochs (up to 2ms) that capture the shape of the AHP of slower spikes may add discriminatory value. Fee (1996b), based on information theory, suggests using 1.3ms waveform segments.

chapter 4

Neuron localisation & classification

4.1 Summary

Extracellular electrodes do not usually provide accurate information about recorded neuron location, nor any indication of cell type (hence the terms ‘unit’ and ‘multiunit’). This chapter describes an algorithm that capitalises on the fixed, closely-spaced site geometry of polytrodes to localise neurons in 3D cortical space. The algorithm was based on a mixed monopole-dipole field model of extracellular spike potentials and was able to generalise to arbitrary neuron orientation, tissue anisotropies, and cell morphology. Independent parameters described the volume conductances in each spatial dimension. Recorded neurons were assumed to lie in front of the polytrode. Levenberg-Marquardt optimisation was used to fit the multisite spike waveform amplitudes to the model functions. Estimated neuron locations emerged as non-overlapping spherical clusters within 150 μ m of the polytrode. Cluster locations moved concordantly with polytrode movements, making the algorithm a useful method for spike sorting unperturbed by electrode drift. Field potential spreads were consistent with the spike shapes and firing patterns of pyramidal cells and interneurons. These preliminary results suggest it is eminently possible to identify both the cortical location and type of multiple single units recorded extracellularly with high density polytrodes.

4.2 Introduction

It is unrealistic to ascertain the exact location of cells recorded with standard tungsten and wire-bundle electrodes, not only due to imprecision in histological track reconstruction, but because multiunit electrodes have the potential to record from thousands of neurons within a $\sim 150\mu\text{m}$ radius sphere, even further along the long axis of large pyramidal neurons (Fatt 1957; Rosenthal et al. 1966; Buchwald et al. 1973; Towe 1973; Drake et al. 1988; Henze et al. 2000; Blanche et al. 2003). Knowledge of the spatial relations between recorded neurons is particularly important for receptive field mapping studies and, for example, in studies of cortical circuits where laminar position may be important. Questions surrounding the nature of cortical 'micromaps' (see section 5.2) motivated the development of improved methods for localising recorded neurons. Precise positioning of the polytrode in the cortical layers (Figure 1.11, Figure 1.12) was the first step. The fixed site geometry and high resolution spike field potential measurements of the three column co-linear polytrode provided the basis for refining the localisation even further. Polytrodes, when combined with precise neuron localisation, are uniquely placed to resolve questions relating to cortical receptive field (RF) map organisation because they integrate high-density multiunit recording (microscale organisation) with contiguous coverage of whole cortical areas (millimetric scale organisation). Development of a biophysically-based field potential model, the focus of this chapter, was the other key element of the proposed method of neuron localisation.

4.2.1 Spike-related field potentials

A logical starting point to begin an exploration of spike-related extracellular field potentials is to consider the major transmembrane currents that are common to all action potentials. In 1952 Hodgkin and Huxley provided a complete empirical model of the cascade of ionic membrane currents underlying action potentials, seminal work for which they were awarded the 1963 Nobel Prize for medicine (with John Eccles). When the membrane reaches threshold, populations of voltage-gated sodium channels open regeneratively, producing a net inward

current at the axon hillock/soma and the sharp initial depolarisation phase of the action potential (Figure 4.1A). A complement of voltage-gated potassium channels repolarise the cell (net outward current), often producing an undershoot referred to as the after hyperpolarisation (AHP). At this point it is useful to distinguish two 'types' of membrane currents, active and passive. Active currents pass through voltage-gated channels, for example at the spike initiation zone during an action potential. According to Kirchoff's first law, currents in a closed system must be balanced by return currents of equal magnitude, at every instant in time, in this case through passive leak channels elsewhere in the cell membrane. In simple terms, the passive return current completes the circuit. As the extracellular space (ECS) has a finite resistance, transmembrane currents generate potential gradients – field potentials – measurable with extracellular electrodes.

The nature of neuronal field potentials has been the subject of investigation for decades, starting with a flurry of technically elegant single-unit experiments in the 1950's and 1960's. Of specific interest to the current modeling efforts, field potential mapping studies (Figure 4.1B, E) reveal that the amplitude, shape, and even polarity of the potential depends, among other things, on the morphology of the cell and the location of the recording electrode relative to the neuron (Fatt 1957; Bishop et al. 1962a, 1962b; Rosenthal et al. 1966; Drake et al. 1988). Figure 4.1C shows the spike field of a motor neuron at the peak of the depolarising phase of an action potential. The active current sink at the cell soma produces a negative potential gradient in the surrounding ECS. At the same time, passive current efflux elsewhere in the cell membrane produces a diffuse positive potential gradient. Therefore, the polarity of the recorded potential (relative to a distant reference electrode) will be positive if the electrode is nearer the site of outward current flow, and negative if it is nearer the site of inward current flow (assuming similar current sink/source densities). The painstakingly detailed measurements of Figure 4.1C were made sequentially in response to repeated antidromic stimuli, and a single cell would take hours to map. If the primary goal is to determine the location of the neuron, it would be easier to simply fill the neuron iontophoretically and reconstruct it in histological sections. However, with the advent of silicon-based multisite electrode arrays capable of making concurrent high resolution field potential measurements (Figure 4.1D), instantaneous mapping of individual or multiple neurons is now practical.

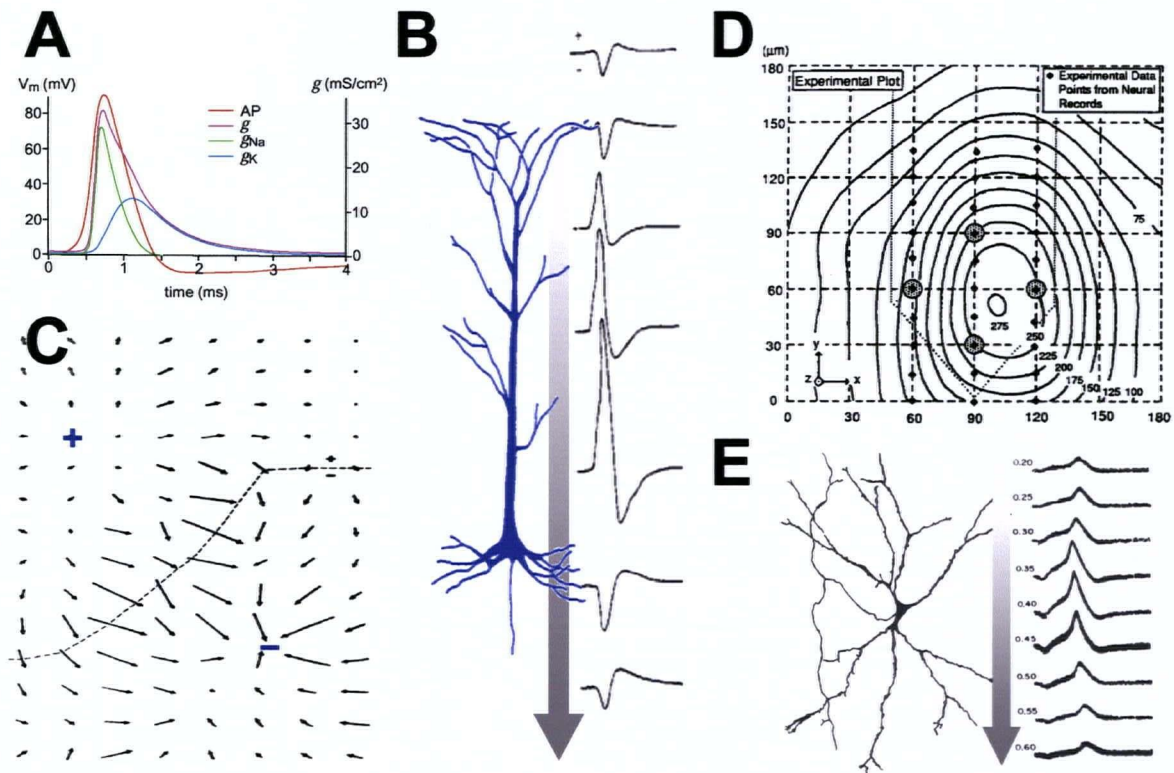


Figure 4.1 Spike field potentials.

(A) Timecourse of membrane conductance changes associated with an action potential (after Hodgkin and Huxley 1952; from Hausser 2000). (B) A single-unit electrode traversing the long axis of a pyramidal neuron exhibits a characteristic pattern of spike-related current sinks and sources (Humphrey 1979). (C) Electric field map evoked by antidromic activation of a motor neuron in the dorsal horn of cat spinal cord (Fatt 1957). The magnitude of the potential gradient is proportional to the vector length. The *dashed line* shows the zero potential boundary dividing the current dipole. (D) Topography of spike field potentials recorded with a silicon-based 'tetrode', assembled from sequential recordings made every 30 μ m (Drake et al. 1988). Note the ellipsoidal isopotential field lines. (E) Field potential of a motor neuron, at the positions indicated (mm). In contrast to putative pyramidal cells (B), neurons with radial processes have monophasic, unipolar spike waveforms, and zero potential beyond the dendritic arbours (Humphrey 1979).

Current flow in the ECS generates a complex field potential around the neuron, whose characteristics depend on the size and geometry of the cell, the distribution of the active and passive membrane conductances, their timecourse, and also the electrical properties of ECS (Humphrey 1979). Although the amplitude of the action potential does not vary significantly with cell size, the input impedance does (McCormick et al. 1985). Since larger neurons have lower input impedance, the peak transmembrane current – and in turn the extracellular

spike amplitude – will be greater for larger cells (Rall 1962). In addition to overall neuron size, the field potentials of stellate shaped neurons having predominantly radial dendrites will differ markedly from that of pyramidal neurons with extensive apical and basal dendrites. Stellate neurons give rise to focal ‘closed fields’ because the current sinks and sources tend to cancel out in such a way that the sink at the soma dominates the recorded potential everywhere except immediately adjacent to each dendrite (Rall 1962; Humphrey 1976). Accordingly, the extracellular potential is unipolar, monophasic (negative) and spatially restricted (Figure 4.1E). In contrast, pyramidal neurons have ‘open fields’ by virtue of their large apical dendrite. A highly focal current sink exists at the axon hillock during the peak of the action potential, whereas the apical and basal dendrites appear extracellularly as distributed current sources (Humphrey 1979, Figure 4.1B).

Independent of the characteristics of the generator potentials, the properties of the ECS and the electrode surface area are also major determinants of the voltages observed extracellularly. Spikes recorded with single unit electrodes are roughly an order of magnitude larger in amplitude than those recorded with multiunit electrodes. This is because the dimensions of the exposed recording tip are small compared to those of the steepest portion of the neuron’s extracellular potential field, as little as $\sim 20\mu\text{m}$ for small cortical neurons (Humphrey 1979). Multiunit electrodes that have dimensions approaching or exceeding this size introduce an isopotential conductor that in effect ‘shorts out’ the potential field, attenuating or even abolishing the spike amplitude. A direct consequence of the sensitivity of single unit electrodes to potential gradients on the scale of microns is that extracellular potentials recorded with these electrodes are heavily influenced by the subtle 3D morphology of the cell. The larger surface area recording sites of multiunit electrodes, including those of polytrodes, ‘spatially average’ these smaller potentials. What remains is a coarse field potential reflecting the principal current sinks and sources and gross cellular morphology, therefore we should not necessarily expect to see spike fields as complex as those portrayed in Figure 4.1B. In the context of field potential modeling for neuron localisation, this is, perhaps counter intuitively, a distinct advantage of using multiunit electrode sites, as it makes possible the development of a generalist model that does not require *a priori* knowledge of the exact morphology of every recorded neuron. Moreover,

knowledge of multiunit electrode properties suggests the appropriate level of detail with which to model spikes field potentials recorded with polytrodes.

The two major properties of the ECS that bear consideration are whether it is a purely resistive medium, or whether it has significant capacitive (reactive) properties (Plonsey and Heppner 1967); and secondly if, to a first approximation, it can be considered a homogenous, isotropic conducting medium. Adopting the nomenclature used by Nicholson (1973), the conductivity is homogeneous if it doesn't vary from place to place and isotropic if it is the same in all spatial directions. The cortex might be inhomogeneous due to differences in cell density (eg. granular vs. agranular cortical layers) but isotropic, homogeneous but anisotropic (eg. parallel with vs. perpendicular to the cortical surface), or perhaps even inhomogeneous and anisotropic. Any of these factors can impose subtle effects on the shape and extent of the field potential distribution (eg. Figure 4.1D).

4.2.2 Extracellular classification of cell type

Intracellular electrophysiologists have the luxury of filling impaled neurons to determine both the position of the soma and the extent of the neuron's dendritic and axonal arbourisations in histological reconstructions. Spine morphology, together with neurotransmitter histochemistry, can then be used to classify the neuron as either inhibitory or excitatory. Comparative electrophysiology of cortical neurons *in vitro* (McCormick et al. 1985) has revealed at least three distinct cell types, regular spiking (RS), bursting, and fast spiking (FS). These three functional classes possess different adaptation responses, maximal firing rates, and current-frequency relationships. FS neurons exhibit virtually no spike frequency adaptation, except with prolonged depolarisation over many seconds (Descalzo et al. 2005). Of particular significance to the present work, RS and bursting cells are generally spiny pyramidal neurons, whereas FS neurons are invariably aspiny or sparsely spiny stellate cells with radial dendrites (McCormick et al. 1985). Since gross differences in cellular anatomy should yield distinctive spike field potentials, this raises the possibility of classifying cell type by modeling the shape and distribution of the extracellular potentials recorded with polytrodes, in addition to spike firing-related measures.

Even though *in vitro* neurons retain the majority of their *in situ* electrophysiological properties (Connors et al. 1982), it is also possible to apply intracellular techniques *in vivo* (e.g. Fregnac et al. 1988; Azouz et al. 1997; Nowak et al. 2003). In a comprehensive comparative study of cat primary visual cortex, Nowak et. al. (2003) delineated four main functional cell types. In addition to verifying the existence of FS and RS cells similar to those found in guinea pig somatosensory cortex *in vitro* (Connors et al. 1982; McCormick et al. 1985), they also identified intrinsically bursting (IB) and chattering (CH) neuron classes. CH neurons were either pyramidal or spiny stellate neurons. IB neurons were present in all cortical layers, but were most abundant in layer 5. Similar discrete morphological and functional classes can even be found in the more primitive three layered turtle visual cortex (Connors and Kriegstein 1986). Across species and brain regions (Connors and Gutnick 1990), it appears cortical neurons constitute a small number of primary functional classes that are heavily influenced by their anatomy (Connors and Regehr 1996; Mainen and Sejnowski 1996), a fact that bodes well for their extracellular classification.

Several intracellular experiments, mostly in cat striate cortex, have mapped the visual RF properties of single neurons, in some instances with their dendritic and axonal projections reconstructed in exquisite detail (Van Essen and Kelly 1973; Kelly and Van Essen 1974; Gilbert and Wiesel 1979; Ahmed et al. 1997; Azouz et al. 1997; Hirsch et al. 2002; Contreras and Palmer 2003; Hirsch et al. 2003). Collectively these studies have examined a number of RF properties, but the small sample population accumulated to date has not yet revealed a consistent association between any single RF property and one of the structural or functional cell types identified here. This highlights the fundamental disadvantage of *in vivo* intracellular recording. It is limited to recording from single cells for short periods of time, prohibiting detailed RF characterisation, and is too unstable to be done in awake behaving animals. If the same detailed morphological identification could be inferred from extracellular multiunit recordings, this could lead to findings that would revolutionise our understanding of cortical circuits.

The desire to classify cell type in extracellular recordings goes back to the early days of single unit recording. Mountcastle (1969) hypothesised that units with “thin” waveforms and high spontaneous activity were non-pyramidal cells, whereas broader waveforms with lower spontaneous activity were likely to be

from pyramidal neurons. Mountcastle's speculation is consistent with the intracellular data and more recent multiunit extracellular work. Using silicon based polytrodes, Bartho and colleagues (2004) applied cross-correlation analysis to thousands of simultaneously recorded cell pairs in rat neocortex and showed that extracellular spike width is a near perfect predictor of whether or not a cell is inhibitory or excitatory. In rabbit somatosensory cortex, suspected inhibitory interneurons have also been categorised extracellularly on the basis of spike width and peripheral response latency (Swadlow 1991, 1994). Therefore, it is possible to divide cortical neurons into two major, albeit heterogeneous, functional groups using simple spike metrics, without necessarily recording from functionally connected cell pairs.

Other extracellularly derivable parameters may also be useful indicators of neural type, such as spike amplitude (Humphrey and Corrie 1978; Gur et al. 1999), the frequency and pattern of spike bursting, and anatomical position (*ie.* cortical layer) (McCormick et al. 1985; Nunez et al. 1993; Gray and McCormick 1996; Azouz et al. 1997; Nowak et al. 2003). The neuron localisation algorithm may therefore play a central role in the eventual neuron classification scheme. Finally, no study has yet systematically explored the discriminatory value of spike field potential spread or symmetry, and this too can be derived from the spike field potential model described later in the text.

4.2.3 Foundational work on neuron localisation

Phil Hetherington, a former postdoctoral fellow in the Swindale laboratory, was the first to pursue the idea of localising neurons with silicon electrode arrays. He reasoned that if spikes from individual neurons were recorded by three or more electrodes, and these electrodes had fixed, known configurations, it should be possible to localise active neurons by triangulation of their spike amplitudes (Hetherington and Swindale 1998; Hetherington et al. 1999). This is analogous to global positioning systems that use differential transmission latencies from orbiting satellites to calculate position and altitude on the earth, only here the triangulation is based on differential spike amplitudes. The algorithm (Hetherington et al. 1999) returned spatially-localised, spherical clusters and, with the right parameters, plausible location predictions (Figure 4.2A). Attempting 3D

source triangulation with a 1D array of electrode sites is inherently ambiguous; the neuron clusters shown to the right could just as likely been located equidistant anywhere along a circle surrounding the polytrode. A 2D array of staggered, more closely spaced sites (Figure 4.2B) removed this lateral ambiguity, and also improved the precision of the location predictions (*ie.* tighter clusters).

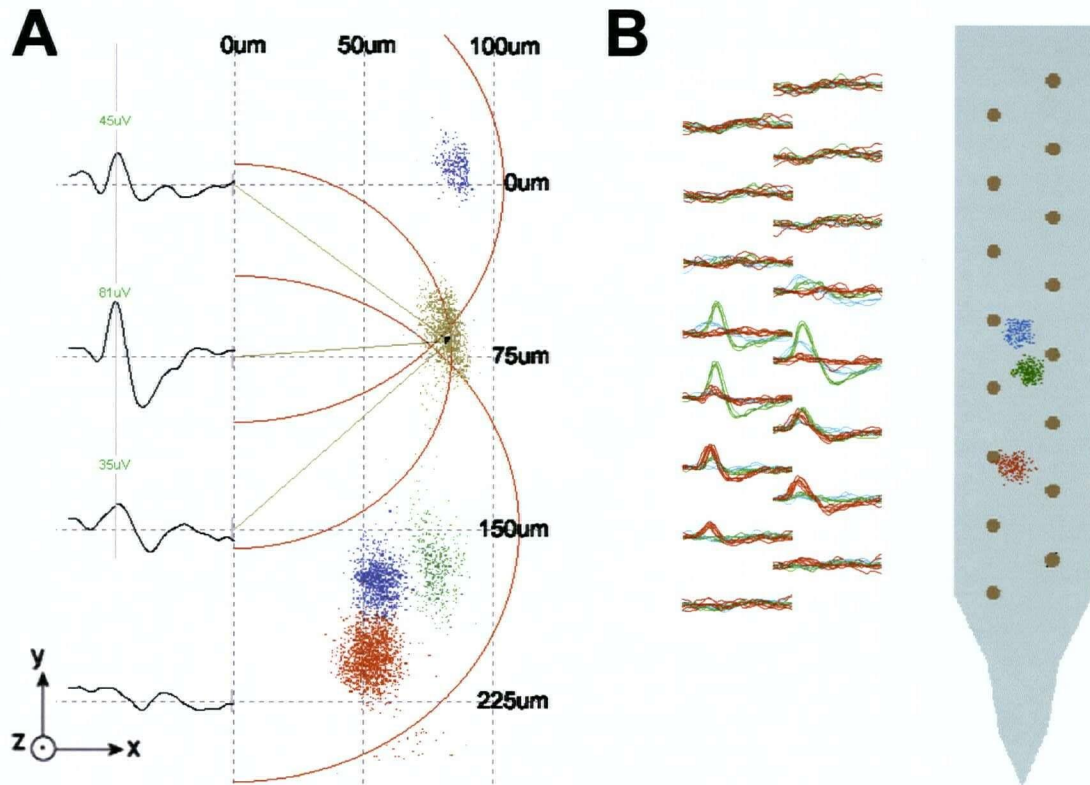


Figure 4.2 Neuron location by spike amplitude triangulation.

(A) Initial work on neuron localisation was based on multiunit recordings made with a four channel polytrode with vertically aligned sites spaced 75μm apart. The amplitudes of each spike across sites (*left panel*) were used to define three spheres (*red lines*); the intersection of these spheres gave an estimate of the source location for that spike. Repeated triangulations for every spike in the recording provided an indication of each neuron's location (*coloured clusters*). (B) A polytrode with more closely spaced electrode sites (65μm apart) in a staggered two-column configuration gave improved localisation and removed the lateral ambiguity inherent in predictions made with the 1D array.

In spite of these promising preliminary results, there were several weaknesses of this approach to neuron localisation. Foremost among these was the inability of the algorithm to generalise. In other words, a given set of parameters and decay functions (*ie.* voltage-distance relationships, spatial decay constants, maximal

spike amplitude, *etc.*) that worked well for one ensemble of neurons failed to work for other neurons. The algorithm either failed to converge or gave unreasonable location predictions with all neurons located up against the polytrode or hundreds of microns away. In other cases the fitting algorithm was not adequately constrained, having too few spike amplitude measurements to converge reliably. Another likely reason why the algorithm failed to generalise was the somewhat arbitrary choice of voltage-distance functions describing the decay of spike amplitudes. No provision was made for aspherical fields (*ie.* the same rate of amplitude decay was assumed in all axes), nor was there any allowance for differences in the size of neurons with respect to the polytrode.

Since the method of triangulation is not formulated in such a way to easily incorporate these sorts of elaborations, this approach was abandoned in favour of more explicit models based on electrostatics. These criticisms are in no way meant to lessen the significance of these first attempts at neuron localisation. On the contrary, this work provided proof of concept that neuron localisation with polytrodes is feasible; it motivated the development of higher density polytrodes with additional sites spread over three columns (in particular the 54 μ m \times 1b design); and finally, it justified the need for the more sophisticated biophysical model presented in this chapter.

Consequently, the objectives of the modeling work described here are two-fold: primarily, to establish a realistic model that describes field potentials recorded with polytrodes, with the aim of localising neurons in 3D cortical space, and with an accuracy comparable to that obtainable with intracellular histology; and secondly to use information derived from this model, in conjunction with spike waveform shape, firing statistics and other functional attributes, to classify neuron type based solely on extracellular multiunit recordings.

4.3 Methods

4.3.1 Model overview

The neuron localisation algorithm can be broken down into several components (Figure 4.3), each of which is detailed in the following subsections.

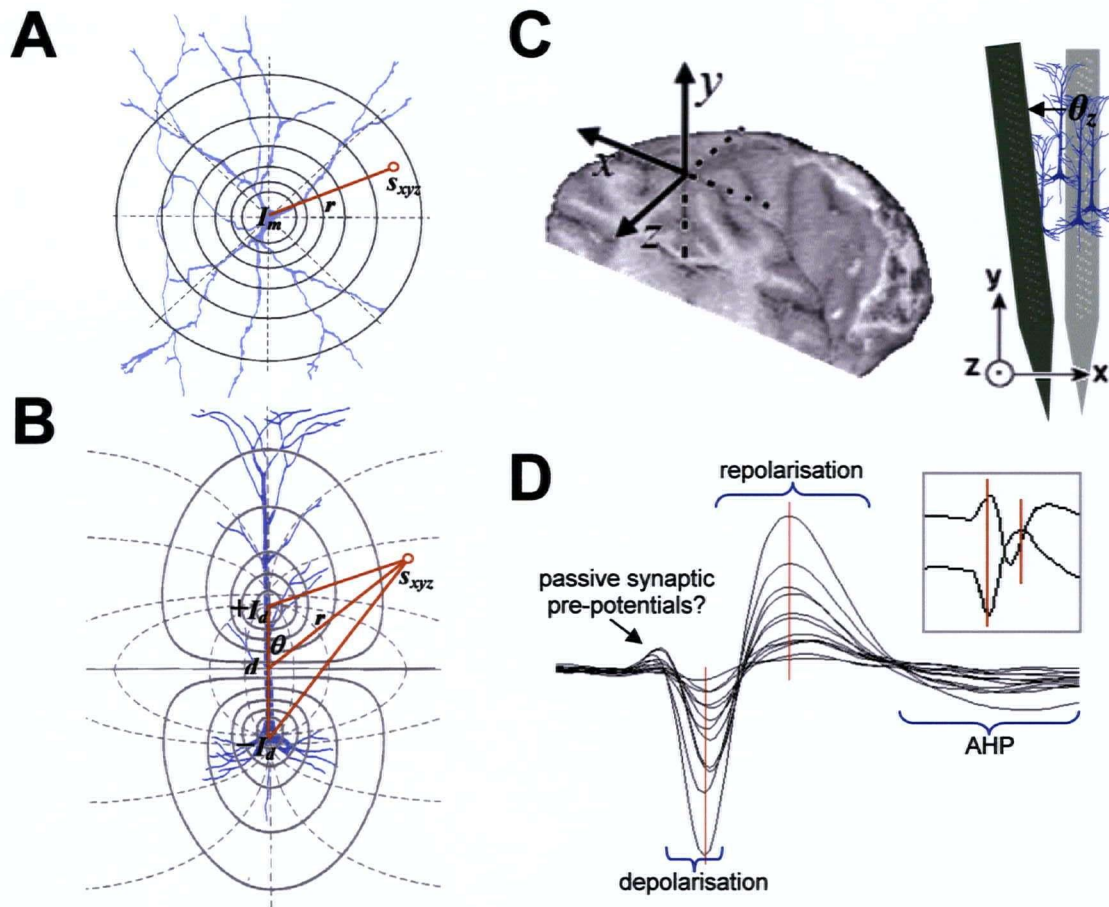


Figure 4.3 Components of the neuron localisation model.

(A) Monopole and (B) dipole electric fields associated with stellate and cortical pyramidal neurons, respectively. Dotted lines indicate the direction of current flow. Solid lines represent isopotential field lines. See Methods for parameter details (note: two possible scenarios are portrayed here, and neither is meant to imply that monopolar fields are synonymous with stellate cells, nor that dipole fields can be equated with pyramidal cells). (C) A rectangular coordinate system was used to relate the position and orientation of the polytrode with the major axes of the brain, or more specifically, the cortical gyri. The x and z axes corresponded to the plane parallel to the cortical layers, and the y axis was perpendicular to this plane. Rotations around the z and x axes (out of the page) were represented by θ_z and θ_x . (D) Superimposed multisite spike waveform traces. Cross-channel spike amplitude measurements were made at an instant in time (red lines). The *inset* shows similar measurements of a dipole field on two adjacent recording sites.

To summarise, a mixed monopole-dipole current source model of extracellular spike field potentials was used (Figure 4.3A, B). Geometric linear transforms were used to account for the arbitrary displacement and orientation (in a Cartesian coordinates) of recorded neurons with respect to the polytrode (Figure 4.3C). The cortical gray matter was assumed to be a purely resistive,

homogeneous conductor, however no assumption was made about the isotropy of conductivity. Independent parameters described the conductivity in each spatial dimension. Recorded neurons were assumed to lie in front of the polytrode. The non-conductive shank of the polytrode acts as an insulator and therefore only neurons in front of the polytrode are recordable (Drake et al. 1988; Anderson et al. 2001). It was also assumed that the extracellular signal decay was isotropic along the two spatial dimensions coplanar with the cortical layers, in order to make it feasible to estimate neuron locations in three-dimensions using a 2D planar electrode array. To calculate neuron locations, the Levenberg-Marquardt gradient descent algorithm (Press et al. 1994) was used to fit the multisite spike waveform amplitudes (Figure 4.3D) to the model functions.

4.3.2 A mixed monopole, dipole field potential model

A useful framework for estimating neuron location is to regard the action potential generated at the axon hillock/soma as a point source[†] and model the extracellular spike field potential around the neuron at an instant in time. Under *quasistationary* conditions (defined below) electrostatic theory can be used to determine the potential at any point in space. Given a single current point source, or *monopole*, the potential ψ_m in an isotropic, homogenous conducting medium is:

$$\psi_m = \frac{I_m}{4\pi\sigma r} \quad (\text{Malmivuo and Plonsey 1995}) \quad (4.1)$$

where I_m is the magnitude of the current source, σ is the conductivity of the medium, and r is the distance (*eg.* of a single electrode site) from the centre of the monopole according to Pythagorean theorem:

$$r = \sqrt{x^2 + y^2 + z^2} \quad (4.2)$$

Note that the potential varies inversely with r , the current flows radially, and the isopotential surfaces are concentric spheres (Figure 4.3A). Although a monopole is not physically realisable due to conservation of charge requirements, this approximation is valid if the return current is sufficiently distant from the axon hillock, as in the case of large polarised cortical pyramidal neurons, or if the passive return current through leak channels is distributed diffusely across the

[†] For convenience the term 'source' refers to both current sinks and current sources. The phase of the action potential determines the direction of current flow at any instant in space and time, and therefore whether a given point in extracellular space is at a net negative potential (sink) or net positive potential (source).

whole cell membrane. In either of these situations the counter currents will not evoke potentials that are large enough to be recordable by a multiunit electrode.

A more general form of equations 4.1 and 4.2 that make no assumption about the isotropy of the tissue in each spatial dimension is:

$$\psi_m = \frac{I_m}{4\pi} \cdot \frac{1}{\sqrt{\sigma_x^2 x^2 + \sigma_y^2 y^2 + \sigma_z^2 z^2}} \quad (4.3)$$

where σ_x , σ_y , and σ_z are the conductivity scalars in the x , y , and z spatial dimensions, respectively (Figure 4.3C). In the striate cortex, neuronal processes run predominantly perpendicular to the cortical gyri (*ie.* the apical dendrites of pyramidal cells are almost parallel to each other). Therefore, it is reasonable to expect that the conductivity is isotropic in the xz plane, but may be different along the y axis:

$$\sigma_x \approx \sigma_z \neq \sigma_y \quad (4.4)$$

The effect of tissue anisotropies is to distort the shape of the isopotential surfaces such that they are no longer spherical, but spheroidal. In the extreme case where the return current is focal and very close to the source, for example a small non-pyramidal cell, ψ_m approaches zero as the current sinks and sources cancel each other. In the intermediate case, two relatively close monopoles of opposite polarity constitute a *dipole*. The radial dependence of a dipole field potential in an isotropic, homogenous conducting medium is given by:

$$\psi_d = \frac{I_d d \cos \theta}{4\pi\sigma r^2} \quad (\text{Malmivuo and Plonsey 1995}) \quad (4.5)$$

where I_d represents the equal magnitude active and passive currents of the two poles, d is distance separating the poles, and θ is the angle (in polar coordinates) between the axis of the dipole aligned with the y axis, and the radial direction of r (Figure 4.3B). Observe that ψ_d is maximal when sampled along the axis of the dipole, when $\theta=0$ ($\cos \theta=1$), and zero when sampled midway between the poles ($\theta=90$, $\cos \theta=0$). Dipole potentials vary inversely with r^2 , decaying even more rapidly in space than monopole potentials. Equation 4.5 holds for $d \ll r$, however the monopole component (equation 4.1) begins to dominate as d increases. Therefore it is reasonable to predict that larger pyramidal neurons will have fields that are best fit by a monotonically decaying monopole distribution, whereas smaller neurons may be better described by a dipole. By the same relation, the contribution of the monopole increases for neurons very close to the

electrode (r approaching d). Spatial averaging of large surface area multiunit electrodes means highly focal dipoles will only be detected for electrode sizes $\ll d$.

As with the monopole, the dipole equation can easily be extended to allow for possible tissue anisotropies:

$$\psi_d = \frac{I_d d \cos \theta}{4\pi} \cdot \frac{1}{\sigma_x x^2 + \sigma_y y^2 + \sigma_z z^2} \quad (4.6)$$

Here again the prediction is that $\sigma_x \approx \sigma_z \neq \sigma_y$.

4.3.3 Coordinate transforms

The dipole model, even without accounting for possible tissue anisotropies, has an intrinsic directional component because its electric field is spatially asymmetric. A monopole field embedded in an anisotropic conductive medium is also aspherical. For a rigid array of electrode sites, as opposed to a single radially defined point in space (r), it is more practical to consider the position and orientation of the polytrode as a single entity with respect to the brain. Likewise, to allow for arbitrary displacement and orientation of a given neuron with respect to the polytrode (Figure 4.3C) it is necessary to reconcile the coordinate system conforming to each neuron with the coordinate system aligned with the polytrode shank. Taking the latter as the principal coordinate system, the following linear transformations can be used for rotations about the z axis:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos \theta_z & -\sin \theta_z & 0 \\ \sin \theta_z & \cos \theta_z & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} \quad (4.7)$$

where (x, y, z) gives position relative to the origin in the rotated coordinate system, and (x', y', z') is the position in unrotated coordinates. Similarly, for rotations about the x axis:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_x & -\sin \theta_x \\ 0 & \sin \theta_x & \cos \theta_x \end{bmatrix} \begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} \quad (4.8)$$

Given that the polytrode is inserted more or less perpendicularly to the cortical gyrus, and the additional imposed constraint that $\sigma_x = \sigma_z$, rotations about the y axis are degenerate and do not need to be considered further. Rotations

around the other two axes can be combined, using matrix multiplication, along with translation for any recording site in the polytrode array:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos \theta_z & \cos \theta_z \sin \theta_x & \sin \theta_x \sin \theta_z \\ \sin \theta_z & \cos \theta_x \cos \theta_z & \sin \theta_x \cos \theta_z \\ 0 & \sin \theta_x & \cos \theta_x \end{bmatrix} \begin{bmatrix} x' - s_x \\ y' - s_y \\ z' - s_z \end{bmatrix} \quad (4.9)$$

where s_x, s_y, s_z are the coordinates of a given recording site.

4.3.4 Extracellular tissue model

As with the CSD analysis (section 1.3.5), the tacit assumption of the field potential equations presented above is that the tissue is a purely resistive medium (ohmic) over the physiologically relevant frequency range (*ie.* 0.1~10,000Hz). This assumption is supported by empirical evidence (appendix A.3), and direct brain impedance measurements (Kay and Schwan 1956; Schwan and Kay 1957a, 1957b). The capacitive (reactive) component of tissue impedance is thus negligible, and under such *quasistatic* conditions all currents and fields vary synchronously (Plonsey and Heppner 1967). Only the tissue conductivity (or resistivity) need be specified in the model, which greatly simplifies matters. Accordingly, the extracellular fields associated with the different frequency components of the spike waveform decay in a similar fashion (Figure B.2), and thus a common field model can be used for both FS and RS neurons.

The other pertinent property of the tissue was alluded to earlier, that of the conductivity (σ) in each spatial location and dimension. In the current model it was assumed that the conductivity of the cortical grey matter was homogenous, unperturbed by either the laminar structure of the cortex or the pial and white matter boundaries. Isotropic conductivity was assumed in the xz plane, but could differ in the y axis, in accordance with the columnar architecture of the cortex (a necessary constraint in any case, given the 2D polytrode).

4.3.5 The neuron localisation algorithm

Model parameters are given in Table 4.1, which includes the constraints and parameter seed values. The model has ten free parameters, three representing neuron location, two for axial rotations, three relating the conductivity in each spatial dimension, and two for the 'intrinsic spike amplitudes' of the monopole

and dipole components (see below). These parameters were constrained as follows: neuron location in the z axis was greater than zero (*ie.* in front of the polytrode shank); rotations about the x and z axes were limited to 90° in either direction; the conductivity scalars in each dimension were also free parameters, with $\sigma_x = \sigma_z$; finally, the upper bounds on the intrinsic spike current amplitudes (ψ_m and ψ_d) were set to 20nA, on the basis that polytrode potentials (presumably some finite distance from the neuron) larger than ~ 1.5 mV have never been observed, and the average resistivity of the cerebral cortex is around $300 \Omega \cdot \text{cm}$ (Freygang and Landau 1955; Ranck 1963). The inclusion of independent monopole and dipole amplitude parameters does not imply the existence of two independent neural sources, but rather allows the algorithm to conform to one model or the other (or both equally) depending on the dominant field distribution. Both field components (equations 4.3 and 4.6) were fitted as:

$$\psi = \psi_m + \psi_d \quad (4.10)$$

in addition to the coordinate system rotations. Note also that $I_d \cdot d$ in equation 4.5 was represented by the single parameter I_d for the localisation algorithm.

The 'bootstrapping' phase of the modeling was on a large representative library ($n=188$) of multisite spike templates extracted with the MTM spike sorting algorithm described in chapter 3. These templates comprised averaged spike waveforms (≥ 100 samples) with appreciable signal on up to thirty recording sites when recorded with the highest density polytrode (54 μ map1b design). Data from the staggered two and three column polytrodes (54 μ map2b, and 1a) were also modeled – with spikes appearing on up to twenty sites – was still sufficient to constrain the fits in most cases. Spike amplitudes (ψ_{xyz}) were determined for each recording site (s_{xyz}) at the initial phase spike peak *or* valley on the maximum amplitude spike channel (*ie.* not the peak-peak spike amplitude on each channel, Figure 4.3D). These two points in time obviously afford the highest SNR, but were also selected because these spike components were qualitatively less variable than other phases of the waveform, such as the AHP, and are reported to conform most closely with intracellular current (Henze et al. 2000).

Separate simulations were also run on unsorted, unaveraged spike waveforms to assess whether the localisation algorithm generated spatially distinct clusters suitable for spike sorting (as in Figure 4.2).

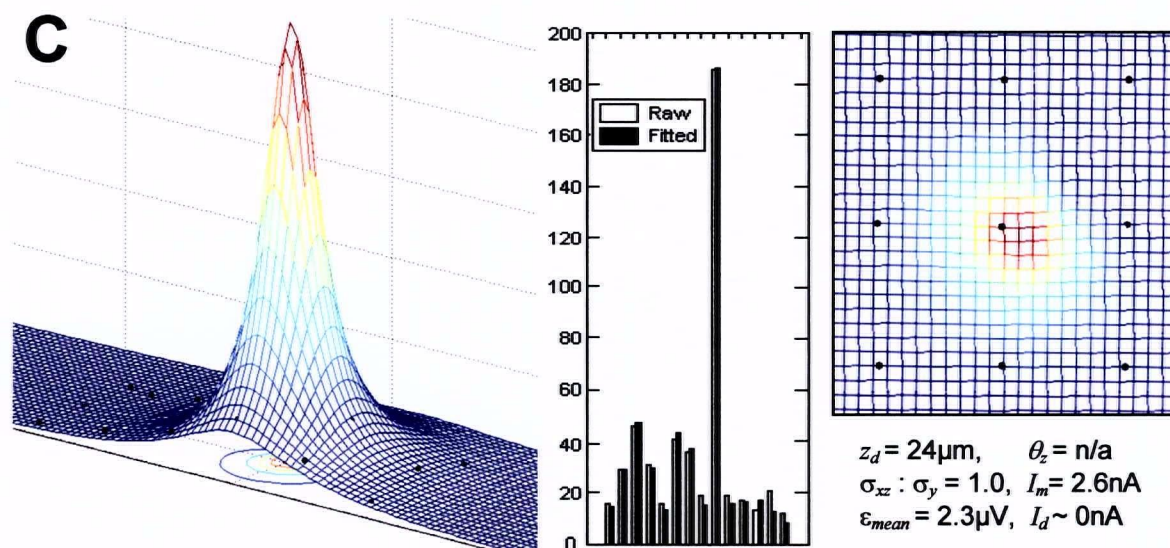
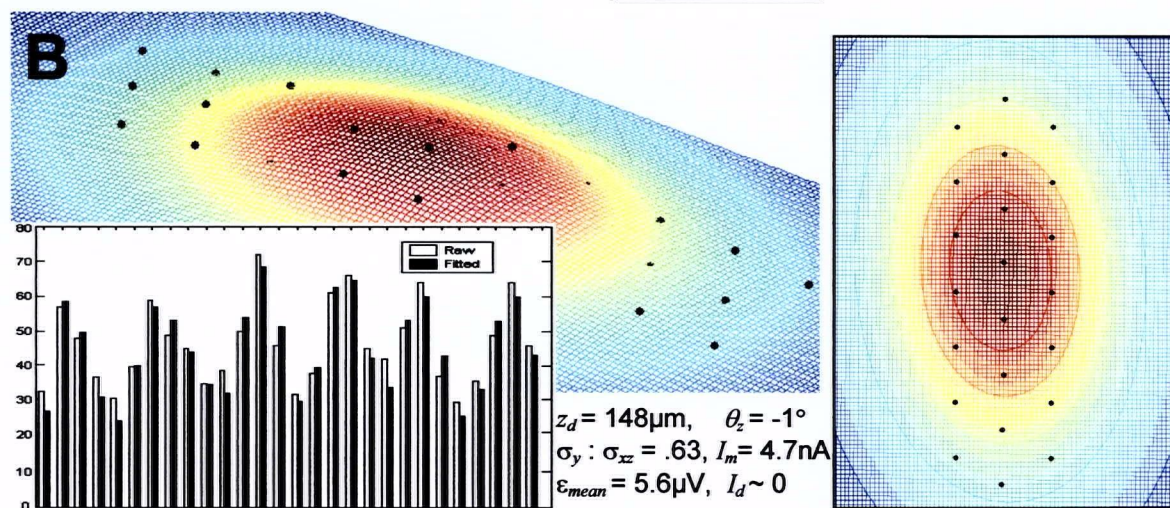
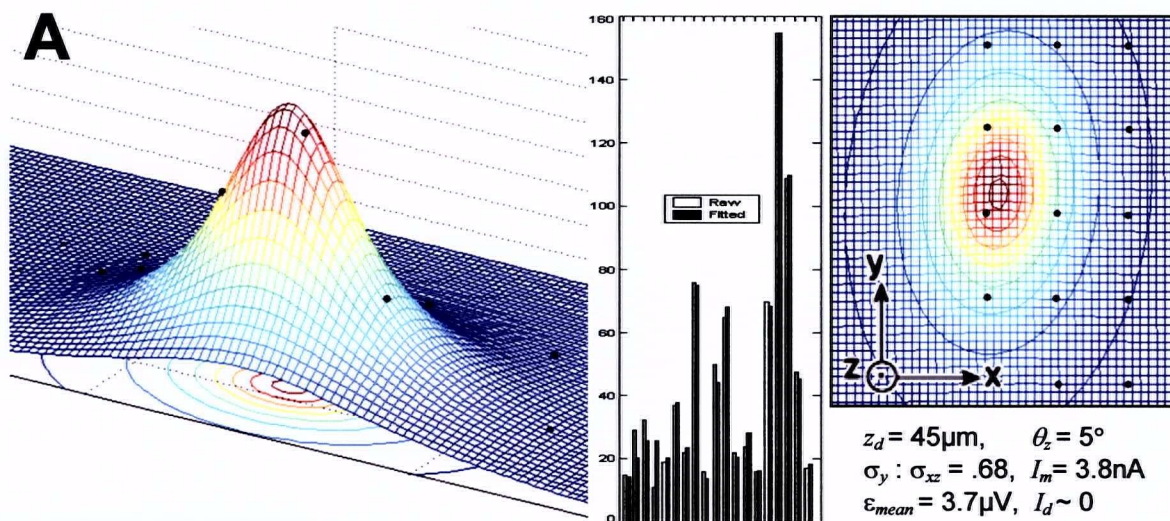
Table 4.1 Neuron localisation model parameters and constraints.

Symbol	Meaning	Constraints	Seed values
s_{xyz}	electrode site coordinates (μm)	fixed, $z = 0$	known
ψ_{xyz}	spike amplitude at site s_{xyz} (μV)	fixed	measured
x, y, z	neuron location, relative to the electrode (μm)	$z \geq 0$	$x = y = i_{max}, z = 50$
θ_z, θ_x	rotation about the z - and x -axes (degrees)	$-75 \leq \theta \leq 75$	0, 0
$\sigma_x, \sigma_y, \sigma_z$	conductivity scalar in each dimension (S.m^{-1})	$\sigma_x = \sigma_z$	$\sigma_x = \sigma_y = \sigma_z = 0.3$
I_m	intrinsic current, monopole component (nA)	≤ 20	3
I_d	intrinsic current, dipole component ($\text{nA} \cdot \mu\text{m}^{-1}$)	≤ 20	3

4.4 Results

4.4.1 Neuron localisation

The majority of neurons in the sample population, around 90%, had fields characteristic of a monopole distribution, with monotonically decaying potentials (Figure 4.4). The model-derived isopotentials of these cells were spherical or prolate spheroids, and in the latter case the long axis was generally aligned with the y axis or had a small rotation about the z axis (though never more than $\pm 60^\circ$). Neurons of average peak spike amplitude ($\sim 150 \mu\text{V}$) were predicted to be around $50 \mu\text{m}$ from the polytrode shank (z_d , Figure 4.4A), and no neuron was farther than $148 \mu\text{m}$ (Figure 4.4B). The most distant neurons had broadly distributed fields, in some instances present on half the recording sites spanning $500 \mu\text{m}$ or more. In contrast, neurons estimated to be less than $40 \mu\text{m}$ from the polytrode were characterised by steep voltage gradients, with appreciable signal on only two or three sites (Figure 4.4C). Neuron translations (xy position) were typically close to the site of maximal amplitude (i_{max}), however a few neurons had estimated locations $165 \mu\text{m}$ beyond the top and bottom row of recording sites, farther than the z or x axis range of prediction but in accord with the prolate field shapes.



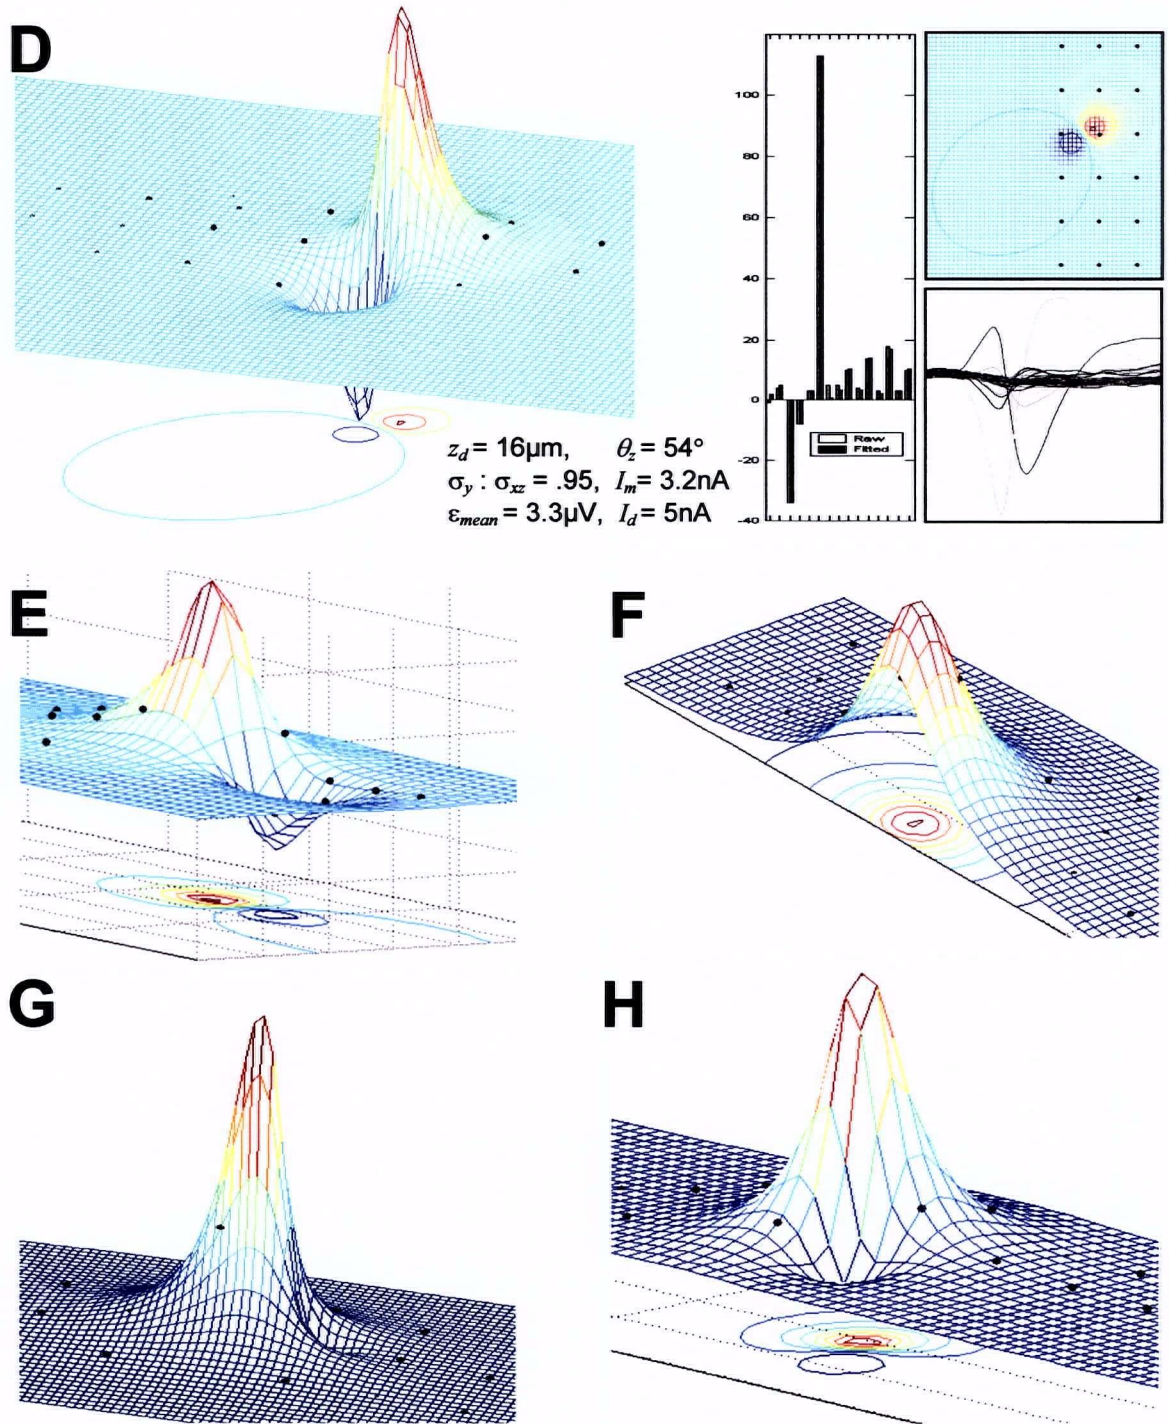


Figure 4.4 Representative model fits.

In each example the coloured mesh plot is the modeled 2D voltage distribution across the polytrode. Sites (*black spheres*) represent actual spike voltage measurements, arranged in their native configuration, with amplitude on the z -axis (*ie.* height above or below the mesh ground plane). Bar charts compare the measured and fitted spike amplitudes (in μV , for all sites with spike amplitudes over $10\mu\text{V}_{\text{pp}}$). *Insets* show the modeled field potentials in the xy plane. Fitted model parameters are shown, in addition to the average fit error per electrode site (ϵ_{mean}). See text for further details

Average fitted values for the conductivity parameters were 0.28 ± 0.12 for σ_y and $0.42 \pm 0.17 \text{ S.m}^{-1}$ for σ_{xz} ($n=188$). The shape of the isopotentials was reflected in the ratio of the conductivities ($\sigma_y : \sigma_{xz}$), hereafter referred to as the apparent conductivity *tensor*. For the majority of cells that had prolate fields, the conductivity tensor ranged from 0.53–0.77. Neurons with spherical fields had a unitary tensor. The very existence of spherical fields suggests that the shape of the field potentials is, at least in part, due to differences in the geometry of the source currents, not, as was assumed, attributable solely to tissue anisotropy. Anatomical and physiological evidence from other studies may lead to the correct interpretation, an issue that is considered further in the Discussion.

Taken together, these results agree with what would be expected given the nature and geometry of monopoles; small amplitude neurons with dispersed fields are distant from the polytrode, whereas large peak amplitude, focal fields are from neurons located very close to the polytrode.

The remaining 10% of neurons studied (19/188) were poorly described by the basic monopole model (with markedly higher fit residuals). These neurons had obvious dipole components (Figure 4.4D, E). Field potentials showing evidence of dipoles usually had inverted waveforms on a few surrounding recording sites, and these neurons were invariably less than $25\mu\text{m}$ from the polytrode, perhaps unsurprising given the inverse squared radial dependence of dipole potentials. The positive pole of the dipole was typically more distributed than the negative pole (*ie.* the main current sink during the initial phase of the action potential). More subtle dipole components were also observed (Figure 4.4G, H). These lacked inverted waveforms, but had signal decays too rapid to be described by a monopole.

Other simulations with simple linear decay constants (*ie.* $\psi \propto -r$) described none of the field distributions well. Fits with functions such as exponentials and Gaussians were similarly poor. Elaborations of the model, such as the addition of quadrupoles (*ie.* $\psi \propto 1/r^3$), did not improve the quality of the fits. Therefore, the mixed monopole-dipole model appears to provide a minimal but adequately detailed description of neuronal field potentials as recorded by large surface area multiunit electrodes.

The neuron localisation algorithm was applied to a recording made at a single depth, with the polytrode inserted perpendicular to the cortical surface (Figure

4.5). Estimated locations for these 82 neurons (a subset of the 188 neuron dataset) were fairly evenly distributed across the polytrode. Predicted z distances were plausible, with a few neurons against the shank, and no neuron further than $130\mu\text{m}$ from the shank. In the xz plane the density of recorded neurons reached a peak $\sim 55\mu\text{m}$ anterior to the polytrode shank, presumably as the volume of recordable neurons increased, and stopped abruptly beyond $\sim 100\mu\text{m}$, an artefact of setting a simple $100\mu\text{V}$ amplitude threshold during spike extraction and sorting.

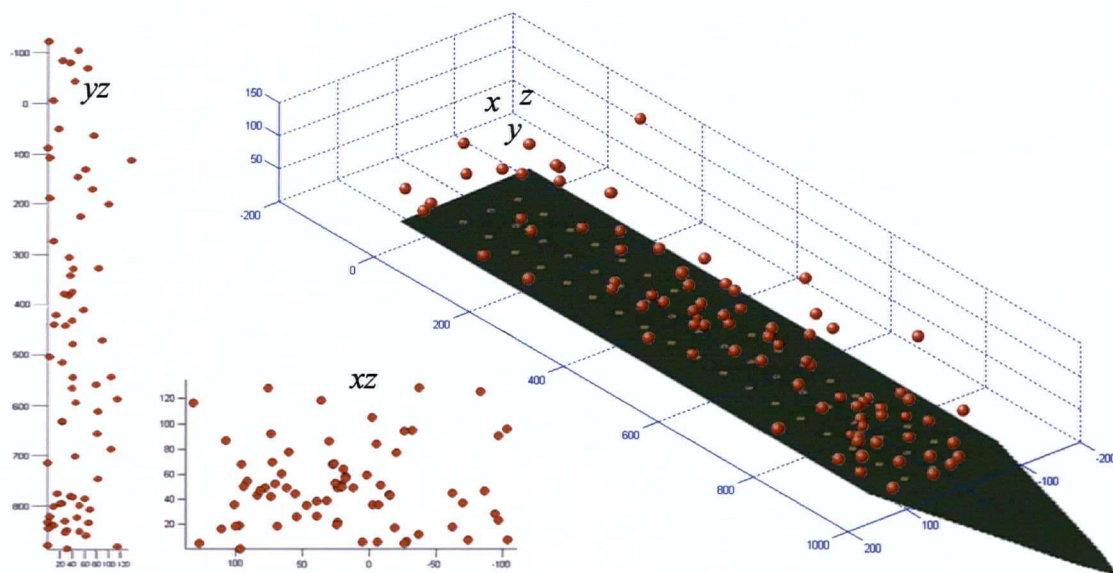


Figure 4.5 Ensemble neuron localisation.

Isometric projection of the estimated 3D location of 82 simultaneously recorded neurons (*red spheres*) relative to a $54\mu\text{m} \times 1\text{mm}$ polytrode. *Insets* are 2D projections of the same neurons in the plane indicated. All dimensions are in microns.

4.4.2 Model precision and stability

Fit errors for all neurons were within noise estimates (*ie.* $\sim 50/\sqrt{100} \approx 5\mu\text{V}$ per channel after averaging), with only slightly higher residuals for the most distal neurons, as might be expected given the increased likelihood of waveform distortion by greater numbers of interposed neurons. Model parameters such as intrinsic spike amplitude, the conductivity tensor, and rotational parameters were statistically independent of neuron location. A weak but statistically significant ($p < 0.05$) inverse correlation was evident between the conductivity parameters

and the two intrinsic current parameters, possibly indicating a degree of redundancy (*ie.* overparameterisation). With the exception of a few neuron z locations up against the polytrode, none of the model parameters were close to their (albeit loose) parameter bounds. There were no systematic differences in the model fits for different polytrode designs (Figure 1.4), each having variable shank widths and electrode site to edge distances, suggesting that any distortion of the isopotentials due to the polytrode shank, if not negligible, were at least similar for all polytrodes. Fits on the 2-column field data did, however, converge less reliably than data derived from the 3-column polytrodes (see below).

Stability analysis was performed for all free model parameters to determine if the predictions were affected by different seed values. Other than varying the number of iterations to converge on a solution, none of the parameters (including neuron location) were dependent on the choice of initial seeds, indicating a well constrained model that was not susceptible to local minima. The only parameters that behaved erratically were the rotational parameters for neurons with radially symmetric field distributions (unitary tensor), unsurprising given that such rotations are degenerate in the isotropic case.

Finally, while few simulations in the test dataset failed to converge (5/188), this was invariably on field data from the 2-column polytrode. In two cases there were fewer significant spike amplitude measurements than free model parameters. In addition, the standard deviations of the two rotational (θ_z, θ_x) and lateral conductivity (σ_{xz}) parameters were higher for the 2-column fits. Forward modeling simulations using 'generative' 2-column data (*ie.* simulated field distributions based on an identical model without additive noise) confirmed this instability. These results highlight the importance of the additional constraints imposed by the third column of electrode sites, and the fact that one of the factors limiting the success of the early neuron localisation algorithm was insufficient high resolution field data. Once the model is bootstrapped, there is of course the possibility of improving the reliability of 2-column fits by fixing certain model parameters to standard values derived from the 3-column model fits. An obvious candidate would be the tensor magnitude, since this should be invariant for a given cortical area, and thus should not vary significantly across recordings.

4.4.3 Spatial clustering for spike sorting

The field potential modeling results reported so far were on average spike templates obtained from the spike sorting algorithm (chapter 3). The original conception of the neuron localisation algorithm was as a new method of generating clusters for spike sorting. Rather than modeling pre-sorted average templates, individual spikes are fitted, and collectively they form clusters in 3D cortical space (Figure 4.2). The centroid of these clusters represent the estimated neuron locations. Like PCA, this is a form of dimension reduction, only here the collapsed dimensions are real cortical coordinates. Since only a single neuron can occupy any one location in the brain, cluster boundaries and isolation distances can be defined in microns, not arbitrary feature-related units.

In the context of spike sorting there are several advantages to this method of spatial clustering, including the fact that spikes need not be sorted in advance. Spikes from a single neuron may vary considerably in amplitude, especially during complex burst firing (Buzsaki et al. 1996). Spike attenuation produces elongated clusters in cross channel amplitude (or PC) feature space (Figure 4.6A), making partitioning of multiple clusters difficult, and thus a major source of spike misclassification (Harris et al. 2000). In comparison, clusters derived from the neuron location algorithm applied to the same set of spikes were more or less spherical (Figure 4.6B). Correlated spike amplitude attenuation, instead of being manifest in the spatial coordinates, was offset by a proportional variation in the intrinsic current parameter I_m , and accordingly the two were highly correlated ($r^2 = 0.93$). The field model thus implicitly 'decorrelates' the spike amplitude variability, analogous to cross-channel whitening procedures that are sometimes used to improve multichannel spike sorting (Emondi et al. 2004).

Another benefit of spatial clustering is the relative ease of tracking neural ensembles in response to electrode drift, or as in the example shown in Figure 4.6C, following a deliberate movement of the polytrode (see also Figure 1.14). The majority of neurons were easily identifiable before and after the movement, with a mean relative shift in neuron position of $\pm 9\mu\text{m}$. Across the whole neuron ensemble there was a small, $5\mu\text{m}$ downward trend in y location, suggesting that insufficient time was allowed for the tissue to settle before restarting the recording. When this vertical shift was accounted for, the average deviation in

post-movement neuron location reduced to only $3\mu\text{m}$. This result also demonstrates that the localisation algorithm gives internally consistent results, and represents an important test of the model's validity.

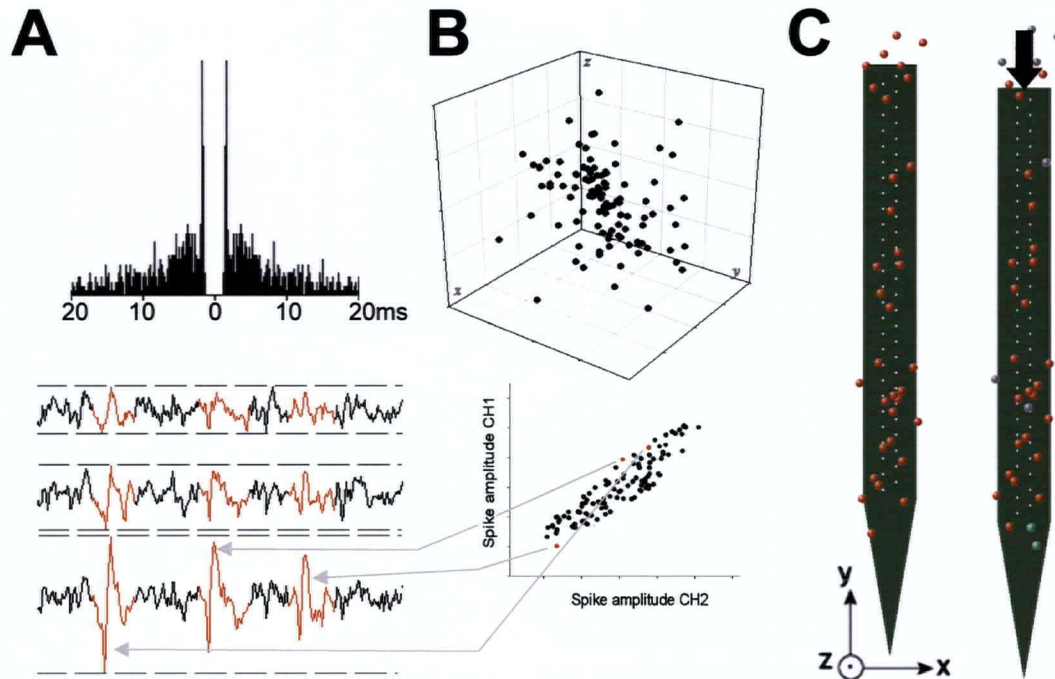


Figure 4.6 Neuron localisation for spike sorting.

(A) A burst firing neuron, with characteristic autocorrelogram, often has attenuating spike amplitudes across electrodes sites. Clusters in cross-channel amplitude space are therefore elongated ($n = 100$ spikes). (B) Clusters in 3D brain space were more spherical. (C) Predicted locations of 33 neurons relative to a $54\mu\text{m}$ ap2b polytrode (left). The polytrode was advanced $75\mu\text{m}$ along the y -axis (right), a non-integer multiple of the vertical electrode site spacing. The absolute and relative locations of most of the neurons (red spheres) remained virtually unchanged. Eight neurons (grey spheres) were not locatable following the movement. Two new active neurons appeared at the bottom (green spheres).

4.4.4 Neuron classification

Starting with the hypothesis that the polarised apical dendrites of cortical pyramidal cells gave rise to extended spike field potentials (open fields), while the radial dendrites of interneurons yield small, focal fields (closed fields), the existing library of sorted spike templates was searched for examples of both these extremes (Figure 4.7A). Qualitatively there was a clear dichotomy between spikes with fields appearing on only one or two recording sites, even on the highest density

polytrode after averaging, and the majority of spikes that were distributed across at least five electrode sites, irrespective of predicted distance from the polytrode (Figure 4.7B).

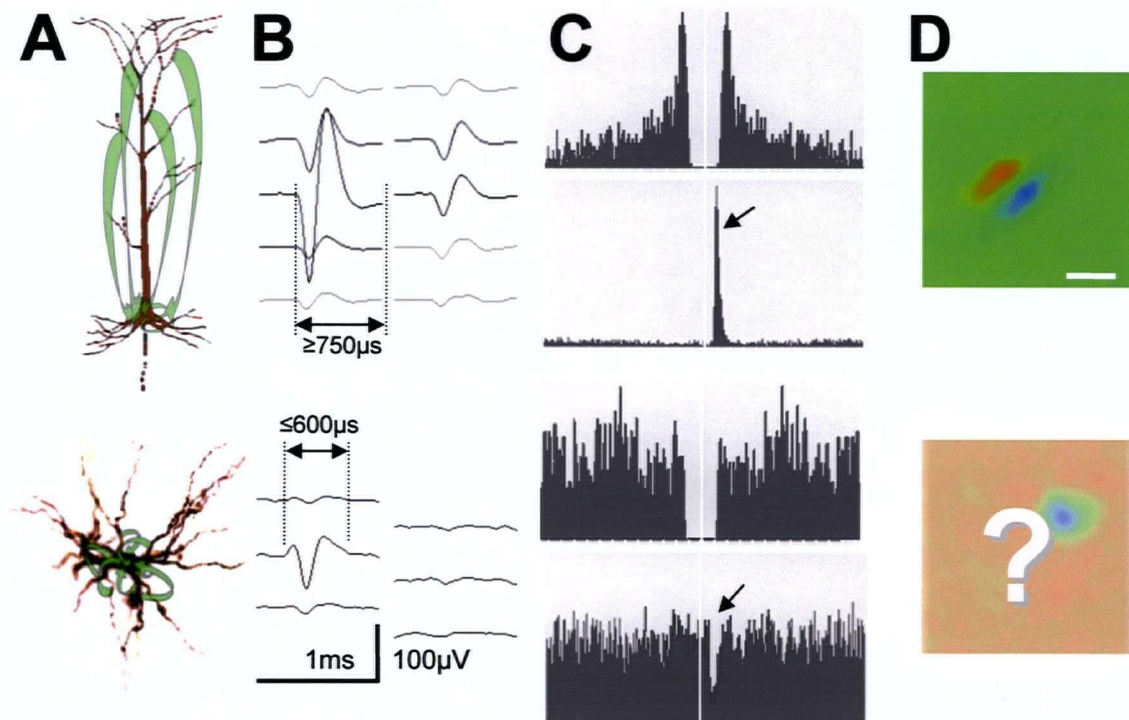


Figure 4.7 Classification of cell type.

Extracellularly recorded units may be classified into two groups, putative pyramidal neurons (*upper panels*) and interneurons (*lower panels*), on the basis of differences in field-potential model parameters, electrophysiological measures, and possibly RF properties. (A) Gross differences in cell morphology are hypothesised to produce (B) substantially different spike amplitude distributions across the polytrode. Spike waveform duration of the (putative) pyramidal neuron was longer than that of the (presumed) interneuron. (C) The auto-correlogram of the pyramidal neuron suggests a burst firing pattern that was absent in the interneuron. Pair-wise cross-correlograms for the pyramidal neuron revealed a monosynaptic excitatory connection; for the interneuron monosynaptic inhibition (*arrows*). Timebase = 1ms, 200µs bins. (D) The pyramidal neuron was an orientation-tuned simple cell (Scalebar = 2°). The RF of the interneuron was not mapped, however the majority of other FS neurons that were mapped had circularly symmetric RFs like that shown here.

Focal spikes invariably had FS waveforms (full-width $\leq 600\mu\text{s}$, $n = 9$), whereas broader fields were characterised by longer spike widths, akin to RS or IB neuron types (full-width $\geq 750\mu\text{s}$ $n = 124$). The association between field size, spike width

types (full-width $\geq 750\mu\text{s}$ $n = 124$). The association between field size, spike width and firing pattern was further evidenced by the profile of the auto-correlograms (ACGs). The latter were typically associated with burst like (Figure 4.7C) or RS, periodic ACGs (*ie.* correlograms with sidebands), whereas FS neurons had consistently flat ACGs. FS neurons also exhibited markedly higher spontaneous (4-13Hz) and maximal sustained firing rates (up to 600Hz) compared with their IB and RS counterparts (0-4Hz, and less than 100Hz, respectively). For FS neurons, the high firing rates were associated with a refractory period less than 1ms.

An analysis of a $\sim 7,000$ pairwise cross-correlograms (CCGs) from three separate recordings revealed evidence of functional monosynaptic excitatory and inhibitory connections (Figure 4.7C). Without exception, RS or IB neurons had excitatory CCGs ($n = 11$), whereas the single currently available example of an inhibitory CCG was for the FS neuron illustrated.

The sample population studied to date is too small to indicate a systematic pattern in the laminar distribution of the putative cell types, however the RF profiles of some neurons were mapped. Most FS neurons (4 of 6) exhibited weak or absent orientation tuning and circularly symmetric RF profiles. The two others were simple cells. Nearly all RS or IB neurons that had clear RF maps (18 of 26) were simple cells with narrow orientation tuning, and in some cases directionally tuned. The remainder were presumably complex cells as their fields were not able to be mapped using reverse correlation to m-sequence noise stimuli. The response of these cells to sparse noise stimuli has not yet been analysed.

4.5 Discussion

The neuron localisation algorithm presented in this chapter combines well-established electrostatic field models with simple coordinate transforms to predict extracellularly, for the first time, 3D neuron location to within a few microns; an accuracy comparable to that obtainable with intracellular histology. Inferring neuron location, or more specifically the location of the dominant neurogenic currents, on the basis of field potential measurements is a classic example of inverse modeling (Plonsey 1969). Here the nature of the volume conductor and the fields were known from prior work, the voltage distribution was measured with the polytrode, but the exact properties and origin of the source currents were

unknown. In theory a given field distribution could be produced by an infinite combination of superimposed sink/source distributions (from any number of neurons). Fortunately, the inclusion of a few physiologically realistic constraints (parameter bounds) and simplifying assumptions was sufficient to predict a single, most probable solution (*ie.* a definitive neuron location).

Broad classification of cell type into pyramidal and interneuron classes, previously the exclusive domain of intracellular recording, can now be predicted by the neuron localisation model on the basis of field potential spread and asymmetry, in addition to spike width (Bartho et al. 2004). This offers the prospect of studying the interactions between different neural types thought to play specific roles in mechanisms of visual RF tuning properties (Alonso and Martinez 1998; Hirsch et al. 2003). Until more neuron types and their functional characteristics have been described, it is too soon to be sure which measures, or combinations of measures, are diagnostic of cell type. Multivariate statistics or cluster analysis should help provide these answers (Nowak et al. 2003).

Using the neuron localisation algorithm as a method of spatial clustering enables polytrodes to track 'constellations' of active neurons whose identification is unperturbed by electrode drift (Figure 4.6). Combining spike shape information with the added constraint that the spatial relationships of the neurons do not change (Figure 1.14), provides a solid criterion for ensuring that the same neurons are being recorded over periods of hours, days, or even months. This may prove particularly useful for chronic studies of the neural correlates of perceptual and motor learning (eg. Gilbert et al. 2001; Paz et al. 2004), where long-term unambiguous identification of multiple neurons is essential.

4.5.1 Model validation and interpretation

In order to obtain indisputable verification of the location predictions, several attempts were made to image patch clamped or intracellularly recorded neurons that were also recorded extracellularly on the polytrode (appendix C). Even without imaging, paired intracellular-extracellular recordings are a technical *tour de force* and rarely attempted (Wehr et al. 1998; Henze et al. 2000), which was some consolation when none of the present experiments provided useful data.

Without such validation, how might the neuron localisation model, and hence the accuracy of predicted neuron locations, be tested? Most importantly, the model yielded reasonable location predictions, in spite of a solitary constraint ($z \geq 0$) imposed on the location parameters. The most distal location predictions were $148\mu\text{m}$ from the polytrode, comparable to the maximum recordable distance of CA1 pyramidal neurons in rat hippocampus using tetrodes (Henze et al. 2000). At the other extreme, neurons with highly focal fields, most likely small interneurons, were always estimated to be in the immediate vicinity of polytrode. The positional invariance of neural ensembles following movement of the polytrode is a good arbiter of the stability and internal consistency of the model. This result is non-trivial because the polytrode was advanced a non-integer multiple of the vertical electrode site spacing, in such a way that the field amplitude distributions across the polytrode array were markedly different (*ie.* not just the same distribution on different sites).

The layered, ordered anatomy of the cortex provides other opportunities to verify the localisation algorithm. Predicted cell densities and incidence of specific cell types should match known laminar distributions (Beaulieu and Colonnier 1985; Peters and Yilmaz 1993; Binzegger et al. 2004). For example, with a larger sample size from multiple penetrations the model should predict a high density of pyramidal neurons in layers II/III and V, the presence of excitatory non-pyramidal (spiny stellate) neurons in layer IV, and FS interneurons throughout all cortical layers. It would also be informative to extract and localise all active neurons (down to $\text{SNR} \approx 1$) at a single recording location, and observe whether the total number of estimated neurons increases as πr^2 (*ie.* the volume of potentially recordable neurons). For recordings made down the medial bank of the striate cortex, predicted neuron orientations (θ_z rotations) should vary with depth as the polytrode traverses the change in radial fibre direction around the cortical gyrus.

Further consideration should also be given to the rationale behind the model assumptions, in particular the evidence that neurons are only recorded in front of the polytrode. Drake and colleagues (1988) did a theoretical study of the influence of the polytrode on recorded field potentials, and showed that the presence of the insulating shank attenuates potentials behind the polytrode by almost 100%. This conclusion was recently confirmed using finite element modeling (Anderson et al. 2001). The assumption that the tissue conductivity is isotropic parallel with the

cortical layers (*ie.* $\sigma_x = \sigma_z$) is supported by the anatomical homogeneity of the cortex in these two dimensions. Vertically (aligned with the columns), one study reported local, layer-dependent changes in isotropy, specifically a discontinuity between the molecular and granular layers of the turtle cerebellum (Okada et al. 1994). The authors emphasised that such boundary conditions could exert an appreciable distortion on extracellular current distributions, and consequently quantitative neuron location predictions (see also Bedard et al. 2004). However, other than layer I, cat cerebral cortex does not have comparable layer-specific transitions in neuronal density (Beaulieu and Colonnier 1985), and thus is unlikely to possess such gross inhomogeneities.

With the exception of BPAPs, which are easily identifiable with template matching, the stationarity of spike-related source currents was assumed. Layer V pyramidal cells of the neocortex (Schiller et al. 1997) and CA1 pyramids in the hippocampus (Gasparini et al. 2004) are known to have two spike initiation zones. Most common are the low-threshold sodium spikes initiated in the initial segment of the axon, but a high-threshold calcium spike may be generated in the distal apical dendrites following high temporally coincident synaptic input. In spite of this possible confound, the low current density of calcium spikes, in addition to the 'closed field' effect of the dendritic tufts makes them unlikely to be recorded extracellularly with multiunit electrodes.

Another way to assess the validity of the model is to compare fitted parameters with known physiological values. The range of intrinsic current parameter values, on the order of a few nanoamps, is consistent with measured peak transmembrane currents in cortical neurons (Plonsey 1969; Malmivuo and Plonsey 1995). Although the σ_x and σ_z conductivity parameters were constrained to be equal, their combined value was free to vary. Average fitted values were $0.42 \pm 0.17 \text{ S.m}^{-1}$, closely approximating empirically derived values of 0.45 (Freygang and Landau 1955) and 0.29 – 0.55 S.m^{-1} (Hoeltzell and Dykes 1979).

Unfortunately, reliable data on adult cat visual cortex conductivity tensors does not exist. It is, however, possible to infer indirectly the conductivity tensors from diffusion weighted magnetic resonance imaging (DTI)[†]. The major anisotropy derived from high resolution 9.4T DTI in cat follows the vertical

[†] DTI measurements do not distinguish between intra- and extracellular compartments, however there is a strong linear relationship between ECS conductivity and the diffusion eigenvalues (Tuch et al. 2001).

(columnar) fibres of the cortical grey matter (Ronen et al. 2003). Quantitatively, the tensor magnitude along this axis, relative to the other two axes, is 1.38 ± 0.12 (Mathieu Ducros, unpublished data). Whereas some studies of other species find no evidence to indicate the cortex is anisotropic (Lehmenkuhler et al. 1993; Nicholson and Tao 1993; Mazel et al. 1998), others, particularly those of stratified brain structures, such as cat (Yedlin et al. 1974) and anuran (Nicholson and Freeman 1975; Rice et al. 1993) cerebellum, rat hippocampus (Mazel et al. 1998), and cat somatosensory cortex (Hoeltzell and Dykes 1979), report significant tissue anisotropies ‡. Where anisotropy has been reported, it is always greatest along the axis of the dominant axonal or dendritic projection (Nicholson and Sykova 1998), and ranges in magnitude from 1.2 ~ 2.9 fold depending on the species, anatomical heterogeneity, and the method used. Ratios greater than one would be expected to distort the field potentials so as to make them oblate along the axis parallel to the cortical pyramids (equations 4.3 and 4.6).

The prevailing assumption throughout this chapter has been that anisotropic field shapes can be equated with anisotropic cortical conductivity. These physiological data pose a dilemma for the current model that predicted the *y*-axis tensors were less than or equal to one (0.53–0.97). Even if the cat DTI data and conductivity measurements from other brain areas are ignored, there must be another explanation for the anisotropic (prolate) field shapes. The existence of neurons with almost spherical fields also puts into question this interpretation, as a homogeneous volume conductor was assumed. Of course it could be postulated that local regions are isotropic (*ie.* the tissue is inhomogeneous), or that spherical fields are from bitufted non-pyramidal or basket neurons orientated perpendicular to the cortical surface. In the latter case the distortion caused by the tissue anisotropy would be offset by the lateral orientation of the neuron, resulting in a spherical field distribution. But these explanations still fail to reconcile the prolate field potentials with the physiological conductivity data.

The existing model is not fundamentally flawed. As a purely descriptive model, the inclusion of semi-independent scalar parameters was both necessary and sufficient to describe the shape of the neural field potentials with great

‡ Many of these studies measured extracellular ion or small molecule diffusion, a quantity referred to as tortuosity. Since ECS conductivity is roughly proportional to the EC volume fraction / tortuosity² (Okada et al. 1994), on this basis if the tortuosity is isotropic, so too is the ECS conductivity.

accuracy. Rather, the problem here is one of interpretation, and suggests that some biophysical property other than conductivity is the cause of the prolate fields. How might this be resolved? The most direct approach would be to make four-electrode impedance measurements (Ivorra et al. 2003) with the polytrode and quantify the conductivity tensors at all cortical depths (as described in appendix A.3). These measurements could be used to fix the model conductivity parameters. If elaborations to allow for tissue inhomogeneities are sufficient to account for the observed field shapes, then the basic form of model can remain unchanged. Alternatively, if some other phenomenon is responsible for the prolate fields, then there will still be some heterogeneity in the field shapes even after compensating for known tissue anisotropies. In this case a likely explanation is that the dominant neurogenic currents are not point sources, but are in fact distributed across the membrane of the neuron, perhaps with the size of the dipole moment proportional to the size of the apical dendrite. As the majority of extracellularly recorded neurons are pyramidal neurons aligned with the cortical columns, this could have created the impression of an anisotropic conductivity tensor along the y axis, even if the tissue was more or less isotropic.

4.5.2 Model refinements

The inverse correlation between the intrinsic current and conductivity parameters suggest a degree of parameter redundancy. To improve the stability of the model it may be desirable to constrain one or both parameters to known physiological values or ranges, but leave the conductivity ratio free to vary.

Since the shape of the neural fields cannot be explained by an anisotropic tissue model, the simplest modification would be to reinstate the full definition of the dipole moment ($I_d d$ in equation 4.5), with d representing the dipole length, and constraining making $I_d = I_m$. This would have the same effect of inverting the y axis conductivity parameter, but with the semantic benefit of keeping the model biophysically realistic. It will probably also be necessary to fix the conductivity tensor to a suitable physiological value, say 1.4.

Although somewhat of a departure from the existing point source models, both could be reformulated as an equivalent cylinder model using methods similar to Rall (1962) and Humphrey (1968). Here the monopole component is

made redundant, and the dipole is replaced by a single concentrated sink representing the soma with a collinear distribution of sources corresponding to the apical dendrite. In this case the extracellular potential is given by:

$$\psi = \frac{1}{4\pi} \cdot \left\{ \sum_{i=1}^n I_i \Delta S r_i - I_s S_s r \right\}, r_i = \frac{1}{\sqrt{\sigma_x^2 x^2 + \sigma_y^2 (y - d_s i)^2 + \sigma_z^2 z^2}} \quad (\text{after Humphrey 1968}) \quad (4.11)$$

where I_s is the instantaneous current density over the soma, S_s is the somatic surface area, d_s is the soma diameter, I_i is the current density over the i th dendritic segment, of surface area ΔS , and n is the total number of segments. Values for S_s could be fixed to typical anatomically derived values for cortical neurons, and $\Delta S \approx S_s/6$, on the basis that each segment length is a multiple of d_s , and the diameter of the apical dendrite is typically $\sim 1/6^{\text{th}}$ the diameter of the soma. To further simply matters, the magnitude of I_i could be distributed evenly amongst the segments, with only the number of segments n free to vary. Observe that the conductivity parameters have been retained, as the (albeit limited) available evidence suggests that tissue anisotropies should be represented. According to this model small pyramidal neurons and stellate shaped neurons would have a few segments, thereby approximating a dipole, whereas large layer V pyramids might have up to two hundred segments. One of the appeals of this approach is that the model may be used to predict the length (or absence) of the apical dendrite, and hence estimate the extent of the neuron's synaptic inputs in the cortical layers akin to filling and reconstructing the cell intracellularly.

Finally, the two theoretical studies that concluded neural sources were not recordable behind the polytrode also predicted that the wide insulating shank distorts the extracellular current flow, accentuating the apparent spike amplitudes by as much as 60% adjacent to the centre of the shank (Drake et al. 1988; Anderson et al. 2001). It may thus be possible to further improve the location predictions by modeling the distortion of the field potentials caused by the presence of the polytrode shank, however there was no obvious evidence of this.

chapter 5

Prospective applications

5.1 A new experimental paradigm

The preceding chapters document the development of methods for large scale neuronal recording, including the design and testing of novel high density polytrodes; supporting hardware and software for high bandwidth continuous data acquisition; signal processing methods for improving waveform fidelity and neuron yield; spike detection and sorting algorithms adapted specifically for polytrodes; and finally the establishment of a biologically realistic field potential model for 3D neuron localisation and estimation of cell type.

Polytrodes have the capacity to monitor most, if not all, active neurons in a local region of brain extending the full length of the polytrode. Both the number of experimental animals and the time needed to adequately sample the neuronal population is thereby greatly reduced. Whereas tetrodes can record from multiple neurons in a local network, high density polytrodes enable studies of the functional connectivity and interactions of neuronal ensembles across whole cortical columns. With the means to determine the precise anatomical location of recorded neurons, and differentiate interneurons from pyramidal cells, polytrodes can be used to test predictions from modeling work in systems neuroscience that have not previously been open to scrutiny. Several such experiments of this nature are described later in this chapter.

One of the challenges of recording simultaneously from large numbers of sensory cortex neurons is devising appropriate stimuli. The usual approach of characterising a single neuron by holding most stimulus dimensions constant and optimal, varying only the stimulus parameter of interest, is not only inherently biased (Towe 1973), but infeasible when recording from an assortment of neurons with a potentially large range of 'optimal' tuning responses. One solution is to present all possible combinations of the relevant stimulus dimensions for a range of different stimuli. While time consuming, this has the added benefit of impartial characterisation of the neurons in the sample population, including those with

small spike amplitudes, highly focal field potentials (Humphrey and Corrie 1978; Humphrey 1979), or low spontaneous firing that might otherwise have been overlooked with sequential single unit recordings.

A recent commentary by Olshausen and Field (2005) argues convincingly that our incomplete understanding of sensory cortices is, in part, due to these sorts of biases in the design and execution of experiments, in addition to impoverished stimuli. In this context, the challenge of stimulating the large numbers of neurons recorded with polytrodes can be viewed as an opportunity. Time traditionally spent hunting for active, well-isolated units can now be used to more fully characterise the response properties of the neuronal ensemble. Since it is routinely possible to obtain stable recordings for periods of several hours (Figure 1.14), an individual population of neurons can be studied with a much broader range of multi-dimensional stimuli, noise stimuli, and, natural scene movies (Figure 5.1).

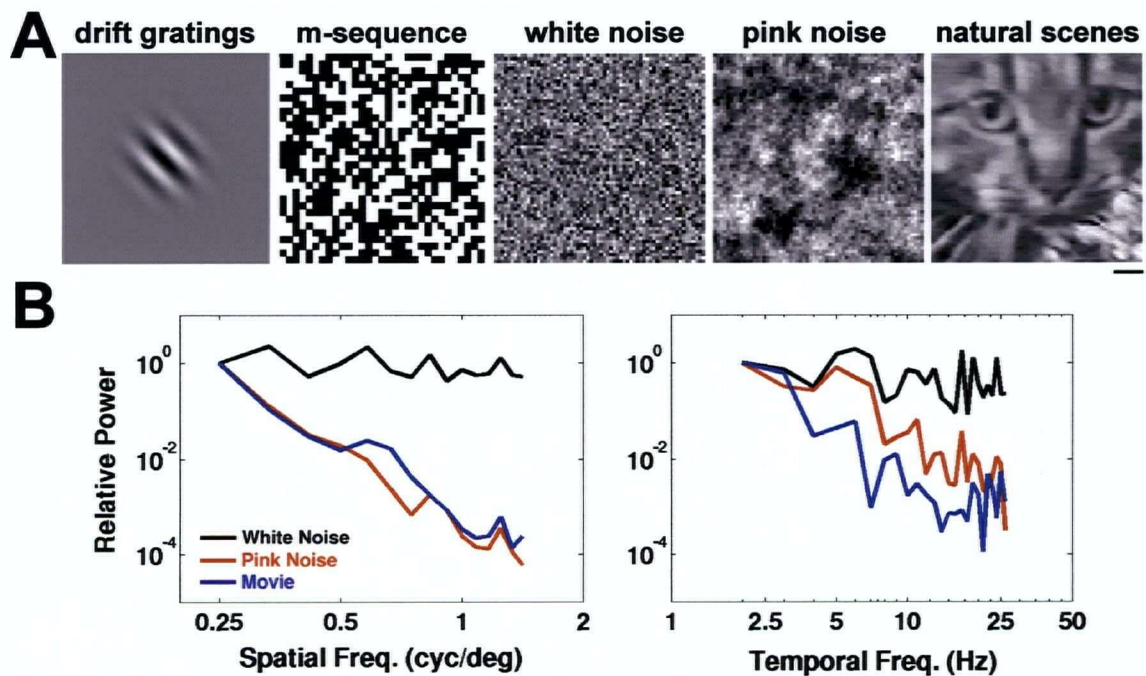


Figure 5.1 Visual stimuli.

(A) A subset of the stimuli presented at each electrode penetration, including 2D gabor drift gratings, binary m-sequence noise, white and pink spatiotemporal noise, and natural scene movies. Scalebar = 1° . (B) White noise is uncorrelated in both space and time, and therefore has a flat power spectrum. Pink noise and natural scenes share similar $1/f$ power spectra, however only the latter contain cohesive phase structure and higher-order correlations characteristic of real world images (Blanche et al. 2003 in collaboration with Nick Lesica, Stanley Lab, Harvard University.).

This exhaustive approach to visual neurophysiology has been embraced in the polytrode recordings made to date. Several days of continuously recorded spike and LFP waveform data (nearly half a terabyte) have been accumulated, containing thousands of neurons from multiple electrode tracks in six acute cat experiments. Each of the area 17 and 18 neurons in this database were stimulated with a rich gamut of visual stimuli (Table 5.1).

Table 5.1 Visual stimuli and purpose.

Stimulus	Purpose
drifting bars	orientation, direction tuning; discrete stimulus for assessing functional connectivity ¹
multidimensional sinusoidal drift gratings, 2D gabors	classical RF tuning properties; orientation, direction, spatial frequency, contrast, and velocity (temporal frequency) tuning; classify simple vs. complex ($f_1:f_0$) ²
rapidly changing static sinusoidal gratings	temporal evolution of orientation ³ and spatial frequency ⁴ tuning using reverse correlation
white noise, pink noise, m-sequence noise, sparse noise	characterise linear spatiotemporal RF (size, shape, position, ON/OFF sub-fields) by reverse correlation ⁵
discrete flashed spots, bars	laminar response latencies ⁶ (relative/absolute), EEG dependence of; rank order coding ⁷ ; study feed-forward shunting inhibition model of contrast invariant tuning ⁸
natural scene movies	study RF differences with higher-order stimuli; non-linear RF models; adaptation under natural viewing ⁹
spontaneous activity (no stimulus)	baseline for all stimuli, control for changes in arousal ¹⁰ ; functional connectivity
dual field, temporally lagged flashed stimuli/m-sequences	<i>in vivo</i> STDP studies ¹¹ , effects of conditioning on RF properties, neuron, laminar specificity; role of BPAPs ¹²

Key references: ¹(Aertsen and Gerstein 1985; Aertsen et al. 1989); ²(Skottun et al. 1991); ³(Ringach et al. 1997); ⁴(Bredfeldt and Ringach 2002); ⁵(Jones and Palmer 1987; Sutter 1991; DeAngelis et al. 1993a, 1993b; Kayser et al. 2003a); ⁶(Best et al. 1986; Saul and Humphrey 1990; Maunsell and Gibson 1992; Nowak et al. 1995; Gawne et al. 1996b); ⁷(Van Rullen and Thorpe 2001); ⁸(Ferster 1988; Ferster and Miller 2000; Miller et al. 2001); ⁹(Stanley et al. 1999; Ringach et al. 2002; Kayser et al. 2003b; Smyth et al. 2003); ¹⁰(Worgotter et al. 1998); ¹¹(Fregnac et al. 1988; Yao and Dan 2001; Fu et al. 2002); ¹²(Stuart and Sakmann 1994; Spruston et al. 1995; Magee and Johnston 1997).

To illustrate the potential utility of polytrodes and the explanatory value of amassing an extensive multiunit visual RF database, several unresolved questions

in visual neuroscience of personal interest – those that provided the original impetus for developing the polytrodes – will now be considered.

5.2 Cellular-scale organization of the visual cortex

The pioneering work of Hubel and Wiesel (1962; 1963) in visual cortex, and Mountcastle (1957) in somatosensory cortex, established the cortical column – a fascicle of cells, about 50 μ m in diameter, running from pia to white-matter – containing cells with similar response properties. They and others since then have shown that many receptive field (RF) properties (*e.g.* eye dominance, RF position, preferred orientation, direction of motion) obey the principle of columnar organization and are mapped in a systematic and orderly fashion across the brains of many species. In recent years, the study of cortical maps has been advanced greatly by the development of optical imaging. With this technique, neural activity is measured as minute changes in reflectance at specific wavelengths of light (Bonhoeffer and Grinvald 1996) in response to the presentation of a stimulus, such as a sine wave grating (Figure 5.2A). Results obtained in this way have revealed maps selective for stimulus orientation, ocular dominance and spatial frequency in cat area 17, which are overlaid in a highly systematic manner (Hubener et al. 1997). Modeling studies can explain the complex spatial interrelationships of these multidimensional maps in terms of continuity – that cells close together in the cortex have similar response properties, which in turn may be due to a requirement that axonal connection lengths be minimized, and completeness – the requirement that all combinations of stimulus properties in individual maps are represented with nearly equal frequency across the map (Swindale et al. 2000).

Recent findings from optical imaging notwithstanding, much of our understanding of the RF organisation of the striate cortex comes from multiple single electrode penetrations (Cynader et al. 1987; Swindale et al. 1987). The quality of these maps is limited by imprecision in determining the position and depth of the electrode, the approximate source location of neurons recorded with multiunit electrodes such as tetrodes (Maldonado et al. 1997), and eye drift or other response non-stationarities during prolonged serial measurements of multiple RF properties. Likewise, because of the finite sampling density with

multiple single-unit mapping studies, and limits in the spatial resolution and depth of optical imaging ($\sim 250\mu\text{m}$), it remains unclear whether the apparent high degree of order shown in the cortical maps is maintained at the level of single cells. If it is, then for which RF properties? Columnar organisation implies a degree of redundancy in the local neuronal population, however neighbouring neurons may reveal very different response properties when characterised in greater detail, particularly for higher-order or naturalistic stimuli (Olshausen and Field 2005). In other words, it is still an open question as to how the myriad RF properties are clustered (or anti-clustered) vertically within a cortical column, and how much redundancy (or conversely how much independence) there is in the information conveyed by adjacent neurons (Gawne et al. 1996a). Uncertainty about the magnitude of RF scatter, both locally (Maldonado and Gray 1996; Maldonado et al. 1997; Hetherington and Swindale 1999), and laterally across several cortical columns, is another unresolved issue that could be addressed using cellular-scale cortical maps generated with polytrodes.

A related, emerging field of visual neurophysiology focuses on the origin of RF tuning properties. Several elegant experiments have examined directly the cortical substrate of orientation tuning by recording simultaneously from visuotopically aligned centre-surround neurons in the lateral geniculate nucleus and simple cells in the visual cortex (Reid and Alonso 1995; Ferster et al. 1996; Alonso et al. 2001). Together these studies have ratified the original feed-forward model of Hubel and Weisel (1962) that the orientation tuning of layer 4 simple cells is due, in large part, to the specific spatial arrangement of the convergent, monosynaptic geniculate afferents. These findings do not preclude a role for local intracortical mechanisms in shaping orientation tuning, and indeed both cortical inhibition (Stillito 1992; Celebrini et al. 1993) and excitation (Ferster 1988) have been implicated in the formation and modulation of orientation tuning width. Studying the dynamics of orientation tuning in different cortical layers (Ringach et al. 1997), and the influence of layer 5/6 neurons on the orientation tuning of upper layer complex cells in primates (Allison et al. 1995), provides further constraints as to the likely mechanisms of orientation tuning. Similar methodologies have been applied to intracortical circuits for RF properties other than orientation selectivity. For example, two recent studies of the functional connectivity between simple and complex cells provide solid experimental support in favour of the hierarchical

model of cortical processing, whereby complex cell RFs derive from the convergent inputs of simple cells with similar orientation preference (Alonso and Martinez 1998; Martinez and Alonso 2001).

Since polytrodes are not limited to recording cell pairs (only single and paired recordings were made in the studies cited above) a rigorous analysis of the existing multiunit database has the potential to build significantly on the existing body of work. Specifically, if the hierarchical simple-complex RF model is robust, then both the substructure and timecourse of a complex cell's spatiotemporal RF should match the RF evolution of the population of functionally connected, visuotopically aligned simple cell afferents. This remains to be demonstrated. The same analyses can be applied to ensembles of functionally connected cells for orientation tuning, direction preference, or for that matter any other RF property of interest. Taking advantage of the cell typing algorithm, it will also be possible to investigate the association of different neuronal classes (McCormick et al. 1985; Azouz et al. 1997; Nowak et al. 2003) thought to play specific roles in mechanisms of RF tuning, such as the role of feed-forward inhibitory interneurons in contrast invariant orientation tuning (Martinez et al. 2002; Contreras and Palmer 2003; Hirsch et al. 2003). Furthermore, by recording from an appreciable proportion of the intracortical circuit, it should be possible to assess the relative contribution of the various neural components involved in a given RF tuning mechanism. Large scale ensemble recordings also provide valuable controls that are often lacking in existing studies. Collectively, these empirical studies will provide information essential for testing and elaborating biologically realistic models of RF tuning in early stage visual pathways (e.g. Troyer et al. 1998; Miller et al. 2001).

Finally, the emphasis is usually placed on increases in firing activity in response to a specific visual stimulus, but since a reliable non-response (or change in baseline activity) also conveys information, the absence of a response is an equally valid metric for generating RF tuning curves. Single unit recording from predominantly inactive neurons is infeasible, but adopting the experimental design described in the previous section (along with continuous data acquisition) makes such studies with alternative response measures eminently possible.

Many of these questions relating to detailed map structure and how RF tuning properties emerge from cortical circuits are impossible to address without substantial numbers of simultaneously recorded neurons of known anatomical

location, characterised in an unbiased fashion with a broad range of visual stimuli. Here the concept of ‘cortical micromapping’ with polytrodes is introduced (Figure 5.2). The central idea of cortical micromapping is to associate neuronal structure (relative and absolute laminar position, cell type, size) with function (RF properties, cortical maps, response latency, timing, functional connectivity, and so forth). This is illustrated for RF shape, cell type and cortical layer in Figure 5.2C, but extends to any combination of functional measures that can be derived from the stimulus ensemble (Table 5.1). Insertion perpendicular to the cortical gyri, either at random positions or into specific features of the cortical map (Figure 5.2A), enables whole cortical columns to be studied simultaneously. A trajectory down the medial bank of the lateral gyrus in cortical areas 17 and 18 (Figure 5.2B) can be used to build cortical micromaps of exceptional detail, in addition to studying the lateral network interactions of several cortical columns or adjacent hypercolumns.

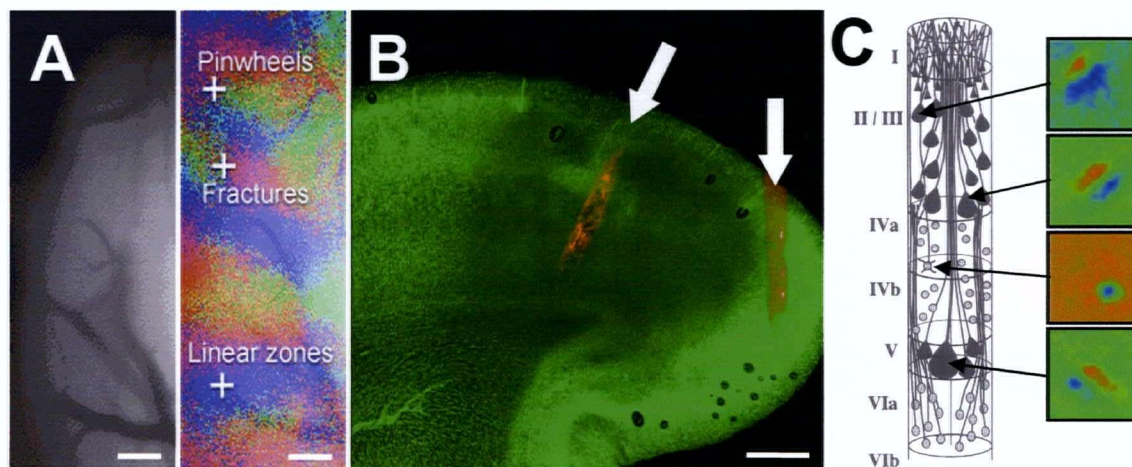


Figure 5.2 Cortical micromapping.

(A) Superior view of an exposed cortex showing the pial surface vasculature (*left panel*). Orientation map from optical imaging of cortical intrinsic signals in response to full-field drift gratings (*right panel*). The orientation vectors are represented by different colours. By co-registering this map with the vasculature it was possible to insert the polytrode into specific regions of interest in the orientation map (+). Scalebars = 1mm. (B) A coronal brain section stained for Nissl substance. Polytrodes were inserted perpendicularly to the surface of the brain for translaminar studies, and vertically down the medial bank of the lateral gyrus to record from neurons across columns within a layer. Scalebar = 500 μ m. (C) Cortical micromapping aims to link structure with function on the cellular scale across large populations of neurons, illustrated here for four neurons and their linear RF profiles.

5.3 Timing and temporal coding in the visual cortex

Models of visual processing generally adopt either a serial (hierarchical) or parallel organization. The hierarchical model, originally proposed by Hubel and Wiesel (1963), postulates that several successive processing steps should separate sensation from perception. In the cat, these successive steps are represented by cortical areas 17, 18 and 19, and by the simple, complex, and hypercomplex functional classes of neurons. With the discovery that distinct functional types of neurons exist in the retina and the lateral geniculate nucleus of the cat, the parallel model of organization gained recognition. X-, Y- and W-type retinal ganglion cells appeared to provide the major inputs to different cortical areas, and lesions or inactivation by cooling of area 17 (Sherk 1978) did not silence neurons in area 18, as would be expected if these areas operated in a hierarchical cascade (Stone et al. 1979). Parallel models place the emphasis on independent and simultaneous processing by modules specialized for different aspects of the visual stimulus.

These models are not mutually exclusive. The modern 'text-book' picture of the organization of the visual cortex is a hybrid of serial and parallel processing streams: parallel channels operating within a hierarchy of cortical areas (DeYoe and Van Essen 1988; Felleman and Van Essen 1991). A limitation of the hybrid model, however, is that it is based solely on anatomical observations. Functional studies (*e.g.* selective lesioning, or reversible inactivation) provide little support for such a model, at least not beyond the striate cortex (Bullier and Nowak 1995). Another way to test whether cortical areas process information in a parallel or serial fashion is to measure latencies with which neurons in different cortical layers or areas respond to visual stimulation. If information is indeed processed serially, then the latency of higher-order areas should be longer than those of the lower-order areas that drive them. In the primate, these experiments reveal that there is a very large range of latencies even within a cortical area, and that latencies overlap extensively across different areas (Raiguel et al. 1989; Maunsell and Gibson 1992; Nowak et al. 1995; Schmolesky et al. 1998). In one comprehensive study, extrastriate visual areas (V2, V3, MT, MST, and FEF) had neuronal responses with roughly the same average latencies, and all were within 6-9ms of activity in V1 (Schmolesky et al. 1998). Latency measurements in the cat visual system lead to similar findings. In the classically defined hierarchical

scheme, the order of activation should be areas 17, 18, 19, PMLS; however the order according to response latency is 18, 17, PMLS, 19 (Dinse and Kruger 1994). Within area 17, latency distributions show a large scatter (22-86 ms), reflecting the laminar position (Best et al. 1986), and the different pathways for X and Y inputs. Some of the response latency variance is probably also due to the non-stationarity of cortical state (*e.g.* changes in EEG) across stimulus repetitions (Arieli et al. 1996; Worgotter et al. 1998; Worgotter and Eysel 2000). Thus, while latency differences do not concur with the anatomical hierarchy across cortical areas, *within* a cortical area they suggest some degree of serial processing.

Clearly more work is needed to assess the relative importance of serial vs. parallel processing within a local cortical circuit. The advantages of making simultaneous neuronal recordings from all cortical layers are obvious, given the confounding effects of various non-stationarities, the dependency on laminar position, and the heterogeneity of cell types (both anatomical and functional). Response latency provides insights that are complementary to the analyses of functional connectivity and RF dynamics described in the previous section.

Recent theoretical work (Gautrais and Thorpe 1998; Van Rullen and Thorpe 2001) goes even further, speculating that response latency may be a form of rapid temporal coding used by the visual system. Using a go-no-go categorization task in which human subjects had to release a button when they detected an animal in a briefly flashed (20ms) natural photograph (previously unseen), Thorpe et al. (1996) showed that the processing required in such a task must be performed in less than 150ms (response time ~220ms). Monkey response times were even faster, only ~160ms (Fabre-Thorpe et al. 1998). Given that these relatively complex visual recognition tasks require processing in many, if not all, visual areas, how might they be performed so rapidly? In order to reach higher-order cortical areas, Thorpe and colleagues argue that the retinal information must go through at least ten processing stages, and mindful of the constraints of real neurons (neuronal integration times, axonal and synaptic delays, maximum firing rates) they proposed that this processing was essentially based on a feed-forward flow of information in which, in any given area, a neuron would rarely generate more than one spike. This model of information processing, dubbed 'rank order coding' (Gautrais and Thorpe 1998; Van Rullen and Thorpe 2001), embodies a precise temporal code in which only the first spike matters – visual information is

encoded in the rank order of discharge – a significant departure from traditional firing rate coding schemes.

Rank order coding, although intriguing and sound in principle, awaits experimental validation. It can be tested explicitly by examining the population response to repeated, briefly presented natural scene movie segments, looking for ordered spike sequences across the neural ensemble. In addition, standard information theoretic tools (Rolls et al. 1997; Panzeri et al. 1999) could be used to quantify the information available within the first few spikes fired in response to the stimulus. Although none of the recordings in the existing database were made simultaneously in multiple visual areas, it is still possible to look for evidence of temporal codes within cortical columns in area 17 and 18. The contribution of relative spike timing (including second and higher order statistics thereof) and the degree of temporal precision (Mainen and Sejnowski 1995; Berry et al. 1997; Borst and Theunissen 1999) in early stimulus representation could also be explored.

5.4 *in vivo* spike timing dependent plasticity

Spike-timing dependent plasticity (STDP) has emerged as one of the principal mechanisms of long term plastic change in the brain (Dan and Poo 2004). Consistent with Hebb's original hypothesis (Hebb 1949), STDP depends on the precise temporal relationship between pre- and postsynaptic spikes (Magee and Johnston 1997; Markram et al. 1997). In general, if the presynaptic neuron fires consistently before the post-synaptic neuron, the connection between the two neurons is potentiated (LTP). The converse is also true; postsynaptic spiking occurring repeatedly before the presynaptic neuron will weaken the synaptic strength (LTD). The temporal window for LTD vs. LTP is narrow, less than ± 50 ms (Bi and Poo 1998). Most of what is known about STDP has come from *in vitro* work, and only recently has there been *in vivo* evidence for STDP (Fregnac et al. 1988; Zhang et al. 1998; Yao and Dan 2001; Fu et al. 2002).

Detailed studies of BPAPs also derive largely from brain slice preparations (Stuart and Sakmann 1994; Spruston et al. 1995; Magee and Johnston 1997). Little attention has been given to *in vivo* analyses of BPAPs let alone their possible role in STDP. Polytrodes offer the unique capability of studying *in vivo* BPAPs, the brain state or stimulus conditions that evoke them, and their putative role as an

associative signal in STDP. Conditioning regimens similar to those of earlier studies will be used to evoke changes in functionally connected cell pairs. These conditioning stimuli comprise either spatially segregated grating stimuli (Fu et al. 2002) presented with a range of temporal delays (5–50ms, 5ms increments), or flashed gratings of different orientation [$\pm 30^\circ$ with respect to the optimal orientation for the column under study (Yao and Dan 2001)]. BPAPs have characteristic moving current sinks (Figure 1.8D) that can be identified using template matching (chapter 3). If BPAPs are indeed an essential mechanism of STDP, as suggested by the *in vitro* data (Magee and Johnston 1997), then a basic prediction is that the occurrence of BPAPs should increase during conditioning periods for any neuron that exhibits STDP modifications. Although BPAPs are a relatively rare observation, presumably because the polytrode must lie in very close proximity to the apical dendrite of an active neuron, it will still be possible to take advantage of the large, heterogeneous recording ensemble to extend the findings of previous studies. Recorded neurons that do not have an orientation preference or retinotopic location matched to the conditioning stimulus serve as important controls, as these (and other) neurons that do not respond with suitably correlated spike trains should not show evidence of STDP. With a sufficiently large sample population, it may also be possible to determine whether the magnitude of STDP change is binary in nature, or rather if the modulation of effective connectivity (or peak shift in RF tuning) is a function of either the proportion of precisely phase-lagged pre- and post-synaptic spikes, or the temporal precision of the spike trains in a given cell pair. It will also be interesting to see if neuron pairs without detectable baseline connectivities prior to conditioning can become functionally connected following induction of STDP. A thorough re-analysis of the repetitive, two-second natural scene movie sequences, focusing on cell pairs that happen to respond to stimulus features with consistent, appropriately phase-lagged spike trains, will provide further data for investigating the properties of *in vivo* STDP under more natural conditions (Froemke and Dan 2002).

5.5 Concluding remarks

Although the experiments described in this chapter come from somewhat disparate lines of research, they all capitalise on the unique recording capabilities of polytrodes and the additional information afforded by neuron localisation and classification. In the longer term the intention is to pool the results from the suite of diverse visual stimuli. Since the data are derived from the same population of neurons, the hope is that a fuller, more unified understanding will be gained of the organisational principals, coding schemes, and mechanisms of adaptation and plasticity that underpin early stage vision, beyond that possible by merely synthesising the findings from independent experiments made in separate laboratories. Whatever the outcome, the prospects for large scale neuronal recording with polytrodes promise to be as exciting as they are numerous.

appendix A

Impedance measurements

A.1 Meter design and calibration

Site impedances were calculated by measuring the relative amplitude and phase of a 1kHz, 10mV_{rms} sinusoidal wave applied sequentially to each recording site. The circuit (Figure A.1) exploits voltage division of an AC signal across a reference resistor in series with a complex impedance to ground to determine the electrode site impedances:

$$\bar{v} = \left(\frac{\bar{Z}}{\bar{Z} + R_{ref}} \right) \bar{v}_s \quad (\text{A.1})$$

where $\bar{Z} = |Z| \angle \phi_z$ is the amplitude and phase of the site impedance

$\bar{v} = |v| \angle \phi_v$ is the amplitude and phase of the signal measured across the site

$\bar{v}_s = |v_s| \angle 0^\circ$ is the sinusoidal source signal

R_{ref} is a 1M Ω reference resistor, roughly equal to the site impedance magnitude, making the voltage divider maximally sensitive to site impedance changes.

Solving for $|Z|$ and ϕ_z gives:

$$|Z| = \frac{|v| R_{ref}}{\sqrt{(|v_s| - |v| \cos \phi_v)^2 + (|v| \sin \phi_v)^2}} \quad (\text{A.2})$$

$$\phi_z = \phi_v + \arctan \left(\frac{|v| \sin \phi_v}{|v_s| - |v| \cos \phi_v} \right) \quad (\text{A.3})$$

Signal amplitudes and phases were obtained from fast Fourier transforms of the test signals (Press et al. 1994). Equation A.2 provides the impedance magnitude. The nature of the impedance (resistive or capacitive) was given by the phase angle from equation A.3. To ensure sites were not damaged, extremely low test currents (< 5nA_{rms}) were applied. Faraday shielding and low-noise op-amps yielded a measurement precision of $\pm 10\text{k}\Omega$. The meter was calibrated at

frequencies between 5Hz and 20kHz using known impedances created with pairs of resistors and capacitors (Figure A.2A). The circuit was linear across this range and showed no evidence of stray capacitances that could invalidate measurements (Figure A.2B).

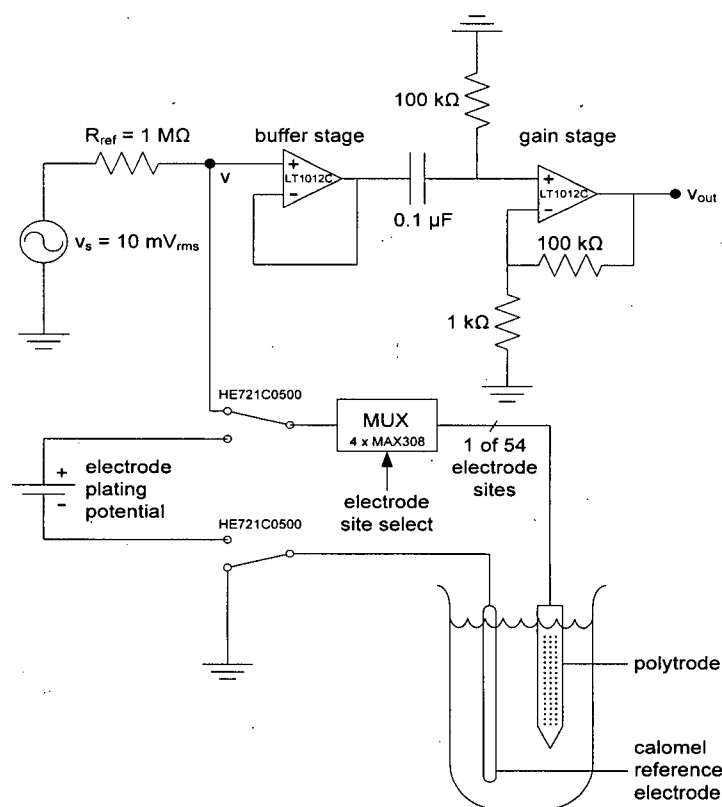


Figure A.1 Impedance meter/electroplater circuit diagram.

Dual-purpose circuit for automated site impedance testing and electrode plating. TTL-level control lines on the DT3010 card switched the multiplexers and selected the mode of operation via two relays (HE721C0500). A standard laboratory function generator, voltage divided to give a 10mV_{rms} AC test signal (V_s), was fed to individual electrode sites via a bank of analog multiplexers (MAX308). The resulting signal was buffered (LT1012C) and passed to a second op-amp (OP37) where it was amplified 100x. The output signal (V_{out}) was digitised, along with the input signal, by one of the DT3010 acquisition cards used for multiunit recording. PC-based software used these signals to calculate the site impedances according to equations (A.2) and (A.3) in the text.

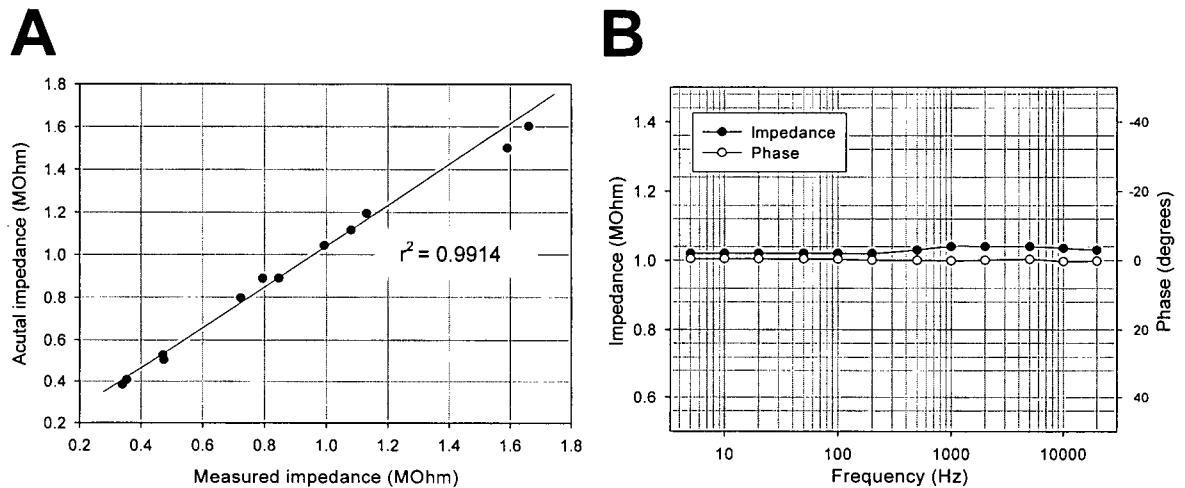


Figure A.2 Impedance meter calibration.

(A) Impedance measurements at 1kHz for various known impedances similar to those of the electrode sites. Minor deviations from the unitary line were within the tolerance specifications of the test capacitors and resistors. (B) Measurements of a test resistor (1.02M Ω) gave accurate readings (—●—) across the test range (5Hz–20kHz). The constant phase angle (—○—) of zero indicates a lack of any parasitic capacitances in the circuit.

The device also incorporates an electroplating mode for electrodeposition of platinum black or gold, a technique useful for lowering and matching site impedances by increasing microscopic surface area. By switching rapidly and automatically between impedance and electroplating modes it was possible to titrate the electrodeposition to achieve a desired impedance (data not shown).

A.2 Site impedance spectroscopy

The impedance meter was used to determine the impedance properties of a typical polytrode recording site across the biologically relevant frequency ranges for spike and LFP recordings (Figure A.3). The high impedance at low frequencies (Figure A.3A) reflects the capacitive nature of the electrode sites and interconnect dielectrics, which necessitates gigaohm input impedance pre-amplifiers for obtaining low-noise recordings of LFP and EEG signals (0.1 ~ 100Hz bandpass).

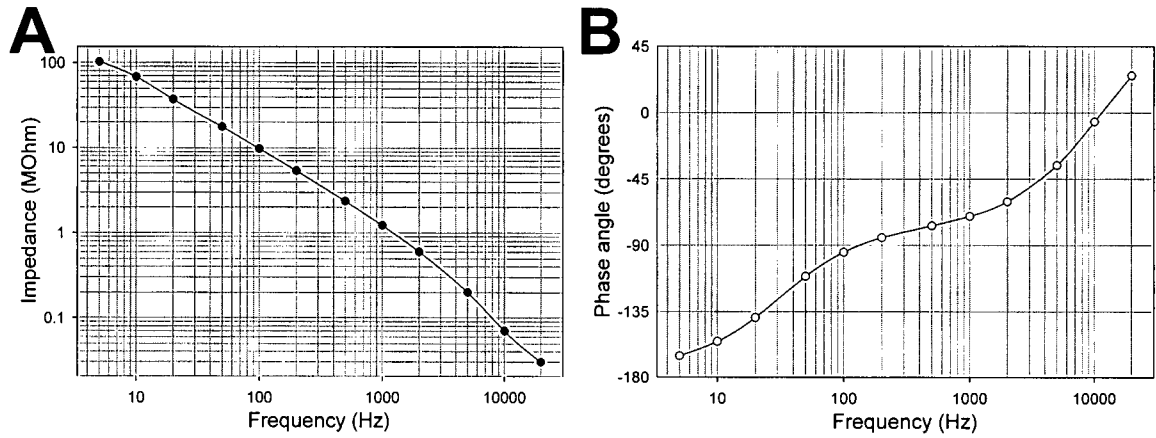


Figure A.3 Impedance spectroscopy of polytrode recording sites.

Across the biologically relevant frequency range, the site impedance magnitude (—●—) decreased, and the phase angle (—○—) increased, as a function of signal frequency.

A.3 Brain impedance spectroscopy

Efforts were made to make *in vivo* brain impedance measurements to determine whether the ECS acts as a complex impedance (differentially filtering spike signals in a frequency dependent manner), as suggested by Bedard and colleagues (2004), or rather if it is a purely resistive medium (signal decay over distance the same for all frequencies). These measurements could also evaluate the isotropy of the ECS conductivity. The two-electrode voltage division circuit used to test polytrode site impedances was too insensitive to measure the impedance characteristics of brain tissue four orders of magnitude smaller than the actual sites. In any case, measurements made in this way would be corrupted by the unknown series polarisation resistance at the electrode-tissue interface (Shalit and Mahler 1966; Shalit and Mahler 1967). A better approach is to use a four-electrode arrangement that annuls the effects of the electrode-electrolyte interface or any mismatch in site impedances. Silicon electrodes with precisely defined site impedances and spacing are well suited for such measurements, capable of a sensitivity better than a 1 Ω .cm (Ivorra et al. 2003). The brain resistivity ρ_f and hence impedance, for a given spatial dimension and frequency, is given by the following relation:

$$\rho_f = \frac{1}{4\pi} \frac{Va(a+b)}{bI} \quad (\text{A.4})$$

where a and b are the inner and outer co-linear electrode site separations, respectively, V is the measured voltage drop across the inner electrodes, and I is current applied to the outer electrodes [after Suesserman and Spelman (1993)].

Control measurements were made in 0.9% NaCl. The measured specific resistivity was $82\Omega\cdot\text{cm}$, very close to the accepted value of $77\Omega\cdot\text{cm}$ for low concentration saline solutions measured at 298K (Parsons 1959). At the time of writing, complex *in vivo* impedance measurements had not yet been made, however changes in electrode site impedance at 1KHz were monitored over time following polytrode insertion and advancement (Figure A.4).

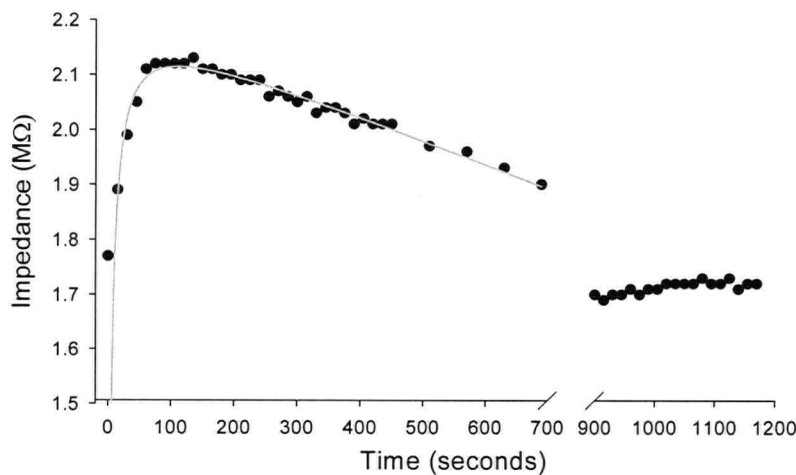


Figure A.4 Site impedance changes following polytrode insertion.

Post-penetration the electrode site impedance was approximately double that measured in saline. Over the next 10 minutes the impedance dropped steadily before stabilising at $\sim 1.7\text{M}\Omega$, 15 minutes (900s) after the initial insertion.

The impedance profile gives an indication of the time course of electrical coupling at the electrode site-tissue interface. It is intriguing (and of practical interest) that this profile qualitatively matches the observed recovery of spike amplitudes following initial insertion and movement of the polytrode. At the time these data were collected it was not possible to record spikes and make site impedance measurements at the same time, so impedances could not be directly correlated with changes in spike amplitude. Concurrent measurements in future experiments should provide a quantitative criteria for deciding how long post-insertion to begin recordings that guarantee stable, consistent spike amplitudes.

Even without direct *in vivo* impedance measurements it is still possible to infer the existence of significant tissue capacitance by studying the frequency content of spike waveforms over distance. Analysis of spike waveform shape (Figure A.5A, B) suggested that the spike waveforms on more distal sites were simply scaled versions of the waveform at i_{max} . A more quantitative measure of the spectral similarity across sites is the spectral coherence:

$$C_{xy}(\omega) = \frac{|\mathfrak{I}_{xy}(\omega)|^2}{\mathfrak{I}_{xx}(\omega)\mathfrak{I}_{yy}(\omega)} \quad (\text{A.5})$$

where \mathfrak{I} is the magnitude of the Fourier component at frequency ω . This equation returns a coefficient between 0 and 1 that indicates the phase independent correlation between two signals $x(t)$ and $y(t)$ at ω . The spectral coherence was high across all sites (Figure A.5C), suggesting that the brain tissue is ohmic over the physiologically relevant frequency range for spike waveforms ($\leq 10\text{kHz}$). Similar results were obtained for all cells examined ($n = 15$), the only exception being fields that were clearly dipoles (with inverted waveforms).

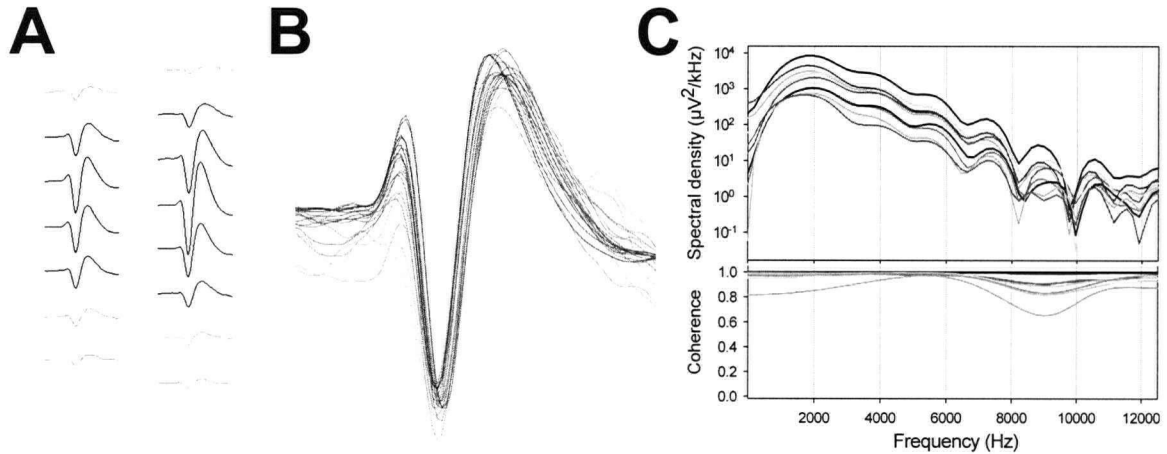


Figure A.5 The brain is a resistive medium.

(A) Average spike waveforms for a typical neuron had qualitatively similar shapes across recording sites. (B) Amplitude-normalised (with respect to i_{max}) waveforms confirmed the similarity of spike shapes. (C) Accordingly, spectral density profiles were similar across sites, differing only in magnitude*. The similar spectral content is reflected in the high spectral coherence between the waveform at i_{max} and the other sites. In (B) and (C) the plot intensity is inversely proportional to the Euclidian distance from the site to the neuron, as determined by the 3D neuron localisation model. * ripple at high frequencies is an artefact of the Fourier transform due to small discontinuities at the edge of the spike templates, magnified here by the logarithmic y-axis scale; the slight dip in spectral coherence between 7 and 11 kHz is probably an artefact of these edge effects.

appendix B

Instrument calibrations

B.1 Electrophysiology system

Further to the characterisation of the polytrodes described in chapter 1, an extensive array of tests and calibrations was made of all components of the electrophysiology apparatus (Figure 1.5). The total system noise with grounded headstage inputs was $\sim 15\mu\text{V}_{\text{pp}}$ (compared with $20\text{--}30\mu\text{V}_{\text{pp}}$ for a polytrode in saline). Once nulled, ADC voltage offsets deviated from zero by less than ± 1 bit over periods of months to years. The electronic noise of the DT3010 acquisition cards with shorted inputs was negligible (Figure B.2A). Shorts and signal coupling between analog channels were excluded (a short was discovered in the filter stage of the FA-I-64 amplifier, producing identical signals on two channels, but this was rectified by removing a stray solder bridge across two surface mount resistors). The software controlled channel selector (MUX-80) did not introduce additional noise, voltage bias or signal attenuation, and unlike previously used rotatory switches, did not cause saturating transient artefacts during channel changes. Careful isolation of digital and analog return lines meant there was no crosstalk from stimulus-related digital inputs apparent on any of the analog channels. Ground loops and 50Hz supply noise, especially problematic for EEG/LFP recordings, was kept below $10\mu\text{V}_{\text{pp}}$ by using a single 'star grounding' arrangement for the cranial reference electrode, stereotaxic frame, and micromanipulator. Faraday shielding (fine copper wire mesh) around the electrode interface board and headstage preamplifiers was needed to remove raster noise emanating from the display monitor.

The net analog gain of 40,000 ($1 \times$ headstage, $5000 \times$ amplifier, $8 \times$ internal ADC gain) was found to be accurate to within 0.03%, or less than ± 1 bit, for all channels. The bandpass characteristics of the analog inputs were linear across the factory-set ranges, and the 5-pole RC filters provided 40dB/decade of signal attenuation outside the specified corner frequencies (Figure B.1A). Since the

electrode sites did not confer any filtering *per se*, it was the amplifiers that dictated the frequency content of extracellularly recorded signals.

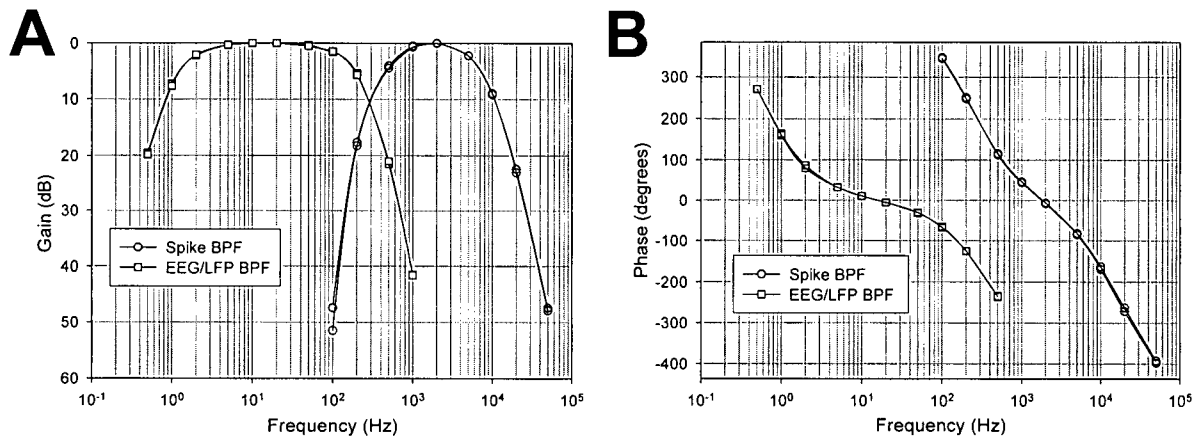


Figure B.1 Filter properties of polytrode sites and amplifiers

Bode plots for the polytrode recording sites (*open symbols*) and the recording system (*filled symbols*) were virtually identical, for both LFP and spike frequency bands. The amplifier filters, not the electrode sites, thus determined the overall filtering of recorded signals.

The phase plots of Figure B.1B deserve special consideration. Signal dispersion due to filter induced phase shifts within either of the frequency bands was, proportionally speaking, insignificant. However, if the exact timing relationship is sought *between* frequency bands, for example spike time with respect to the phase of a 3Hz delta wave, then these differential phase shifts are not insignificant and should be corrected for (*ie.* by subtracting ~69ms from the spike time, in this example).

The power spectral density (PSD) of the spikes and noise were also studied (Figure B.2). Electronic noise, measured with grounded ADC inputs, was broadband but had negligible power. The frequency dependence of the *in vivo* noise, of background neural activity without spikes, was essentially the same as the amplifier BPFs, with energy proportional to $\sim f^{-2}$ beyond a peak of 1kHz. The average spike spectrum derived from 148 randomly selected neurons had a broad maximum at 1.5kHz, and decayed monotonically at a rate proportional to f^{-4} . RS neurons (spike duration $\geq 750\mu\text{s}$) differed markedly from FS neurons (spike duration $\leq 500\mu\text{s}$) in the frequency content of their waveforms (Figure B.2B). Accordingly, RS neurons had peak spectral energy at 1kHz, with a fall-off of f^{-3} ,

whereas FS neurons had a broader peak at 4kHz and a shallower frequency dependence. Similar PSD plots were obtained for 7 other RS and 4 other FS neurons that were examined individually. These findings are compatible with those from a similar study of rat somatosensory neurons (Fee et al. 1996b), however the FS cells in cat cortex appear to be considerably faster than those reported for rat. Finally, the selection of 6kHz corner frequencies for the spike amplifiers appears to be ideal, providing negligible filtering of RS waveforms and minimal attenuation of the high frequency signal components of FS neurons.

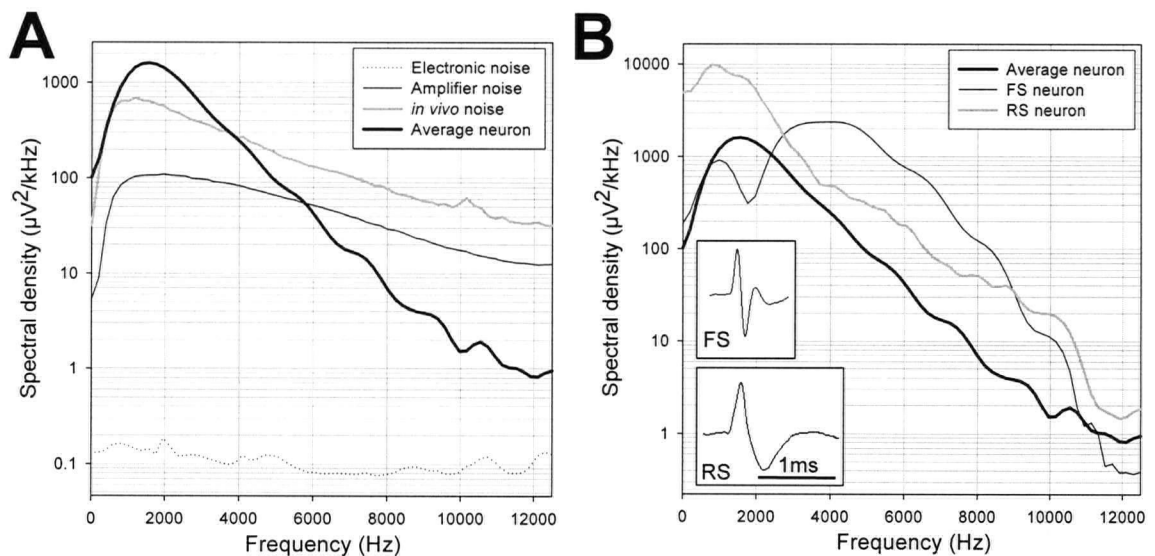


Figure B.2 PSD of spikes, neural noise and the recording system.

(A) Spectral analysis from DC up to the Nyquist frequency for a 25kHz sample rate (see text for details). (B) A RS and FS neuron with similar amplitudes (*insets*) have different frequency spectra due to differences in waveform shape. The average spike spectrum ($n = 148$) is dominated by the RS spectrum since FS neurons were much less prevalent.

B.2 Visual stimulus displays

In visual neurophysiology, as in visual psychophysics, the characteristics of the stimulus display monitor can influence experimental results. Whether or not a given feature or limitation of the display is relevant depends on the experimental question under study. Studies of contrast tuning and adaptation require a linear luminance profile, whereas a wide dynamic range (10 or more bits, Haberich and Lingelbach 1980) and low persistence phosphor (Di Lollo et al. 1997) are desirable

for studying perceptual thresholds and their neural correlates. Since many of the present experiments deal with analyses of visual response latencies, temporal receptive field profiles, and precise spike timing between neurons on the order of milliseconds, the refresh rate and potential for flicker-induced modulation of responses were of particular concern.

All experiments reported in this thesis used a high quality Sony Trinitron 200sf monitor for stimulus display. Prior to these experiments a photodiode was used to measure the phosphor dynamics and raster latency, and a precision photometer (Minolta LS-100) was used to calibrate the screen luminance. On-screen controls (brightness, contrast, colour temperature, etc.) were fixed at values that gave a qualitatively optimal display (close to factory settings). Calibration and measurement results are summarised in Figure B.3. The luminance profile was typical for a cathode ray tube monitor, able to be linearised across the full dynamic range with software-based gamma correction. Display uniformity was excellent (± 0.33 cd.m² stdev). The monitor could generate a 100Hz refresh rate, with a scan raster latency proportional to the vertical screen position (making it a simple matter to correct for this delay in physiological latency studies, if necessary). Phosphor 'on' responses were virtually instantaneous, and 'off' responses followed a double exponential decay ($\tau_{fast} = 1.2\text{ms}$, $\tau_{slow} = 4.9\text{ms}$).

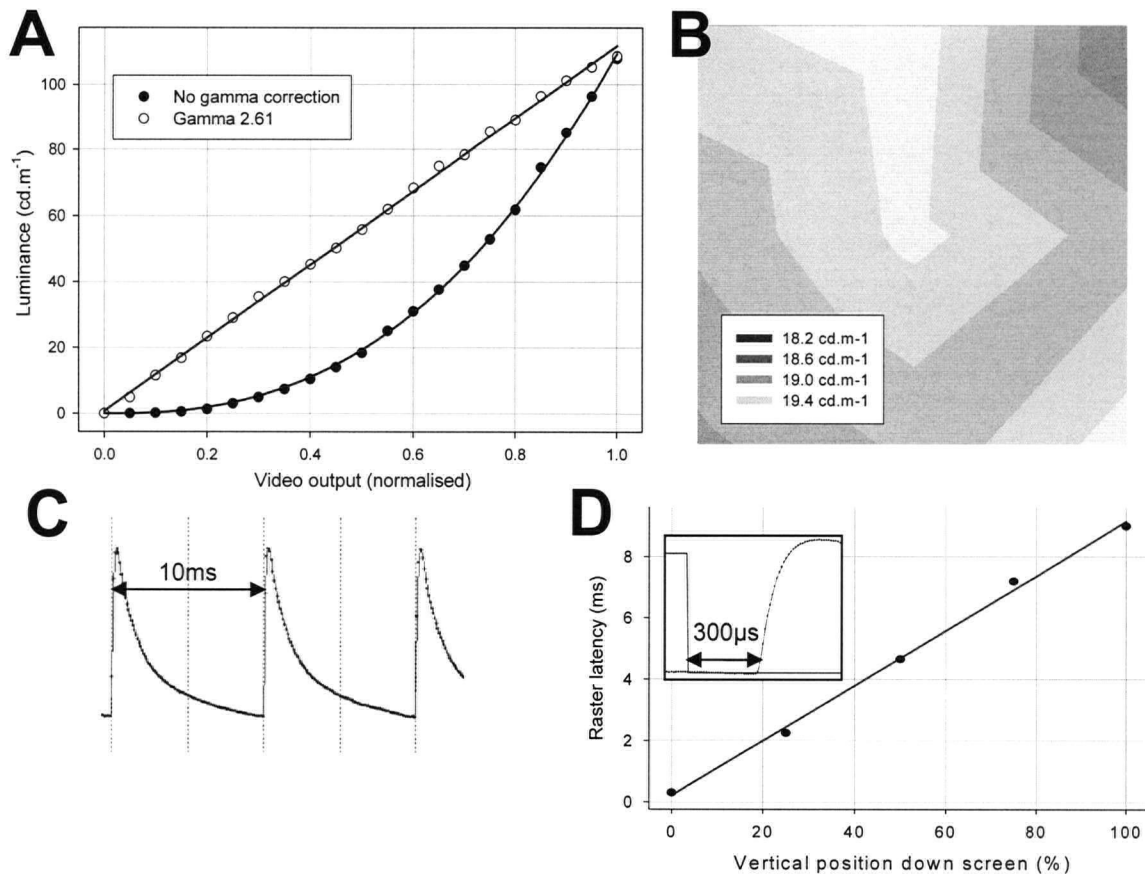


Figure B.3 Sony 200sf monitor calibration.

Calibration results for the visual stimulus display. (A) Luminance profiles with (o) and without (●) software gamma correction, $\gamma=2.61$ yielding a linear display. (B) Screen uniformity at ~20% output was based on a 3x3 array of spot measurements made at regular intervals across the screen. (C) Phosphor dynamics for two frames at a 100Hz refresh rate (up is increased luminance). (D) Most of the frame interval was evenly distributed across the vertical scan raster. A fixed 'front-porch' of 300 μs , the delay between the frame bit changing and the actual onset of luminance in the top-left hand corner of the screen, is evident (*inset*).

At screen refresh rates below 75Hz, screen flicker is perceptually noticeable. Retinal, geniculate, and even cortical neurons are known to phase-lock at these frequencies (Wollman and Palmer 1995), an unavoidable consequence of using conventional raster-based display monitors for visual stimuli. In spite of the relatively high refresh rate of the Sony monitor, an appreciable number of cortical neurons exhibited phase locking to the screen refresh at 100Hz (Figure B.4). This was somewhat surprising as until very recently (after these observations were made) it had not been reported in the literature (Williams et al. 2004). Observe

that these neurons were not geniculate afferents, since many were recorded in the infragranular cortical layers, and some had simple-cell RFs.

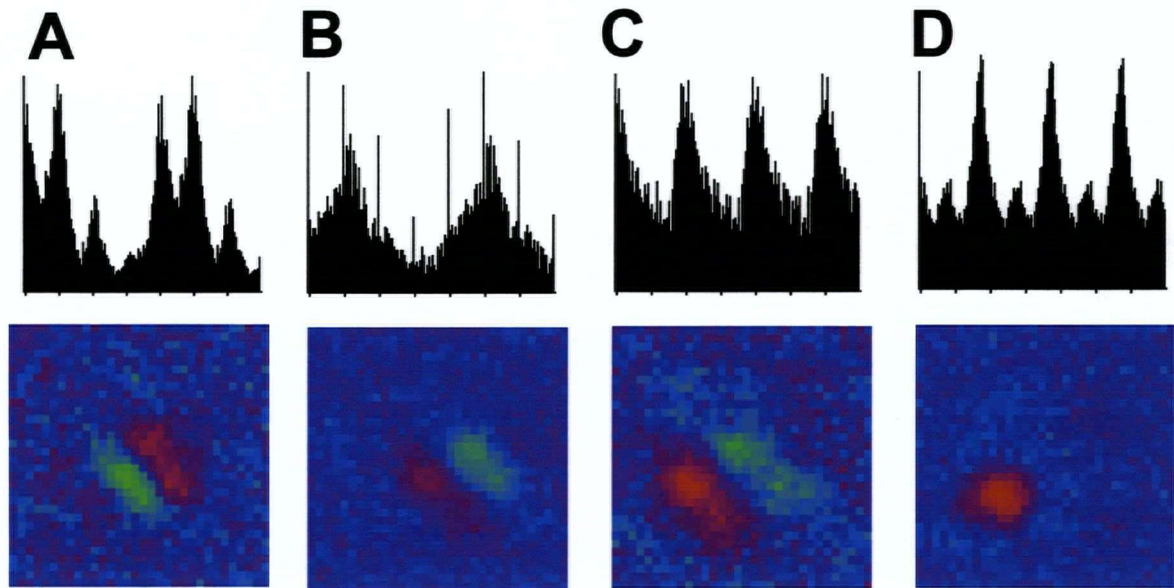


Figure B.4 Phase locking of cortical neurons to the 100Hz screen refresh.

Cortical neurons revealed varying degrees of entrainment that was independent of the visual stimulus, but always phase-locked to the 100Hz vertical refresh. For each of the four examples shown here, the *upper panels* show the peri-stimulus time histograms (10ms timebase, $t_0 = \text{Vsync pulse}$), and the *lower panels* show the spatial receptive fields ($12 \times 12^\circ$). (A) and (B) simple cell responses to an m-sequence stimulus with a 40ms inter-stimulus interval (ISI). Superimposed on the 40ms ISI modulation is clear evidence of entrainment at 10ms (100Hz). (C) and (D) two other neurons driven by m-sequence stimulation with a 20ms ISI. The simple cell was modulated by the ISI, but not the screen refresh, while the neuron with the circular receptive field was modulated by both.

My specific interest in precise neuron timing, and the fact that single-unit cross-correlograms would be irrevocably contaminated by the 100Hz refresh modulation, provided the impetus to invest in a display system capable of higher refresh rates. Of the few suitable commercially available systems, a Radeon 9800 graphics card (ATI, Montreal, QC) and a 200Hz refresh display (Iiyama Pro 454, Nagano, Japan) were chosen. Similar tests and calibrations were run on this monitor, and a comparison of both monitors is given in Table B.1. To date this display has not yet been used for experiments, although it is unlikely that cortical neurons will be modulated by flicker at 200Hz. If they are, this would be a surprising and interesting physiological finding in itself.

Table B.1 Comparison of stimulus displays.

Specification	Sony Trinitron 200sf	Iiyama Pro 454
Diagonal screen size (dot pitch)	44cm (0.25mm)	50cm (0.25mm)
Min/max luminance	0.03 / 108 cd.m ²	0.001 / 119 cd.m²
Gamma correction / linearity	2.61 / $r^2=0.999$	2.52 / $r^2=0.999$
Max refresh rate @ 800x600 resolution	120Hz	200Hz
Phosphor dynamics (rise / fall) †	≤ 200μs / 7ms	≤ 200μs / 3ms
Vertical raster period @ max refresh ‡	0.3 + % v_{posn} × 8.7 ms	0.45 + %v_{posn} × 4.35 ms
Display uniformity @ 20% luminance	±0.33 cd.m² (stdev)	±0.69 cd.m ² (stdev)

Parameters relevant to visual neurophysiology for the two display monitors. The **bold** text highlights the monitor with the better performance. Both monitors functioned according to specifications, but in all aspects but one (uniformity) the Iiyama monitor was superior to the Sony. † rise-time to peak luminance, fall-time to 5% of peak luminance. ‡ used to calculate the actual stimulus onset (with respect to the vertical sync video signal) as a function of the vertical position of the stimulus on the screen, v_{posn} .

appendix C

Validation of neuron localisation

The neuron localisation model (chapter 4) has a solid theoretical basis, is well constrained by the multisite field potential data, and makes plausible predictions of neuron location that shift concordantly with movements of the polytrode. Neuron classification measures explored thus far are consistent with those from intracellular studies. Nevertheless, the potential utility of this method warrants some form of independent validation, ideally from simultaneous imaging and recording of active neurons around the polytrode. This could corroborate the most basic prediction of the algorithm: does the 3D location and orientation of recorded neurons match those imaged *in vivo*? If not, where does the algorithm fail (neurons very near or far from the polytrode, or specific cell types or morphologies)? Are the model assumptions reasonable? Specifically, are neurons only recorded in front of the polytrode? Is the tissue conductivity, and therefore signal decay, isotropic for the two axes co-planar with cortical layers (*ie.* in front of and across the polytrode shank)? Can we assume a single, stationary dipole for each neuron at the axon hillock/soma, or do polytrodes sometimes record potentials from spiking dendrites or axons? Does the cortical layer or local cell density (assumed to be homogenous and isotropic on the millimetric scale) influence neuron location predictions? Do the phase-lagged dipoles moving up the polytrode shank (Figure 1.8) always indicate a BPAP? Finally, does the distribution of field potential spreads unambiguously distinguish pyramidal from stellate cells, as suggested by their waveforms and other physiological properties?

A progression of increasingly sophisticated techniques was used to address these questions, starting with routine histological methods, followed by *in vitro* brain slice recordings combined with confocal imaging, and most recently *in vivo* two photon (2p) imaging. Months of development and many experiments later, none of these approaches were ultimately successful to the point of resolving the questions raised above. Nonetheless this appendix documents the progress that was made towards future experiments to this end.

Early efforts investigated the feasibility of applying standard histological methods used routinely in single unit studies. The idea was simple: identify the polytrode position in the brain, either by current lesioning or staining of the electrode track, and stain the surrounding neurons. The cresyl violet stained section in Figure C.1A shows how naive this approach was. Even if the electrode track was precisely locatable, a veritable syncytium of cell bodies fill the field of view of a 50 μ m thick brain section, and polytrodes can record from any of the thousands of neurons in a tissue volume hemisphere within a ~150 μ m radius. The high cell density almost guaranteed that a predicted location would coincide with a cell body, but since there was no way to know which of the stained cells were active, registration of predicted locations with Nissl sections was meaningless.

On the reasonable assumption that the majority of recorded neurons were large pyramidal cells, the pyramid-specific antibody stain SMI-32 was explored next (Figure C.1B). Although more promising than cresyl violet, the capricious nature of this stain, indiscriminate labelling of both active and inactive neurons, and the inability to label non-pyramidal cells made SMI-32 of limited value for validating the localisation algorithm.

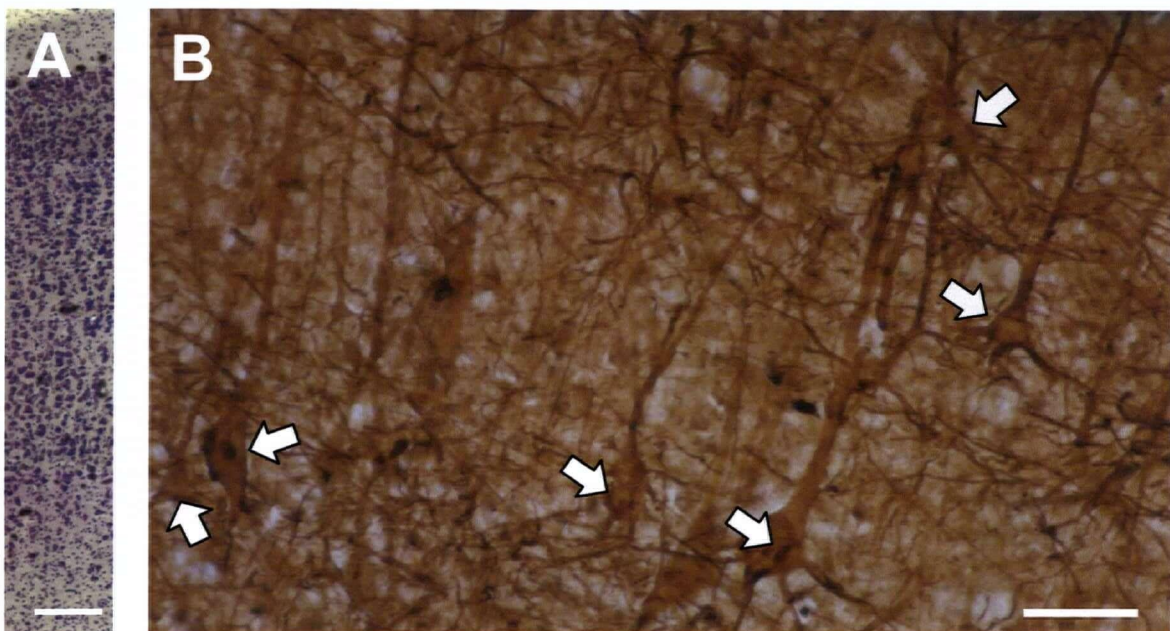


Figure C.1 Histological validation of the neuron localisation algorithm.

(A) Coronal cortical section, 50 μ m thick, stained with cresyl-violet. Dense labelling of neuronal and glial nuclei demonstrates why this method of validation was untenable. (B) SMI-32 stains large pyramidal neurons (*arrows*). Scalebars = 100 μ m and 50 μ m.

The second attempt combined simultaneous intracellular and extracellular recording with confocal imaging of neuron location in acute brain slice preparations. The inherent problems of histological validation were avoided because the intracellular electrode gave an unambiguous record of when the impaled neuron fired, in addition to precise control of that cell's activity under current clamp. Coronal brain slices, 350-400 μ m thick, were placed in a modified perfusion chamber with a side port to enable lateral insertion of the polytrode into the slice (Figure C.2A). Blind whole-cell patch-clamping was used initially, however since the microscope was inverted and had a maximum imaging depth of ~125 μ m, intracellular electrodes were found to be more suitable for recording neurons deep into the slice, ignoring those encountered *en route*. With some practice it was possible to routinely impale neurons under and near to the polytrode by 'dead reckoning'. Confocal reconstruction of iontophoretically filled neurons afforded a much better estimate of neuron location with respect to the polytrode than was possible with histological methods, in addition to information about cell type and orientation relative to the polytrode (Figure C.2B, C).

The inability to image cells deep in the slice was a major constraint, as was the poor axial resolution at depths beyond 50 μ m. More problematic, however, was the realisation that the polytrodes did not reliably couple to the tissue when surrounded by the fluid of the perfusion bath. Extracellular spikes that were occasionally recorded were of low amplitude and unsuitable for modeling. Obtaining good quality data with any one of these techniques – stable intracellular recordings, *in vitro* multiunit recordings, or fluorescent imaging of dye-filled cells in unfixed brain slices – is technically demanding. Achieving all three concurrently proved elusive.

The final, most ambitious endeavour involved *in vivo* polytrode recordings combined with 2p imaging in anesthetised rat (Figure C.3A). The major impasse of the *in vitro* experiments could be alleviated; polytrodes record reliably *in vivo*, and 2p microscopes are capable of high resolution imaging 500 μ m or more in the intact brain (Svoboda et al. 1999; Oheim et al. 2001). Two different techniques were established for determining the location of spiking neurons. In the initial experiments whole cell patch-clamping was used to record and fill neurons close to the polytrode (Figure C.3B). Subsequent experiments used bulk loading of calcium green dextran (Figure C.3C, D), with the aim of correlating the

extracellularly recorded spike trains with line-scanned traces of spike-related calcium transients from the ensemble of nearby active neurons.

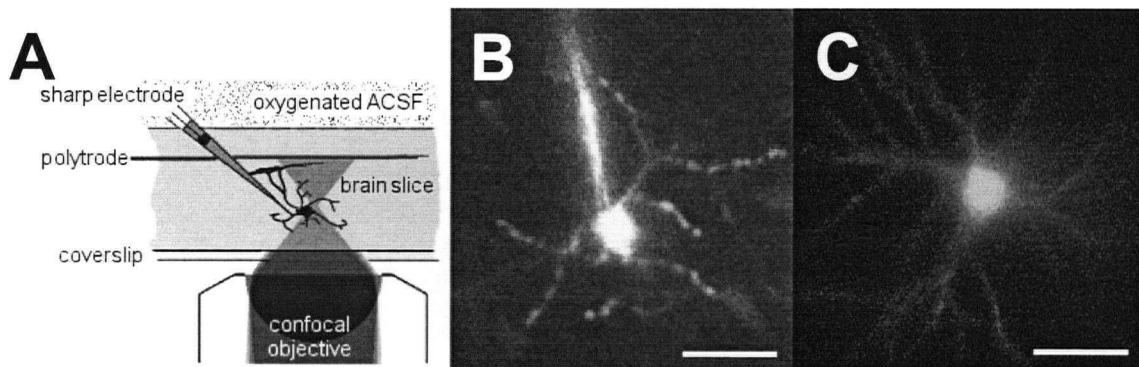


Figure C.2 *in vitro* validation of the neuron localisation algorithm.

(A) Experimental setup for simultaneous intra- and extracellular recording of single neurons imaged on a confocal microscope. (B) Imaging 54 μm into the slice, the dye filled neuron (Alexa 488, Molecular Probes) and its processes were clearly visible, including the intracellular microelectrode. (C) 108 μm into the slice, images were degraded due to attenuation of laser power and scatter of emission light by the unfixed brain tissue, however the location of the soma was still discernable. Scalebars = 50 μm .

Most of the technical difficulties associated with these experiments were surmountable. Firstly, a small stereotaxic frame had to be incorporated into the microscope stage. It was necessary to insert the polytrode into the brain very close to the microscope objective, in addition to making whole cell recordings (or imaging calcium transients) of neurons within $\sim 150\mu\text{m}$ of the polytrode shank. These physical constraints meant the epoxy around the bond area of the silicon probe had to be thinned so it did not obscure the field of view. In order to image the position of the polytrode fluorescent DiI was applied to the probe shank. Since it wasn't possible to use a glass coverslip to minimise brain movements and insert the polytrode at the same time, agar dissolved in artificial CSF was substituted. This gave access to the polytrode, but was not entirely effective at removing cardio ballistic brain movements. One solution in future might be to synchronise image acquisition to the phase of the ECG signal.

At the time these preliminary experiments were undertaken, the 2p microscope was still in development. Headscan mirrors unsuitable for infrared wavelengths and an inefficient photomultiplier tube meant the maximum imaging depth was only $\sim 100\mu\text{m}$, less for more weakly fluorescent bulk-loaded cells. These

critical components have since been upgraded, so this *in vivo* approach remains the most promising ways to validate the neuron localisation algorithm.

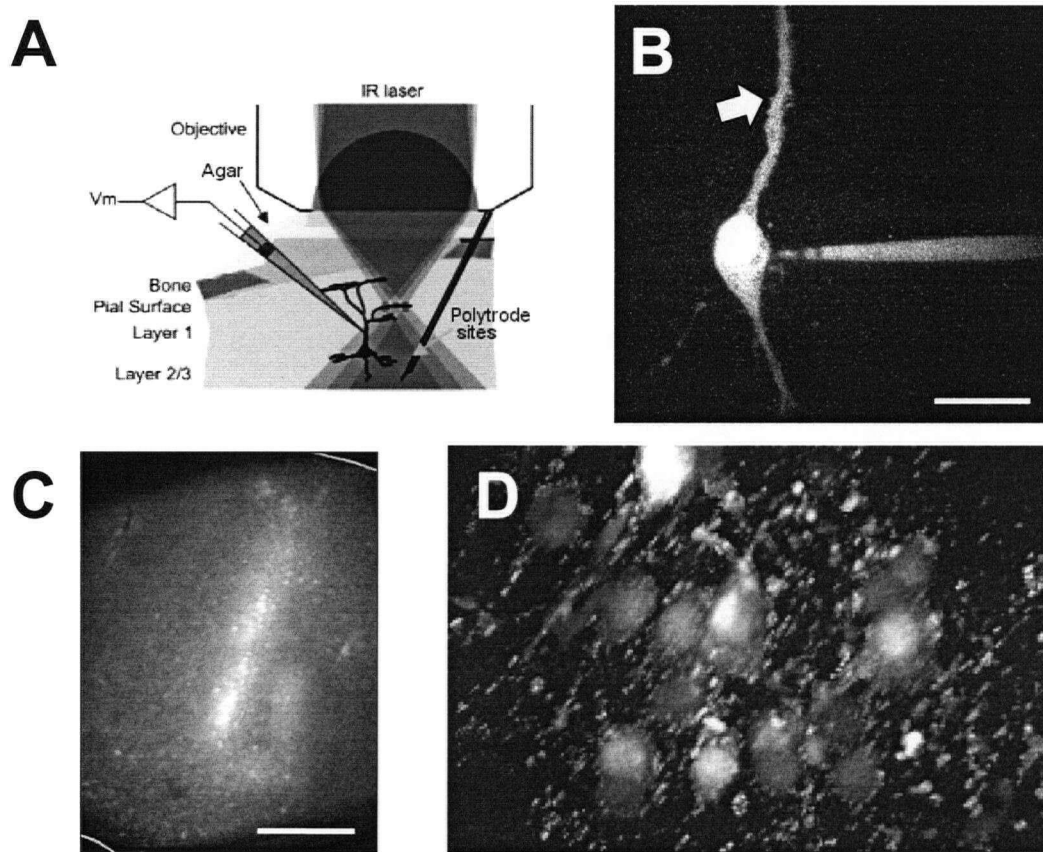


Figure C.3 *in vivo* validation of the neuron localisation algorithm.

(A) Experimental setup, modified from Svoboda et al. (1999). (B) The precise position and orientation of the patch-clamped neuron relative to the polytrode, and its detailed morphology, could be determined. The microscope had a resolving power capable of imaging individual spines $80\mu\text{m}$ into the brain (*arrow*). Scalebar = $30\mu\text{m}$. (C) Bulk loading of cortical neurons with the calcium indicator calcium green dextran (10,000 MW, Molecular Probes). Scalebar = $500\mu\text{m}$. (D) High power 3D reconstruction of neurons bulk loaded with calcium green.

references

- Abeles M, Goldstein M.** Multispikes train analysis. *Proc IEEE*, 1977; 65: 762-73.
- Aertsen AM, Gerstein GL.** Evaluation of neuronal connectivity: sensitivity of cross-correlation. *Brain Res*, 1985; 340: 341-54.
- Aertsen AM, Gerstein GL, Habib MK, Palm G.** Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J Neurophysiol*, 1989; 61: 900-17.
- Ahmed B, Allison JD, Douglas RJ, Martin KA.** An intracellular study of the contrast-dependence of neuronal activity in cat visual cortex. *Cereb Cortex*, 1997; 7: 559-70.
- Allison JD, Casagrande VA, Bonds AB.** The influence of input from the lower cortical layers on the orientation tuning of upper layer V1 cells in a primate. *Vis Neurosci*, 1995; 12: 309-20.
- Alonso JM, Martinez LM.** Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat Neurosci*, 1998; 1: 395-403.
- Alonso JM, Usrey WM, Reid RC.** Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *J Neurosci*, 2001; 21: 4002-15.
- Anderberg MR (1973)** *Cluster analysis for applications*. Academic Press, New York.
- Anderson DJ, Najafi K, Tanghe SJ, Evans DA, Levy KL, Hetke JF, Xue XL, Zappia JJ, Wise KD.** Batch-fabricated thin-film electrodes for stimulation of the central auditory system. *IEEE Trans Biomed Eng*, 1989; 36: 693-704.
- Anderson DJ, Oweiss KG, Bierer SM.** Sensor arrays in the micro-environment of the brain. *Acoustics, Speech, and Signal Processing*, 2001. Proceedings. 2001 IEEE International Conference on, 2001; 6: 3433-6 vol.6.
- Arieli A, Sterkin A, Grinvald A, Aertsen A.** Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 1996; 273: 1868-71.
- Atiya AF.** Recognition of multiunit neural signals. *IEEE Trans Biomed Eng*, 1992; 39: 723-9.
- Azouz R, Gray CM.** Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci*, 1999; 19: 2209-23.

Azouz R, Gray CM, Nowak LG, McCormick DA. Physiological properties of inhibitory interneurons in cat striate cortex. *Cereb Cortex*, 1997; 7: 534-45.

Bai Q, Wise KD. Single-unit neural recording with active microelectrode arrays. *IEEE Trans Biomed Eng*, 2001; 48: 911-20.

Ball GH, Hall DJ (1965) *ISODATA, A novel method of data analysis and pattern classification*. Stanford Research Institute, Menlo Park, CA.

Bar-Gad I, Ritov Y, Vaadia E, Bergman H. Failure of identification of overlapping spikes from multiple neuron activity causes artificial correlations. *J Neurosci Methods*, 2001; 107: 1-13.

Bartho P, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsaki G. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol*, 2004; 92: 600-8.

Beaulieu C, Colonnier M. 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*, 1985; 339: 166-70.

Bedard C, Kroger H, Destexhe A. Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophys J*, 2004; 86: 1829-42.

Bergman H, DeLong MR. A personal computer-based spike detector and sorter: implementation and evaluation. *J Neurosci Methods*, 1992; 41: 187-97.

Berry MJ, Warland DK, Meister M. The structure and precision of retinal spike trains. *Proc Natl Acad Sci U S A*, 1997; 94: 5411-6.

Best J, Reuss S, Dinse HR. Lamina-specific differences of visual latencies following photic stimulation in the cat striate cortex. *Brain Res*, 1986; 385: 356-60.

Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 1998; 18: 10464-72.

Bierer JA, Middlebrooks JC. Auditory cortical images of cochlear-implant stimuli: dependence on electrode configuration. *J Neurophysiol*, 2002; 87: 478-92.

Bierer SM, Anderson DJ. Multi-channel spike detection and sorting using an array processing technique. *Neurocomputing*, 1999; 26-27: 947-56.

Binzegger T, Douglas RJ, Martin KA. A quantitative map of the circuit of cat primary visual cortex. *J Neurosci*, 2004; 24: 8441-53.

Bishop PO, Burke W, Davis R. The identification of single units in central visual pathways. *J Physiol*, 1962a; 162: 409-31.

Bishop PO, Burke W, Davis R. The interpretation of the extracellular response of single lateral geniculate cells. *J Physiol*, 1962b; 162: 451-72.

Blanche TJ, Hetherington PA, Rennie CJ, Spacek MA, Swindale NV. Model-based 3D cortical neuron localization and classification with silicon electrode arrays. 32nd Annual Meeting of the Society for Neuroscience, 2003.

Bonhoeffer T, Grinvald A (1996) Optical Imaging Based on Intrinsic Signals in *Brain Mapping: The Methods*. Academic Press, Inc., pp. 55-97.

Borst A, Theunissen FE. Information theory and neural coding. *Nat Neurosci*, 1999; 2: 947-57.

Bragin A, Hetke J, Wilson CL, Anderson DJ, Engel J, Jr., Buzsaki G. Multiple site silicon-based probes for chronic recordings in freely moving rats: implantation, recording and histological verification. *J Neurosci Methods*, 2000; 98: 77-82.

Bredfeldt CE, Ringach DL. Dynamics of spatial frequency tuning in macaque V1. *J Neurosci*, 2002; 22: 1976-84.

Brown GD, Yamada S, Sejnowski TJ. Independent component analysis at the neural cocktail party. *Trends in Neurosci.*, 2001; 24: 54-63.

Brown KT, Tasaki K. Localization of electrical activity in the cat retina by an electrode marking method. *J Physiol London*, 1961; 158: 281-95.

Buchwald JS, Grover F. Amplitudes of background fast activity characteristic of specific brain sites. *J Neurophysiol*, 1970; 33: 148-59.

Buchwald JS, Holstein SB, Weber DS (1973) Multiple unit recording: Technique, Interpretation, and Experimental Applications in *Bioelectric Recording Techniques*. eds. Thompson, R. F. & Patterson, M. M. Academic Press, New York, Vol. Part A, pp. 201-42.

Bullier J, Nowak LG. Parallel versus serial processing: new vistas on the distributed organization of the visual system. *Curr Opin Neurobiol*, 1995; 5: 497-503.

Buzsaki G. Large-scale recording of neuronal ensembles. *Nat Neurosci*, 2004; 7: 446-51.

Buzsaki G, Horvath Z, Urioste R, Hetke J, Wise K. High-frequency network oscillation in the hippocampus. *Science*, 1992; 256: 1025-7.

Buzsaki G, Kandel A. Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat. *J Neurophysiol*, 1998; 79: 1587-91.

Buzsaki G, Penttonen M, Nadasdy Z, Bragin A. Pattern and inhibition-dependent invasion of pyramidal cell dendrites by fast spikes in the hippocampus in vivo. *Proc Natl Acad Sci U S A*, 1996; 93: 9921-5.

Campbell PK, Jones KE, Huber RJ, Horch KW, Normann RA. A silicon-based, three-dimensional neural interface: manufacturing processes for an intracortical electrode array. *IEEE Trans Biomed Eng*, 1991; 38: 758-68.

Celebrini S, Thorpe S, Trotter Y, Imbert M. Dynamics of orientation coding in area V1 of the awake primate. *Vis Neurosci*, 1993; 10: 811-25.

Chen J, Wise KD, Hetke JF, Bledsoe SC, Jr. A multichannel neural probe for selective chemical delivery at the cellular level. *IEEE Trans Biomed Eng*, 1997; 44: 760-9.

Churchland PS, Sejnowski TJ. Perspectives on cognitive neuroscience. *Science*, 1988; 242: 741-5.

Churchland PS, Sejnowski TJ (1992) *The computational brain*. MIT Press, Cambridge, Mass.

Claverol-Tinture E, Nadasdy Z. Intersection of microwire electrodes with proximal CA1 stratum-pyramidale neurons at insertion for multiunit recordings predicted by a 3-D computer model. *IEEE Trans Biomed Eng*, 2004; 51: 2211-6.

Comon P. Independent component analysis - a new concept. *Sig. Proc.*, 1994; 36: 287-314.

Connors BW, Gutnick MJ. Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci*, 1990; 13: 99-104.

Connors BW, Gutnick MJ, Prince DA. Electrophysiological properties of neocortical neurons in vitro. *J Neurophysiol*, 1982; 48: 1302-20.

Connors BW, Kriegstein AR. Cellular physiology of the turtle visual cortex: distinctive properties of pyramidal and stellate neurons. *J Neurosci*, 1986; 6: 164-77.

Connors BW, Regehr WG. Neuronal firing: does function follow form? *Curr Biol*, 1996; 6: 1560-2.

Contreras D, Palmer LA. Response to contrast of electrophysiologically defined cell classes in primary visual cortex. *J Neurosci*, 2003; 23: 6936-45.

Csicsvari J, Henze DA, Jamieson B, Harris KD, Sirota A, Bartho P, Wise KD, Buzsaki G. Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J Neurophysiol*, 2003; 90: 1314-23.

Cui X, Wiler J, Dzaman M, Altschuler RA, Martin DC. In vivo studies of polypyrrole/peptide coated neural probes. *Biomaterials*, 2003; 24: 777-87.

Cynader MS, Swindale NV, Matsubara JA. Functional topography in cat area 18. *J Neurosci*, 1987; 7: 1401-13.

Dan Y, Poo MM. Spike timing-dependent plasticity of neural circuits. *Neuron*, 2004; 44: 23-30.

DeAngelis GC, Ohzawa I, Freeman RD. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J Neurophysiol*, 1993a; 69: 1091-117.

DeAngelis GC, Ohzawa I, Freeman RD. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J Neurophysiol*, 1993b; 69: 1118-35.

Descalzo VF, Nowak LG, Brumberg JC, McCormick DA, Sanchez-Vives MV. Slow adaptation in fast-spiking neurons of visual cortex. *J Neurophysiol*, 2005; 93: 1111-8.

DeYoe EA, Van Essen DC. Concurrent processing streams in monkey visual cortex. *Trends Neurosci*, 1988; 11: 219-26.

Di Lollo V, Seiffert AE, Burchett G, Rabeeh R, Ruman TA. Phosphor persistence of oscilloscopic displays: a comparison of four phosphors. *Spat Vis*, 1997; 10: 353-60.

DiCarlo JJ, Lane JW, Hsiao SS, Johnson KO. Marking microelectrode penetrations with fluorescent dyes. *J Neurosci Methods*, 1996; 64: 75-81.

Diedrich A, Charoensuk W, Brychta RJ, Ertl AC, Shiavi R. Analysis of raw microneurographic recordings based on wavelet de-noising technique and classification algorithm: wavelet analysis in microneurography. *IEEE Trans Biomed Eng*, 2003; 50: 41-50.

Dinse HR, Kruger K. The timing of processing along the visual pathway in the cat. *Neuroreport*, 1994; 5: 893-7.

Drake KL, Wise KD, Farraye J, Anderson DJ, BeMent SL. Performance of planar multisite microprobes in recording extracellular single-unit intracortical activity. *IEEE Trans Biomed Eng*, 1988; 35: 719-32.

Emondi AA, Rebrik SP, Kurgansky AV, Miller KD. Tracking neurons recorded from tetrodes across time. *J Neurosci Methods*, 2004; 135: 95-105.

Ensell G, Banks DJ, Richards PR, Balachandran W, Ewins DJ. Silicon-based microelectrodes for neurophysiology, micromachined from silicon-on-insulator wafers. *Med Biol Eng Comput*, 2000; 38: 175-9.

Fabre-Thorpe M, Richard G, Thorpe SJ. Rapid categorization of natural images by rhesus monkeys. *Neuroreport*, 1998; 9: 303-8.

Fatt P. Electric potentials occurring around a neurone during its antidromic activation. *J Neurophysiol*, 1957; 20: 27-60.

Fee MS, Mitra PP, Kleinfeld D. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J Neurosci Methods*, 1996a; 69: 175-88.

Fee MS, Mitra PP, Kleinfeld D. Variability of extracellular spike waveforms of cortical neurons. *J Neurophysiol*, 1996b; 76: 3823-33.

Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, 1991; 1: 1-47.

Ferster D. Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *J Neurosci*, 1988; 8: 1172-80.

Ferster D, Chung S, Wheat H. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature*, 1996; 380: 249-52.

Ferster D, Miller KD. Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci*, 2000; 23: 441-71.

Forgy E. Cluster analysis of multivariate data: efficiency vs. interpretability of classifications. *Biometrics*, 1965; 21: 768.

Forster C, Handwerker HO. Automatic classification and analysis of microneurographic spike data using a PC/AT. *J Neurosci Methods*, 1990; 31: 109-18.

Fregnac Y, Shulz D, Thorpe S, Bienenstock E. A cellular analogue of visual cortical plasticity. *Nature*, 1988; 333: 367-70.

- Freund TF, Buzsaki G.** Interneurons of the hippocampus. *Hippocampus*, 1996; 6: 347-470.
- Freygang WH, Jr., Landau WM.** Some relations between resistivity and electrical activity in the cerebral cortex of the cat. *J Cell Physiol*, 1955; 45: 377-92.
- Froemke RC, Dan Y.** Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 2002; 416: 433-8.
- Fromherz P** (2003) Neuroelectronic Interfacing: Semiconductor Chips with Ion Channels, Nerve Cells, and Brain in *Nanoelectronics and Information Technology*. ed. Waser, R. Wiley-VCH Verlag, Berlin, pp. 781-810.
- Fu YX, Djupsund K, Gao H, Hayden B, Shen K, Dan Y.** Temporal specificity in the cortical plasticity of visual space representation. *Science*, 2002; 296: 1999-2003.
- Fung SH, Burstein D, Born RT.** In vivo microelectrode track reconstruction using magnetic resonance imaging. *J Neurosci Methods*, 1998; 80: 215-24.
- Gasparini S, Migliore M, Magee JC.** On the initiation and propagation of dendritic spikes in CA1 pyramidal neurons. *J Neurosci*, 2004; 24: 11046-56.
- Gautrais J, Thorpe S.** Rate coding versus temporal order coding: a theoretical approach. *Biosystems*, 1998; 48: 57-65.
- Gawne TJ, Kjaer TW, Hertz JA, Richmond BJ.** Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cereb Cortex*, 1996a; 6: 482-9.
- Gawne TJ, Kjaer TW, Richmond BJ.** Latency: another potential code for feature binding in striate cortex. *J Neurophysiol*, 1996b; 76: 1356-60.
- Gilbert CD, Sigman M, Crist RE.** The neural basis of perceptual learning. *Neuron*, 2001; 31: 681-97.
- Gilbert CD, Wiesel TN.** Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature*, 1979; 280: 120-5.
- Gray CM, Maldonado PE, Wilson M, McNaughton B.** Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods*, 1995; 63: 43-54.
- Gray CM, McCormick DA.** Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, 1996; 274: 109-13.

Gray CM, Singer W. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci U S A*, 1989; 86: 1698-702.

Green JD. A simple microelectrode for recording from the central nervous system. *Nature*, 1958; 182: 962.

Gur M, Beylin A, Snodderly DM. Physiological properties of macaque V1 neurons are correlated with extracellular spike amplitude, duration, and polarity. *J Neurophysiol*, 1999; 82: 1451-64.

Haberich FJ, Lingelbach B. Psychophysical measurements concerning the range of visual perception. *Pflugers Arch*, 1980; 386: 141-6.

Hammond P, MacKay DM. Modulatory influences of moving textured backgrounds on responsiveness of simple cells in feline striate cortex. *J Physiol*, 1981; 319: 431-42.

Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol*, 2000; 84: 401-14.

Hausser M. The Hodgkin-Huxley theory of the action potential. *Nat Neurosci*, 2000; 3 Suppl: 1165.

Hebb D (1949) *The organization of behaviour*. Wiley, New York.

Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsaki G. Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol*, 2000; 84: 390-400.

Hetherington PA, Blanche TJ, Swindale NV. Prediction of neuron 3D cortical location by triangulation of spike signals from silicon electrode recording arrays. 29th Annual Meeting of the Society for Neuroscience, 1999.

Hetherington PA, Swindale NV. Real-time, in vivo prediction of neuron cortical location by triangulation of spike signals from silicon electrode recording arrays. 28th Annual Meeting of the Society for Neuroscience, 1998; 24: 895.

Hetherington PA, Swindale NV. Receptive field and orientation scatter studied by tetrode recordings in cat area 17. *Vis Neurosci*, 1999; 16: 637-52.

Hetke JF, Lund JL, Najafi K, Wise KD, Anderson DJ. Silicon ribbon cables for chronically implantable microelectrode arrays. *IEEE Trans Biomed Eng*, 1994; 41: 314-21.

- Heynen AJ, Bear MF.** Long-term potentiation of thalamocortical transmission in the adult visual cortex in vivo. *J Neurosci*, 2001; 21: 9801-13.
- Hirsch JA, Martinez LM, Alonso JM, Desai K, Pillai C, Pierre C.** Synaptic physiology of the flow of information in the cat's visual cortex in vivo. *J Physiol*, 2002; 540: 335-50.
- Hirsch JA, Martinez LM, Pillai C, Alonso JM, Wang Q, Sommer FT.** Functionally distinct inhibitory neurons at the first stage of visual cortical processing. *Nat Neurosci*, 2003; 6: 1300-8.
- Hodgkin AL, Huxley AF.** A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 1952; 117: 500-44.
- Hoeltzell PB, Dykes RW.** Conductivity in the somatosensory cortex of the cat -- evidence for cortical anisotropy. *Brain Res*, 1979; 177: 61-82.
- Hoffman KL, McNaughton BL.** Coordinated reactivation of distributed memory traces in primate neocortex. *Science*, 2002; 297: 2070-3.
- Hoogerwerf AC, Wise KD.** A three-dimensional microelectrode array for chronic neural recording. *IEEE Trans Biomed Eng*, 1994; 41: 1136-46.
- Horowitz P, Hill W** (1989) *The Art of Electronics*. Cambridge University Press, Cambridge.
- Hubel DH, Wiesel TN.** Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol*, 1962; 160: 106-54.
- Hubel DH, Wiesel TN.** Shape and arrangement of columns in cat's striate cortex. *J Physiol*, 1963; 165: 559-70.
- Hubener M, Shoham D, Grinvald A, Bonhoeffer T.** Spatial relationships among three columnar systems in cat area 17. *J Neurosci*, 1997; 17: 9270-84.
- Hulata E, Segev R, Ben-Jacob E.** A method for spike sorting and detection based on wavelet packets and Shannon's mutual information. *J Neurosci Methods*, 2002; 117: 1-12.
- Hulata E, Segev R, Shapira Y, Benveniste M, Ben-Jacob E.** Detection and sorting of neural spikes using wavelet packets. *Phys Rev Lett*, 2000; 85: 4637-40.
- Humphrey DR.** Re-analysis of the antidromic cortical response. I. Potentials evoked by stimulation of the isolated pyramidal tract. *Electroencephalogr Clin Neurophysiol*, 1968; 24: 116-29.

Humphrey DR (1976) Neural networks and systems modeling in *Biological foundations of Biomedical Engineering*. ed. Kline, J. Little, Brown & Co., Boston.

Humphrey DR (1979) Extracellular, single unit recording methods in *Electrophysiological Techniques*. Society for Neuroscience, Bethesda, Maryland, pp. 199-259.

Humphrey DR, Corrie WS. Properties of pyramidal tract neuron system within a functionally defined subregion of primate motor cortex. *J Neurophysiol*, 1978; 41: 216-43.

Iriarte J, Urrestarazu E, Valencia M, Alegre M, Malanda A, Viteri C, Artieda J. Independent component analysis as a tool to eliminate artifacts in EEG: a quantitative study. *J Clin Neurophysiol*, 2003; 20: 249-57.

Ivorra A, Gomez R, Noguera N, Villa R, Sola A, Palacios L, Hotter G, Aguilo J. Minimally invasive silicon probe for electrical impedance measurements in small animals. *Biosens Bioelectron*, 2003; 19: 391-9.

Jog M, Connolly C, Kubota Y, Iyengarm D, Garrido L, Harlan R, Graybiel A. Tetrode technology: advances in implantable hardware, neuroimaging, and data analysis techniques. *J Neurosci Methods*, 2002; 117: 141-52.

Johnston D, Magee JC, Colbert CM, Cristie BR. Active properties of neuronal dendrites. *Annu Rev Neurosci*, 1996; 19: 165-86.

Joliffe IT, Morgan BJ. Principal component analysis and exploratory factor analysis. *Stat Methods Med Res*, 1992; 1: 69-95.

Jones JP, Palmer LA. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J Neurophysiol*, 1987; 58: 1187-211.

Kaiser JF. On a simple algorithm to calculate the 'energy' of a signal. *Proc IEEE ICASSP*, 1990; 1: 381-4.

Kandel A, Buzsaki G. Cellular-synaptic generation of sleep spindles, spike-and-wave discharges, and evoked thalamocortical responses in the neocortex of the rat. *J Neurosci*, 1997; 17: 6783-97.

Kay CF, Schwan HP. Specific resistance of body tissues. *Circ Res*, 1956; 4: 664-70.

Kayser C, Kording KP, Konig P. Learning the nonlinearity of neurons from natural visual stimuli. *Neural Comput*, 2003a; 15: 1751-9.

Kayser C, Salazar RF, Konig P. Responses to natural scenes in cat V1. *J Neurophysiol*, 2003b; 90: 1910-20.

Kelly JP, Van Essen DC. Cell structure and function in the visual cortex of the cat. *J Physiol*, 1974; 238: 515-47.

Kettenmann H, Grantyn R, eds. (1992) *Practical electrophysiological methods: a guide for in vitro studies in vertebrate neurobiology*. John Wiley & Sons, New York.

Kewley DT, Hills MD, Borkholder DA, Opris IE, Maluf NI, Storment CW, Bower JM, Kovacs GT. Plasma-etched neural probes. *Sensors and Actuators A: Physical*, 1997; A58: 27-35.

Kim KH, Kim SJ. Neural spike sorting under nearly 0-dB signal-to-noise ratio using nonlinear energy operator and artificial neural-network classifier. *IEEE Trans Biomed Eng*, 2000; 47: 1406-11.

Kovacs GT, Storment CW, Halks-Miller M, Belczynski CR, Jr., Della Santina CC, Lewis ER, Maluf NI. Silicon-substrate microelectrode arrays for parallel recording of neural activity in peripheral and cranial nerves. *IEEE Trans Biomed Eng*, 1994; 41: 567-77.

Lamme VA, Super H, Landman R, Roelfsema PR, Spekreijse H. The role of primary visual cortex (V1) in visual awareness. *Vision Res*, 2000; 40: 1507-21.

Laubach M. Wavelet-based processing of neuronal spike trains prior to discriminant analysis. *J Neurosci Methods*, 2004; 134: 159-68.

Lehmenkuhler A, Sykova E, Svoboda J, Zilles K, Nicholson C. Extracellular space parameters in the rat neocortex and subcortical white matter during postnatal development determined by diffusion analysis. *Neuroscience*, 1993; 55: 339-51.

Letelier JC, Weber PP. Spike sorting based on discrete wavelet transform coefficients. *J Neurosci Methods*, 2000; 101: 93-106.

Lewicki MS. Bayesian modeling and classification of neural signals. *Neural Comput*, 1994; 6: 1005-30.

Lewicki MS. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 1998; 9: R53-78.

Lloyd SP. Least squares quantization in PCM. *IEEE Trans Information Theory*, 1982; 28: 129-37.

Magee JC, Johnston D. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science*, 1997; 275: 209-13.

- Mahalanobis PC.** On the generalized distance in statistics. *Proc Nat Inst Sci (India)*, 1936; 2: 49-55.
- Mainen ZF, Sejnowski TJ.** Reliability of spike timing in neocortical neurons. *Science*, 1995; 268: 1503-6.
- Mainen ZF, Sejnowski TJ.** Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 1996; 382: 363-6.
- Maldonado PE, Godecke I, Gray CM, Bonhoeffer T.** Orientation selectivity in pinwheel centers in cat striate cortex. *Science*, 1997; 276: 1551-5.
- Maldonado PE, Gray CM.** Heterogeneity in local distributions of orientation-selective neurons in the cat primary visual cortex. *Vis Neurosci*, 1996; 13: 509-16.
- Malmivuo J, Plonsey R** (1995) *Bioelectromagnetism*. Oxford University Press, New York.
- Maragos P, Kaiser JF, Quatieri TF.** On amplitude and frequency demodulation using energy operators. *IEEE Trans Signal Process*, 1993; 41: 1532-50.
- Markram H, Lubke J, Frotscher M, Sakmann B.** Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 1997; 275: 213-5.
- Martinez LM, Alonso JM.** Construction of complex receptive fields in cat primary visual cortex. *Neuron*, 2001; 32: 515-25.
- Martinez LM, Alonso JM, Reid RC, Hirsch JA.** Laminar processing of stimulus orientation in cat visual cortex. *J Physiol*, 2002; 540: 321-33.
- Massart DL, Kaufman KK** (1983) *The interpretation of analytical chemical data by the use of cluster analysis*. John Wiley & Sons, New York.
- Maunsell JH, Gibson JR.** Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol*, 1992; 68: 1332-44.
- Mazel T, Simonova Z, Sykova E.** Diffusion heterogeneity and anisotropy in rat hippocampus. *Neuroreport*, 1998; 9: 1299-304.
- McCormick DA, Connors BW, Lighthall JW, Prince DA.** Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol*, 1985; 54: 782-806.
- McGill KC, Dorfman LJ.** High-resolution alignment of sampled waveforms. *IEEE Trans Biomed Eng*, 1984; 31: 462-8.

McNaughton BL, O'Keefe J, Barnes CA. The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci Methods*, 1983; 8: 391-7.

Meijering EH. A chronology of interpolation: From ancient astronomy to modern signal and image processing. *Proc IEEE*, 2002; 90: 319-42.

Mensinger AF, Anderson DJ, Buchko CJ, Johnson MA, Martin DC, Tresco PA, Silver RB, Highstein SM. Chronic recording of regenerating VIIIth nerve axons with a sieve electrode. *J Neurophysiol*, 2000; 83: 611-5.

Metha AB, Crane AM, Rylander HG, 3rd, Thomsen SL, Albrecht DG. Maintaining the cornea and the general physiological environment in visual neurophysiology experiments. *J Neurosci Methods*, 2001; 109: 153-66.

Millecchia R, McIntyre T. Automatic nerve impulse identification and separation. *Comput Biomed Res*, 1978; 11: 459-68.

Miller KD, Pinto DJ, Simons DJ. Processing in layer 4 of the neocortical circuit: new insights from visual and somatosensory cortex. *Curr Opin Neurobiol*, 2001; 11: 488-97.

Mitzdorf U. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev*, 1985; 65: 37-100.

Mitzdorf U, Singer W. Prominent excitatory pathways in the cat visual cortex (A 17 and A 18): a current source density analysis of electrically evoked potentials. *Exp Brain Res*, 1978; 33: 371-94.

Mountcastle VB, Davies PW, Berman AL. Response properties of neurons of cat's somatic sensory cortex to peripheral stimuli. *J Neurophysiol*, 1957; 29: 374-407.

Mountcastle VB, Talbot WH, Sakata H, Hyvarinen J. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *J Neurophysiol*, 1969; 32: 452-84.

Mukhopadhyay S, Ray GC. A new interpretation of nonlinear energy operator and its efficacy in spike detection. *IEEE Trans Biomed Eng*, 1998; 45: 180-7.

Musial PG, Baker SN, Gerstein GL, King EA, Keating JG. Signal-to-noise ratio improvement in multiple electrode recording. *J Neurosci Methods*, 2002; 115: 29-43.

Najafi K, Hetke JF. Strength characterization of silicon microprobes in neurophysiological tissues. *IEEE Trans Biomed Eng*, 1990; 37: 474-81.

Najafi K, Ji J, Wise KD. Scaling limitations of silicon multichannel recording probes. *IEEE Trans Biomed Eng*, 1990; 37: 1-11.

Najafi K, Wise K, Mochizuki T. A high-yield ic-compatible multichannel recording array. *IEEE Trans Electron Devices*, 1985; 32: 1206-11.

Najafi K, Wise KD. An implantable multielectrode array with on-chip signal processing. *IEEE J Solid-State Circuits*, 1986; 21: 1035-44.

Nicholson C. Theoretical analysis of field potentials in anisotropic ensembles of neuronal elements. *IEEE Trans Biomed Eng*, 1973; 20: 278-88.

Nicholson C, Freeman JA. Theory of current source-density analysis and determination of conductivity tensor for anuran cerebellum. *J Neurophysiol*, 1975; 38: 356-68.

Nicholson C, Sykova E. Extracellular space structure revealed by diffusion analysis. *Trends Neurosci*, 1998; 21: 207-15.

Nicholson C, Tao L. Hindered diffusion of high molecular weight compounds in brain extracellular microenvironment measured with integrative optical imaging. *Biophys J*, 1993; 65: 2277-90.

Niparko JK, Altschuler RA, Xue XL, Wiler JA, Anderson DJ. Surgical implantation and biocompatibility of central nervous system auditory prostheses. *Ann Otol Rhinol Laryngol*, 1989; 98: 965-70.

Norlin P, Kindlundh M, Mouroux A, Yoshida K, Hofmann UG. A 32-site neural recording probe fabricated by DRIE of SOI substrates. *J Micromech Microeng*, 2002; 12: 414-9.

Nowak LG, Azouz R, Sanchez-Vives MV, Gray CM, McCormick DA. Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quantitative analyses. *J Neurophysiol*, 2003; 89: 1541-66.

Nowak LG, Munk MH, Girard P, Bullier J. Visual latencies in areas V1 and V2 of the macaque monkey. *Vis Neurosci*, 1995; 12: 371-84.

Nunez A, Amzica F, Steriade M. Electrophysiology of cat association cortical cells in vivo: intrinsic properties and synaptic responses. *J Neurophysiol*, 1993; 70: 418-30.

Nyquist H. Certain topics in telegraph transmission theory. *Trans AIEE*, 1928; 47: 617-44.

- Obeid I, Wolf PD.** Evaluation of spike detection algorithms for a brain machine interface application. *IEEE Trans Biomed Eng*, 2003; 51: 905-11.
- Oheim M, Beauprepaire E, Chaigneau E, Mertz J, Charpak S.** Two-photon microscopy in brain tissue: parameters influencing the imaging depth. *J Neurosci Methods*, 2001; 111: 29-37.
- Okada YC, Huang JC, Rice ME, Tranchina D, Nicholson C.** Origin of the apparent tissue conductivity in the molecular and granular layers of the in vitro turtle cerebellum and the interpretation of current source-density analysis. *J Neurophysiol*, 1994; 72: 742-53.
- Olshausen BA, Field DJ** (2005) What is the other 85% of V1 doing? in *Problems in Systems Neuroscience.*, eds. Sejnowski, T. J. & van Hemmen, L. Oxford University Press., Oxford.
- Oram MW, Wiener MC, Lestienne R, Richmond BJ.** Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol*, 1999; 81: 3021-33.
- Panzeri S, Treves A, Schultz S, Rolls ET.** On decoding the responses of a population of neurons from short time windows. *Neural Comput*, 1999; 11: 1553-77.
- Parsons** (1959) *Handbook of electrochemical constants*. Academic Press.
- Paz R, Wise SP, Vaadia E.** Viewing and doing: similar cortical mechanisms for perceptual and motor learning. *Trends Neurosci*, 2004; 27: 496-503.
- Peters A, Yilmaz E.** Neuronal organization in area 17 of cat visual cortex. *Cereb Cortex*, 1993; 3: 49-68.
- Plonsey R** (1969) *Bioelectric phenomena*. McGraw-Hill, New York.
- Plonsey R, Heppner DB.** Considerations of quasi-stationarity in electrophysiological systems. *Bull Math Biophys*, 1967; 29: 657-64.
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP** (1994) *Numerical Recipes in Fortran 77: The Art of Scientific Computing*. Cambridge University Press, Cambridge, Mass.
- Quian Quiroga R, Garcia H.** Single-trial event-related potentials with wavelet denoising. *Clin Neurophysiol*, 2003; 114: 376-90.
- Quirk MC, Blum KI, Wilson MA.** Experience-dependent changes in extracellular spike amplitude may reflect regulation of dendritic action potential back-propagation in rat hippocampal pyramidal cells. *J Neurosci*, 2001; 21: 240-8.

- Quiroga RQ, Nadasdy Z, Ben-Shaul Y.** Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput*, 2004; 16: 1661-87.
- Raiguel SE, Lagae L, Gulyas B, Orban GA.** Response latencies of visual cells in macaque areas V1, V2 and V5. *Brain Res*, 1989; 493: 155-9.
- Rall W.** Electrophysiology of a dendritic neuron model. *Biophys J*, 1962; 2(2)Pt 2: 145-67.
- Ranck JB, Jr.** Specific impedance of rabbit cerebral cortex. *Exp Neurol*, 1963; 7: 144-52.
- Rathnasingham R, Kipke DR, Bledsoe SC, Jr., McLaren JD.** Characterization of implantable microfabricated fluid delivery devices. *IEEE Trans Biomed Eng*, 2004; 51: 138-45.
- Rebrik SP, Wright BD, Emondi AA, Miller KD.** Cross-channel correlations in tetrode recordings: implications for spike sorting. *Neurocomputing*, 1999; 26-27: 1033-8.
- Reid RC, Alonso JM.** Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 1995; 378: 281-84.
- Rice ME, Okada YC, Nicholson C.** Anisotropic and heterogeneous diffusion in the turtle cerebellum: implications for volume transmission. *J Neurophysiol*, 1993; 70: 2035-44.
- Ringach DL, Hawken MJ, Shapley R.** Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, 1997; 387: 281-4.
- Ringach DL, Hawken MJ, Shapley R.** Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *J Vis*, 2002; 2: 12-24.
- Rizzolatti G, Camarda R.** Influence of the presentation of remote visual stimuli on visual responses of cat area 17 and lateral suprasylvian area. *Exp Brain Res*, 1977; 29: 107-22.
- Roelfsema PR, Lamme VA, Spekreijse H.** Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nat Neurosci*, 2004; 7: 982-91.
- Rolls ET, Treves A, Tovee MJ, Panzeri S.** Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *J Comput Neurosci*, 1997; 4: 309-33.
- Ronen I, Kim KH, Garwood M, Ugurbil K, Kim DS.** Conventional DTI vs. slow and fast diffusion tensors in cat visual cortex. *Magn Reson Med*, 2003; 49: 785-90.

Rosenthal F, Woodbury JW, Patton HD. Dipole characteristics of pyramidal cell activity in cat postcruciate cortex. *J Neurophysiol*, 1966; 29: 612-25.

Rousche PJ, Normann RA. Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex. *J Neurosci Methods*, 1998; 82: 1-15.

Rousche PJ, Normann RA. Chronic intracortical microstimulation (ICMS) of cat sensory cortex using the Utah Intracortical Electrode Array. *IEEE Trans Rehabil Eng*, 1999; 7: 56-68.

Saul AB, Humphrey AL. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J Neurophysiol*, 1990; 64: 206-24.

Schafer RW, Rabiner LR. A digital signal processing approach to interpolation. *Proc IEEE*, 1973; 61: 692-702.

Schiller J, Schiller Y, Stuart G, Sakmann B. Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *J Physiol*, 1997; 505 (Pt 3): 605-16.

Schmidt EM. Computer separation of multi-unit neuroelectric data: a review. *J Neurosci Methods*, 1984; 12: 95-111.

Schmolesky MT, Wang Y, Hanes DP, Thompson KG, Leutgeb S, Schall JD, Leventhal AG. Signal timing across the macaque visual system. *J Neurophysiol*, 1998; 79: 3272-8.

Schwan HP, Kay CF. Capacitive properties of body tissues. *Circ Res*, 1957a; 5: 439-43.

Schwan HP, Kay CF. The conductivity of living tissues. *Ann N Y Acad Sci*, 1957b; 65: 1007-13.

Shalit M, Mahler Y. Electrode polarization effects in brain impedance measurements. *Electroencephalogr Clin Neurophysiol*, 1967; 22: 591.

Shalit MN, Mahler Y. Brain impedance measurements by the use of small bipolar needle electrodes. *J Appl Physiol*, 1966; 21: 1237-42.

Shannon CE. Communication in the presence of noise. *Proc Inst Radio Eng*, 1949; 37: 10-21.

Sherk H. Area 18 cell responses in cat during reversible inactivation of area 17. *J Neurophysiol*, 1978; 41: 204-15.

Skottun BC, De Valois RL, Grosz DH, Movshon JA, Albrecht DG, Bonds AB. Classifying simple and complex cells on the basis of response modulation. *Vision Res*, 1991; 31: 1079-86.

Smith JO (2004) *Digital Audio Resampling Home Page*. <http://www-ccrma.stanford.edu/~jos/resample/>.

Smith SW (1997) *The Scientist and Engineer's Guide to Digital Signal Processing*. California Technical Publishing, San Diego.

Smyth D, Willmore B, Baker GE, Thompson ID, Tolhurst DJ. The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J Neurosci*, 2003; 23: 4746-59.

Sneath PHA, Sokal RR (1973) *Numerical taxonomy*. W.H. Freeman, San Francisco.

Snider RK, Kabara JF, Roig BR, Bonds AB. Burst firing and modulation of functional connectivity in cat striate cortex. *J Neurophysiol*, 1998; 80: 730-44.

Spence AJ, Hoy RR, Isaacson MS. A micromachined silicon multielectrode for multiunit recording. *J Neurosci Methods*, 2003; 126: 119-26.

Spruston N, Schiller Y, Stuart G, Sakmann B. Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science*, 1995; 268: 297-300.

Stanley GB, Li FF, Dan Y. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J Neurosci*, 1999; 19: 8036-42.

Starr A, Wise K, Csongradi J. An evaluation of photoengraved microelectrodes for extracellular single-unit recording. *IEEE Trans Biomed Eng*, 1973; 20: 291-3.

Stillito AM. GABA mediated inhibitory processes in the function of the geniculo-striate system. *Prog Brain Res*, 1992; 90: 349-84.

Stone J, Dreher B, Leventhal A. Hierarchical and parallel mechanisms in the organization of visual cortex. *Brain Res*, 1979; 180: 345-94.

Stuart GJ, Sakmann B. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, 1994; 367: 69-72.

Suesserman MF, Spelman FA. Quantitative in vivo measurements of inner ear tissue resistivities: I. In vitro characterization. *IEEE Trans Biomed Eng*, 1993; 40: 1032-47.

Sutter EE. The fast m-Transform: a fast computation of cross-correlations with binary m-sequences. *SIAM J. Comput.*, 1991; 20: 686-94.

Svoboda K, Helmchen F, Denk W, Tank DW. Spread of dendritic excitation in layer 2/3 pyramidal neurons in rat barrel cortex in vivo. *Nat Neurosci*, 1999; 2: 65-73.

Swadlow HA. Efferent neurons and suspected interneurons in second somatosensory cortex of the awake rabbit: receptive fields and axonal properties. *J Neurophysiol*, 1991; 66: 1392-409.

Swadlow HA. Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: axonal properties, sensory receptive fields, and subthreshold synaptic inputs. *J Neurophysiol*, 1994; 71: 437-53.

Swindale NV, Matsubara JA, Cynader MS. Surface organization of orientation and direction selectivity in cat area 18. *J Neurosci*, 1987; 7: 1414-27.

Swindale NV, Shoham D, Grinvald A, Bonhoeffer T, Hubener M. Visual cortex maps are optimized for uniform coverage. *Nat Neurosci*, 2000; 3: 822-6.

Takahashi K, Matsuo T. Integration of multi-microelectrode and interface circuits by silicon planar and three-dimensional fabrication technology. *Sensors and Actuators*, 1984; 5: 89-99.

Takahashi S, Anzai Y, Sakurai Y. Automatic sorting for multi-neuronal activity recorded with tetrodes in the presence of overlapping spikes. *J Neurophysiol*, 2003; 89: 2245-58.

Thompson RF, Patterson MM, eds. (1973) *Bioelectric recording techniques. Part A. Cellular processes and brain potentials.* Academic Press, New York.

Thorpe S, Fize D, Marlot C. Speed of processing in the human visual system. *Nature*, 1996; 381: 520-2.

Towe AL (1973) Sampling Single Neuron Activity in *Bioelectric Recording Techniques*. eds. Thompson, R. F. & Patterson, M. M. Academic Press, New York, Vol. A, pp. 79-93.

Townsend G, Peloquin P, Kloosterman F, Hetke JF, L.S. L. Recording and marking with silicon multichannel electrodes. *Brain Res Protocols*, 2002; 9: 122-9.

Trachtenberg JT, Chen BE, Knott GW, Feng G, Sanes JR, Welker E, Svoboda K. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 2002; 420: 788-94.

- Troyer TW, Krukowski AE, Priebe NJ, Miller KD.** Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *J Neurosci*, 1998; 18: 5908-27.
- Tuch DS, Wedeen VJ, Dale AM, George JS, Belliveau JW.** Conductivity tensor mapping of the human brain using diffusion tensor MRI. *Proc Natl Acad Sci U S A*, 2001; 98: 11697-701.
- Van Essen D, Kelly J.** Correlation of cell shape and function in the visual cortex of the cat. *Nature*, 1973; 241: 403-5.
- Van Rullen R, Thorpe SJ.** Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Comput*, 2001; 13: 1255-83.
- Vetter RJ, Williams JC, Hetke JF, Nunamaker EA, Kipke DR.** Chronic neural recording using silicon-substrate microelectrode arrays implanted in cerebral cortex. *IEEE Trans Biomed Eng*, 2004; 51: 896-904.
- Vinje WE, Gallant JL.** Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *J Neurosci*, 2002; 22: 2904-15.
- Vornov JJ, Tasker RC, Coyle JT.** Direct observation of the agonist-specific regional vulnerability to glutamate, NMDA, and kainate neurotoxicity in organotypic hippocampal cultures. *Exp Neurol*, 1991; 114: 11-22.
- Wehr M, Pezaris JS, Sahani M.** Simultaneous paired intracellular and tetrode recordings for evaluating the performance of spike sorting algorithms. *Neurocomputing*, 1998; 26-27: 1061-8.
- Weiland JD, Anderson DJ.** Chronic neural stimulation with thin-film, iridium oxide electrodes. *IEEE Trans Biomed Eng*, 2000; 47: 911-8.
- Wheeler BC, Heetderks WJ.** A comparison of techniques for classification of multiple neural signals. *IEEE Trans Biomed Eng*, 1982; 29: 752-9.
- Wheeler BC, Smith SR.** High-resolution alignment of action potential waveforms using cubic spline interpolation. *J Biomed Eng*, 1987; 10: 47-53.
- Williams PE, Mechler F, Gordon J, Shapley R, Hawken MJ.** Entrainment to video displays in primary visual cortex of macaque and humans. *J Neurosci*, 2004; 24: 8278-88.
- Wilson MA, McNaughton BL.** Dynamics of the hippocampal ensemble code for space. *Science*, 1993; 261: 1055-8.

- Wise KD, Najafi K.** Microfabrication techniques for integrated sensors and microsystems. *Science*, 1991; 254: 1335-42.
- Wollman DE, Palmer LA.** Phase locking of neuronal responses to the vertical refresh of computer display monitors in cat lateral geniculate nucleus and striate cortex. *J Neurosci Methods*, 1995; 60: 107-13.
- Worgotter F, Eysel UT.** Context, state and the receptive fields of striatal cortex cells. *Trends Neurosci*, 2000; 23: 497-503.
- Worgotter F, Suder K, Zhao Y, Kerscher N, Eysel UT, Funke K.** State-dependent receptive-field restructuring in the visual cortex. *Nature*, 1998; 396: 165-8.
- Yang X, Shamma SA.** A totally automated system for the detection and classification of neural spikes. *IEEE Trans Biomed Eng*, 1988; 35: 806-16.
- Yao H, Dan Y.** Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron*, 2001; 32: 315-23.
- Yedlin M, Kwan H, Murphy JT, Nguyen-Huu H, Wong YC.** Electrical conductivity in cat cerebellar cortex. *Exp Neurol*, 1974; 43: 555-69.
- Yoon TH, Hwang EJ, Shin DY, Park SI, Oh SJ, Jung SC, Shin HC, Kim SJ.** A micromachined silicon depth probe for multichannel neural recording. *IEEE Trans Biomed Eng*, 2000; 47: 1082-7.
- Zhang LI, Tao HW, Holt CE, Harris WA, Poo M.** A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 1998; 395: 37-44.
- Zhang PM, Wu JY, Zhou Y, Liang PJ, Yuan JQ.** Spike sorting based on automatic template reconstruction with a partial solution to the overlapping problem. *J Neurosci Methods*, 2004; 135: 55-65.
- Zhu Z, Lin K, Kasamatsu T.** Artifactual synchrony via capacitance coupling in multi-electrode recording from cat striate cortex. *J Neurosci Methods*, 2002a; 115: 45-53.
- Zhu Z, Wang Y, Xu XZ, Li CY.** A simple and effective method for obtaining stable in vivo whole-cell recordings from visual cortical neurons. *Cereb Cortex*, 2002b; 12: 585-9.
- Zouridakis G, Tam DC.** Multi-unit spike discrimination using wavelet transforms. *Comput Biomed Res*, 1997; 27: 9-18.