
Automatic Classification of Neural Data

JUAN MARTÍNEZ GÓMEZ

Thesis submitted for the degree of
Doctor of Philosophy

July 19, 2011

Supervisor: Prof. Rodrigo Quian Quiroga
Co-Supervisor: Dr. Matias J. Ison



University of
Leicester

Department of Engineering

Abstract

Automatic Classification of Neural Data

JUAN MARTÍNEZ GÓMEZ

In this thesis we present a new solution for an automatic classification of the single-neuron activity. The study of the computational role of individual neurons underlying different cognitive process is a gold standard in Neuroscience. This type of analysis is done first, by recording the extracellular spikes of the neurons near the tip of a microelectrode and second, by isolating the spikes of the recorded cells based on the similarity of their shapes using a method called spike sorting.

In recent years, important advances in microelectrode technology allow us now to perform massive parallel recordings using a high number of channels with the possibility to study the activity of large ensembles of neurons at a time. However, this fascinating opportunity introduces at the same time a challenge for the efficient and fast analysis of this data.

In this research work, we address this problem by developing a new implementation for unsupervised spike sorting that improves the performance of a widely-used spike sorting algorithm, increasing the number of automatically identified neurons. Moreover, we developed a new testing platform which generates simulations of extracellular recordings including challenging conditions such as realistic noise, multi-unit activity -spikes of distant neurons impossible to be identified as single units- or the presence of neurons with low firing rates.

In summary, the results presented here provide contributions to the development of automated and efficient quantitative frameworks for the analysis of multiple-channel recordings that help us to understand single-neuron population codes.

Acknowledgements

I would like thank very specially to Professor Rodrigo Quian Quiroga, supervisor of this work, for giving me the opportunity to learn and whose valuable dedication and essential help has made possible the completion of this thesis.

I would like to thank to Dr. Matias J. Ison, co-supervisor of this work, for his important contribution to this research and for being always available to teach and help.

To all my friends in the Neuroengineering Lab and the Department of Engineering for being always generous and making of this an amazing and fun experience.

To Priya who turned Leicester into a home for creating special experiences and for being a positive part of my life.

Last, to my mother Rosalía for her loyal support and for making everything possible.

This work was supported by the Engineering and Physical Sciences Research Council.

©Juan Martínez Gómez, 2011

This thesis is copyright material and no quotation from it may be published without proper acknowledgement

List of Publications

Journals

Martinez J., Pedreira C., Ison M.J. and Quian Quiroga R. (2009). “Realistic simulation of extracellular recordings”. *J. Neurosci. Methods* 184(2) pp 285-293.

Pedreira C., **Martinez J.**, Ison M.J. and Quian Quiroga R. “How many neurons can we detect with current spike sorting algorithms”. *J. Neurosci. Methods* (in revision).

Alarcón G.*, **Martinez J.***, Kerai S.V, Lacruz E.M, Quian Quiroga R., Selway R.P., Richardson M.P., Garcia Seoane J.J, Valentín A. “Interictal single neuron activity recorded in vivo in human focal epilepsy ”. (to be submitted).

* *Equal contribution*

Martinez J., Ison M.J., Pedreira C. and Quian Quiroga R. “An improved algorithm for unsupervised spike sorting”. *J. Neurosci. Methods* (to be submitted).

Ison M.J., **Martinez J.** and Quian Quiroga R. “Automatic classification of single- and multi-unit activity from extracellular recordings”. *IEEE Trans. Biomed. Eng* (to be submitted).

Peer reviewed proceedings

Pedreira C., **Martinez J.** and Quian Quiroga R. “Neural Prostheses: linking brain signals to prosthetic devices”. *ICROS-SICE International joint conference*. Fukuoka, Japan. August 2009.

Poster presentations

Ison M.J., **Martinez J.** and Quian Quiroga R. “Automatic classification of single- and multi-unit activity from extracellular recordings”. *Workshop on Spike Train Measures and Their Applications to Neural Coding*. Plymouth, UK. June 2010.

Mazzoni A., Kayser C., Murayama Y., **Martinez J.**, Quian Quiroga R., Logothetis N.K. and Panzeri S. “Primary visual cortex encodes complementary information about naturalistic movies at different temporal scales”. *Research in Encoding And Decoding of Neural Ensembles*. Santorini, Greece. June 2010.

Martinez J., Pedreira C., Ison M.J. and Quian Quiroga R. “Realistic simulation of extracellular recordings”. *INCF UK Node Congress*. Edinburgh, UK. February 2010.

Martinez J., Pedreira C., Ison M.J. and Quian Quiroga R. “An improved algorithm for unsupervised spike sorting”. *6th FENS Forum of European Neuroscience*. Geneva, Switzerland. July 2008.

Contents

Abstract	i
Acknowledgements	ii
List of Publications	iii
Table of Contents	v
List of Figures	viii
List of Tables	x
I Introduction	1
1 Motivation	2
1.1 Organization of the Thesis	5
2 Theoretical Background	7
2.1 Techniques to Study the Brain Activity	7
2.2 Neuronal Action Potentials	11
2.3 Extracellular Recordings	13
2.4 Isolation of Single Neuron Activity	15
2.5 The Spike Sorting Method	18
2.6 Wave_clus	21
2.7 Summary of Chapter 2	25
II Results	27
3 Simulation of Extracellular Recordings	28

3.1	Introduction	28
3.2	Materials and Methods	30
3.2.1	Generation of Background Noise	31
3.2.2	Generation of Multi-unit Activity	33
3.2.3	Generation of Single-unit Activity	33
3.2.4	Real Data	34
3.2.5	Spike Detection and Sorting	34
3.3	Results	35
3.3.1	Real Data	35
3.3.2	Simulation of the Background Noise Activity	36
3.3.3	Matching of the Noise Amplitude Distribution	36
3.3.4	Matching of the Noise Power Spectra	40
3.3.5	Addition of the Single-unit and Multi-unit Activity	41
3.3.6	Testing of the Spike Sorting Algorithm	43
3.4	Discussion	45
3.4.1	Hybrid Simulation Approach	45
3.4.2	Multi-unit Activity	46
3.4.3	Background Noise	47
3.5	Summary of Chapter 3	49
4	Automatic Spike Sorting	50
4.1	Introduction	50
4.2	Materials and Methods	51
4.2.1	Improvements	54
4.2.2	Simulated Data	56
4.2.3	Algorithm Optimization and Validation	56
4.3	Results	57
4.3.1	Optimization	57
4.3.2	Example of Use of Criterion 1	58
4.3.3	Example of Use of Criterion 2	58

<i>CONTENTS</i>	vii
4.3.4 General Results	60
4.4 Discussion	63
4.5 Summary of Chapter 4	65
III Conclusions	67
5 General Conclusions	68
5.1 Simulation of Extracellular Recordings	69
5.2 Automatic Spike Sorting	70
6 Suggestions for Future Work	72
6.1 Improvements on Simulation of Extracellular Recordings	72
6.2 Spike Sorting on Multi-Electrode Data	73
IV Appendices	75
A Simulations	76
A.1 Extracellular Recordings	76
A.2 Simulations for Testing the Spike Sorting Algorithm	77

List of Figures

2.1	Example of spike shape variability	13
2.2	Example of single-unit isolation	16
2.3	Spike sorting procedure	19
2.4	GUI of Wave_clus showing the outcome of the automatic spike sorting	25
2.5	GUI of Wave_clus showing the outcome of the manual spike sorting	26
3.1	Overview of the simulation approach	31
3.2	Real data	37
3.3	Real spikes	38
3.4	Voltage signal and amplitude distribution of the noise	39
3.5	Levels of Gaussian noise	42
3.6	Example of a 2 minute simulation	44
3.7	Result of the automatic clustering	45
4.1	Limitations of the former algorithm	52
4.2	Example of use of criterion 1	59
4.3	Example of use of criterion 2	61
4.4	General results	62
4.5	Scatter plot firing rate vs. peak amplitude	63
4.6	Outcome of the former and new algorithm for different firing rate and peak amplitude bands	64
A.1	Scaling exponents	78

A.2 Simulated clusters	79
A.3 Power density spectrum	80
A.4 Amplitude distribution	81

List of Tables

3.1	Simulation and spectral features	35
3.2	Performance of the unsupervised spike sorting algorithm	45
4.1	Performance of the former and new algorithm	57

Part I

Introduction

Chapter 1

Motivation

One of the most attractive challenges in science is to understand how the brain works. How do different performances of individual nerve cells give rise to perception, movement, learning or memory? How do patterns of connections between neurons are related to behaviour? How are these neurons and their connections modified by experience? A well established methodology to address these questions is to study the extracellular activity of the individual neurons underlying cognitive processes.

Recording extracellularly from single cells, by positioning conductive microelectrodes near their cell body, permits us to see groups of neurons in action. Single-cell recordings have been performed in different regions of the brain in order to associate the role of individual neurons with specific mental processes such as vision, perception, movement, memory or consciousness.

In the visual cortex, located in the occipital lobe, neurons were found to encode basic image features such as local orientation at particular spatial locations (Hubel and Wiesel, 1962; Gallant et al., 1993), which contribute to the understanding of the neural computations during the early stages of the visual processing. Also, recordings in the inferior temporal lobe (IT) led to the discovery of neurons that process more complex visual stimuli such as faces or objects (Gross et al., 1969, 1972; Desimone et al., 1984; Logothetis et al., 1994; Tanaka, 1996), leading to the association of IT to the final stages of visual processing.

Important advances in the study of movement have been made by means of analysing the extracellular activity of individual neurons in different brain regions. For example, recordings in the arm area of the motor cortex in monkeys had contributed to the understanding of how the neural coding underlying movement direction (Georgopoulos et al., 1986) or the intention of movement (Snyder et al., 1997). These findings had led to the development of exciting and promising clinical applications that might help in the treatment of paralysed patients. This relatively new and popular discipline called Brain-machine interface, studies the activity of large population of neurons in specific areas of the brain related to motor action (e.g. motor cortex, posterior parietal cortex) to extract optimal commands that make possible the control of prosthetic limbs (Nicoletis, 2001; Serruya et al., 2002; Andersen et al., 2004).

In the late seventies, extracellular recordings started to be done in the human brain (Halgren et al., 1978) for clinical reasons. These studies showed that high-level representation areas like medial temporal lobe (MTL) (Fried et al., 1997), respond to different images of faces, animals, objects or scenes (Fried et al., 1997; Kreiman et al., 2000). Recent investigations, also recording the activity of single cells in the human MTL, showed neurons very selective activated by pictures of given individuals, landmarks or even by letter strings with their names, suggesting a role of MTL neurons in the transformation of complex visual perceptions into long-term and more abstract memories (Quiñones Quiroga et al., 2005).

The analysis of single-cell responses given a particular cognitive process, such the ones described above is done by first isolating the observed action potentials (spikes) into the responses of different putative neurons. The spike shapes associated with each neuron depend mainly on its morphology and its spatial location with respect to the recording electrode (Gold et al., 2006); therefore an electrode will be able to record in principle different spike shapes for each neuron in the surroundings. Spike shape features are extensively used for the classification of single-cell activity. The grouping of spikes shapes into clusters

based on the similarity of their features is called spike sorting (Lewicki, 1998; Quian Quiroga, 2007).

As described by Buzsaki (2004) the number of sorted cells (up to 20) is much smaller than the number of theoretically recordable neurons (~ 1000). The remaining cells may be silent neurons (Shoham et al., 2006), which only fire when very specific stimulus are presented, like images of individuals (Quian Quiroga et al., 2005, 2008) or when an animal is in specific locations in an environment ("place cells") (O'Keefe, 1976, 1979; Brun et al., 2002). These neurons could also be dead or damaged due to short and long term effects of electrode insertion on surrounding tissue (Rousche and Normann, 1998; Claverol-Tinture and Nadasdy, 2004) or remain undetected due to limitations in the current spike sorting algorithms (Pedreira et al., *under revision*). Improvements in electrode and spike sorting technology are therefore an indispensable step for a rigorous study of the role played by single neurons in different cognitive processes.

Advances in microelectrode technologies allow us now to analyse the activity of hundreds of neurons using extracellular multielectrode recordings (Nicoletis and Ribeiro, 2002). This, together with the increasing storage capacity, gives the opportunity to process massive amounts of data produced by large ensemble of neurons during large periods of time. The tremendous capacity of these new technological platforms for the analysis of large-scale data sets will be fully optimised with the use of unsupervised spike sorting algorithms. Automatic clustering will definitely increase the speed of the detection and sorting of single-units and at the same time will improve the outcome performance avoiding the errors due to subjectivity in the manual intervention.

In recent years, significant improvements have been made in the development of spike sorting algorithms. One relevant example is Wave_clus (Quian Quiroga et al., 2004), a widely used method which has been proved to identify neurons with very sparse activity in the human MTL (Quian Quiroga et al., 2005, 2008). This algorithm does not depend on assumptions about the structure of the data and has good performance even with few spikes per cluster, as in the case of

sparse neurons. Although Wave_clus provides a first unsupervised solution, some neurons have been only identified by further intervention of an expert operator, changing sometimes the quality of the data depending on which person performed the analysis. In large-scale recordings using multiple electrodes, supervised spike sorting turns to be a very time consuming task and at the same time increments the human errors, decreasing the optimal performance of the spike separation (Harris et al., 2000) with the number of channels. An efficient automatic spike sorting algorithm is crucial for the analysis of large-scale population of neurons contributing also to the development of optimal on-line spike sorting frameworks, where manual intervention is not possible.

The development of robust tools for the analysis of extracellular recordings strongly depends on the quality of the test and measurement platform used. The development of simulation environments providing realistic and challenging scenarios are crucial for the optimal design of spike sorting algorithms. Features such as coloured background noise, multiunit activity, sparse single-unit activity or electrode drifts are challenging components typically found in real extracellular recordings. Synthetic data sets able to generate this type of features would provide with very realistic frameworks in which develop robust unsupervised spike sorting algorithms. Moreover, the combination of both simulations and automated analysis tools might be used to create platforms that allow us to study for example the effects of different recording conditions in neural population coding (Quiñan Quiroga, 2009) or electrode designs (Robinson, 1968).

1.1 Organization of the Thesis

The organization of this thesis is as follows. In part I, I present the motivations that lead to the development of this research and I introduce relevant matters to this work. Part II is dedicated to expose the results obtained in this thesis. In particular, In Chapter 3 I present the development of a new simulation platform in which I tested a new solution for unsupervised sorting, described in Chapter

4. Finally, in Part III I summarise the discussions of the corresponding findings and I describe the suggestions for future work.

Chapter 2

Theoretical Background

In this chapter I introduce the theoretical background concerning the research work described in this thesis. In this introduction, I cover some of the key areas relevant to the automatic classification of the extracellular activity of single-cells. I first present a brief overview of the most common techniques used to explore the brain activity as well as some of their main characteristics and limitations. I later focus on some of the key aspects of the extracellular potentials and their role in the classification of single-cells as well as the technology used for recording from individual neurons. Finally, I introduce the concept of spike sorting and I describe the methodology as well as its crucial role in the analysis of the activity of individual neurons.

2.1 Techniques to Study the Brain Activity

There are several methods to explore the brain activity. Each of the existing techniques is characterised by a specific temporal and spatial resolution. In experiments on a neuronal network level, the ideal smallest temporal resolution would correspond to the time scale of an action potential which is in the order of milliseconds. On the other hand, the desired spatial resolution depends on the type of scientific question that one is interested to address.

At the moment, none of the current methods is capable of providing the

optimal spatiotemporal resolution to study the brain at all topographical levels. Therefore, different methods are used (or a combination of them) depending of the level of resolution required in each study.

One of the most popular techniques used by Neuroscientists is the Electroencephalography (EEG). This method measures the electrical activity of the brain by placing multiple electrodes on the scalp of the subject. Each EEG electrode records the average extracellular activity of a cortical area on the order of 10cm^2 (Buzsaki, 2006). The spatial distribution of the electrodes, covering all the areas of the cortical surface, together with its high temporal resolution (within the millisecond scale) makes this technique very attractive to study the simultaneous activation of different areas during the course of an experiment. However, EEG recordings do not have enough spatial resolution to study what single neurons do.

A similar technique, measuring in this case the magnetic fields of the brain instead of its electrical activity, is the Magnetoencephalography (MEG). This method records the activity of the brain from outside the skull and therefore, as in the EEG, it is also a non-invasive technique although in this case, no electrodes are placed on the scalp of the subject. Recording from the magnetic activity of the brain provides a better spatial resolution ($<1\text{cm}$) than the obtained from the electric fields since the magnetic signal is not scattered by the inhomogeneities of the skull and the scalp (Hamalainen et al., 1993). The MEG has also the same good temporal resolution as the EEG, offering in principle a better spatiotemporal compromise. However, despite of the good spatial resolution is still not accurate enough to study brain computations at the single-neuron level.

Probably the most popular non-invasive technique to study the brain in a cognitive level is the functional magnetic resonance imaging (fMRI). fMRI measures tissue perfusion, blood-volume changes and changes in the concentration of oxygen in the brain. When neurons fire, the blood flow in the nearby regions increases as a response to increments in their energy consumption. These biophysical phenomena are used as an indirect technique to record from the neu-

ronal activity. A widely used measure in fMRI to quantify this phenomenon is the blood-oxygen-level dependence (BOLD) contrast mechanism (Ogawa et al., 1990). Among the principal advantages of using the BOLD signal are its non-invasive nature and the possibility to record from the activity of the entire neuronal networks in different brain areas at relatively high spatial resolution. Studies in animals, reported resolutions as high as $0.25 \times 0.25 \text{ mm}^2$ with a slice thickness of 2 mm (Goense et al., 2007). However, using the fMRI signal presents also some limitations which might be crucial depending on the type of research question to be addressed. Recording from the BOLD signal, with peaks obtained after a couple of seconds of the start of neuronal spiking activity (Logothetis, 2008), introduces important delays if one wants to study how information is processed and propagated through the neuronal substrate. This limitation will not allow the study of particular cognitive processes in which their temporal resolution is lower than the response of the BOLD signal.

At this point, it is important to remark that methods such as EEG, MEG or fMRI, despite of their particular spatiotemporal limitations, still provide a major advantage: the non-invasive access to the human brain. This, of course, has great relevance in a scientific level but also tremendous implications in clinical neuroscience and medical diagnostic in areas such as epilepsy or autism. Although, these techniques provide valuable insights about the local neural activity underlying different brain functions, they do not provide direct information about how individual neurons and their collective behaviour give rise to cognitive process. It is therefore necessary to use methods capable of measuring the spiking activity of single cells such as optical imaging or extracellular recordings.

During recent years, optic methods have been increasingly used in neuroscience. They provide the best spatial resolution, up to the micrometer scale, while covering at the same time relative large areas of the cortical surface. With this technique, the detection is done by sensing changes in the optical properties of the brain tissue using photon-detecting arrays (Denk et al., 1990). The temporal resolution usually depends on the molecular transducer, used to trans-

form changes in the membrane potentials into optical signals, and also on the photodetection device. Up to date, the best optical technique is the recent two-photon laser scanning microscopy (2-PLSM). The method uses a powerful laser which is capable of recording signals from deeper layers into the neocortex using a microscope that captures the photons generating an accurate three dimensional image (Grinvald and Hildesheim, 2004). Despite its powerful spatiotemporal resolution, the 2-PLSM is still being developed to obtain a high enough temporal resolution to identify single neurons. Moreover, the 2-PLSM signal, as in some of the techniques mentioned earlier, does not provide a direct measure of the spiking activity of the local neurons and in some cases it is difficult to interpret the optical output and isolate its neuronal nature from other non-neuronal sources of fluorescence.

Recording from individual neurons using intracellular electrophysiology allows the analysis of neural activity with high spatiotemporal resolution. In this method, the measurements are typically done in brain slices or anesthetized animals. This method allows the study of the neural activity at different topographical levels, from single cells to neuronal networks. Moreover, using this technique it is possible to inject different drugs into a particular cell and study the effects of different chemicals on the local neuronal circuits. Also, it is possible to establish relationships between a type of cell (i.e. pyramidal, interneuron) and its corresponding role in a particular neural computation. However, despite the tremendous potential of this technique and extensive ongoing research, intracellular electrophysiology only obtains signals from one cell at a time and usually does not allow the study of the neural mechanisms of cognitive processes during behaviour or connectivity with other areas of the brain.

To date, the most widely used method to record directly from groups of individual neurons in action is with extracellular electrophysiology (Hubel and Wiesel, 1962). Extracellular electrophysiology consists in sensing the action potentials of neurons near their cell body by using a voltage-sensing microelectrode, which consists of a microwire in the scale of micrometers insulated except in the

tip (Hubel, 1957). This intracortical probe records an electrical signal with contributions from mainly two different sources of neural activity: i) the Local Field Potentials (LFPs) and ii) the spiking activity of neurons. LFPs are represented by the slow fluctuations of the recorded signal, typically associated with the input activity of neurons nearby the electrode (Buzsaki, 2006). On the other hand, the action potentials represent the output activity of the cell. In the next sections we describe in more detail the role of the extracellular signals for the study of the brain activity.

2.2 Neuronal Action Potentials

Extracellular action potentials convey a great amount of information for the study of neural computation by single-cells. The timing at which these extracellular spikes occur, together with their frequencies of discharge (firing rate) form the base of a numerous scientific discoveries about how populations of individual neurons codify a specific cognitive process (Averbeck et al., 2006; Quiñan Quiroga and Panzeri, 2009). Spike times are used to quantify how much information about a certain stimuli is encoded by single-neurons (Rieke et al., 1997). Depending on the type of coding, this information might be carried by a significant change in the number of spikes fired within a particular time window, independently of when they occur (Adrian, 1928; Rieke et al., 1997) or by a particular time pattern of each spike train (Optican and Richmond, 1987; Victor and Purpura, 1996) or even by irregularity modulations in the firing rates (Softky and Koch, 1992; Shadlen and Newsome, 1998).

The precise timing or frequency of a spike train are essential for the analysis of single-cell activity. Nevertheless, the waveform of these action potentials is at the same time important for the identification of different units present during an extracellular recording. It is assumed that, near the tip of an intracortical microelectrode, each neuron exhibits action potentials whose waveforms differ in amplitude and shape. These differences are mainly due to the position and

geometry of the recording electrode and the spatial distributions of currents in the neuron (Fee et al., 1996b; Gold et al., 2006).

If one considers a cell as a point of charge and the extracellular medium as an isotropic volume conductor of resistivity ρ , the extracellular potential Φ propagates with an inverse r -law following the Coulomb's law

$$\Phi = \frac{\rho I}{4\pi r} \quad (2.1)$$

where I is a point source of current and r is the distance from the cell to the recording electrode (Plonsey, 1969; Holt and Koch, 1999; Gold et al., 2006). However, at distances near the cell body, this amplitude-distance relationship might not be consistent with the experimental results presented by Henze et al. (2000) in which simultaneous intracellular and extracellular recordings showed a rapid decrease in the spike amplitude with distance from the neuron. Moreover, the basic theoretical approach mentioned earlier does not explain the waveform shape variability also found in these types of experiments and it does not include any knowledge about the morphology of the neuron or possible inhomogeneities within the extracellular medium (Gold et al., 2006).

Following the existing knowledge about dendritic neurons (Ramon y Cajal, 1909), the effects of the cell morphology on the shape of extracellular potentials were first studied by Rall (1962) using mathematical models to analyse the distribution and interaction of currents in soma and dendrites.

Further works on the electrophysiological properties of extracellular potentials, identified different sources of spike waveform variability, extrinsic and intrinsic to the neuron (Fee et al., 1996b). Extrinsic variability is associated to the background noise produced mainly by distant neurons and other sources of high-frequency electrical activity such as fast synaptic currents (Farrant et al., 1994). On the other hand, the sources of intrinsic variability are believed to be associated for example to spatial distributions of currents in the soma and in the dendritic tree of the particular cell, electrical interactions between nearby neu-

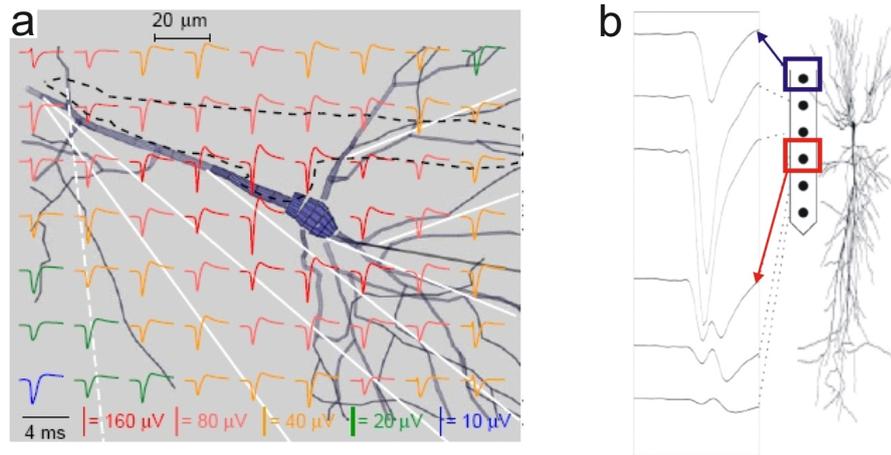


Figure 2.1: Examples of variability in the recorded spike shapes of a single neuron. (a) Computational model of a dendritic pyramidal cell in which the extracellular waveform generated in the soma has been computed at different locations within the extracellular space (Gold et al., 2006). (b) Real extracellular recordings from the activity of a pyramidal cell showing the variability of the wave shape in each one of the recording channels (Harris et al., 2000).

rons or burst firing activity (McCormick et al., 1985; Henze et al., 2000). These electrophysiological phenomena were analysed by means of highly-detailed computational models of dendritic neurons in order to study the origins of the extracellular waveforms and its relation with the intracellular action potentials (Holt and Koch, 1999; Gold et al., 2006). These powerful mathematical models allow the simulation of many types of neurons with different patterns of dendritic branching. These synthetic neurons are typically enclosed on an extracellular space, described as a purely ohmic conductor volume (Plonsey, 1969; Holt and Koch, 1999) in which the action potential waveforms are computed on different locations within this medium (Figure 2.1a).

2.3 Extracellular Recordings

In single-cell recording experiments, the intracortical electrode placed in the area of interest, records the activity of all the spiking neurons in the vicinity of the electrode tip. The electrode impedance is the main feature for single neuron recordings and it is determined by the material of the electrode and its diameter.

The smaller the diameter, the higher the impedance and the smaller the radius of the cylinder from which the electrode can obtain recordable cells.

In some experiments when small electrodes are needed (i.e. to minimize tissue injury during the insertion), to lower its impedance the tip has to be plated with for example platinum black or gold (Robinson, 1968). In a typical electrode with a diameter of approximately 12-25 μm , the activity of a single neuron is recordable at distances up to 50 μm away from the tip (Buzsaki, 2004). A sphere with such a radius contains approximately 140 neurons, which is the upper theoretical limit of recordable neurons for a single electrode. From 50 up to approximately 140 μm , the single-cell activity recorded by the electrode has low levels of signal-to-noise ratio that makes it not possible to identify the differences in the spike shapes of each individual neuron. These spikes are known as multi-unit activity and also contain valuable information about local computations (Stark and Abeles, 2007).¹

In recent years, the electrode technology has considerably improved, moving more and more towards probes with multiple channels capable of performing massive parallel recordings from large ensemble of neurons. One of the most popular probes is the Utah Array, developed by Rousche and Normann (1998) with 100 silicon microelectrodes. The distribution of the channels in a matrix of 10×10 makes this type of probe suitable for recording from neurons located in a single layer, approximately up to 1.5 mm beneath the cerebral cortex. Another popular multiarray of microelectrodes is the Michigan Probe (Drake et al., 1988; Hetke et al., 1994). This electrode, with channels distributed lengthwise along the shank allows the simultaneous recording of neurons within cortical columns. This type of probe is able to record extracellular action potentials from different parts of the same neuron to study, for instance, the relationship between the recording sites and the spike waveform (Henze et al., 2000).

¹Ison et al., (to be submitted) presented a quantitative approach to measure the differences between single- and multi-unit activity comparing the classification performance of a set of experimental recordings using several features including the first moments of the spike shape amplitude distribution or the proportion of ISI smaller than 3ms.

After an extracellular recording is performed, a second step, which includes the separation of the single-cell activity from the background noise, has to be done using spike sorting algorithms. As mentioned earlier, from the spike data, those waveforms generated by far-away neurons will typically present a lower signal-to-noise ratio and will be grouped as multi-unit.

2.4 Isolation of Single Neuron Activity

A crucial step in the analysis of spike trains is the correct isolation of the activity of individual cells. Due to the arguments described in Section 2.2, the characteristic spike shape of each neuron provides specific features that allow their separation. These differences in the spike shapes are used by spike sorting algorithms to classify the detected extracellular action potentials into groups and identify the activity of the different putative neurons present during the experiment (Quiñan Quiroga, 2007). Figure 2.2 shows an example of an extracellular recording scenario (left) in which one multi-unit and two single units of similar amplitude, detected from the continuous data (right-top blue plot), present differences in their spike shapes which are used by spike sorting algorithms for classification.

Under some conditions, the variability of the spike shapes might turn into a major cause of sorting errors since different recorded extracellular waveforms might be generated by the same neuron. Some cells, during states of bursting activity, may produce spikes with significant differences in their amplitudes (McCormick et al., 1985; Henze et al., 2000). In some experiments especially in chronic implants on freely behaving animals; possible electrode drifts might place the recording channel on different locations within the same neuron. Simultaneous intracellular and extracellular recordings showed how extracellular spikes generated by the same neuron may dramatically change, not only their amplitude but also their shapes if recorded from different locations (Harris et al., 2000; Henze et al., 2000). Figure 2.1b shows the differences between the spike

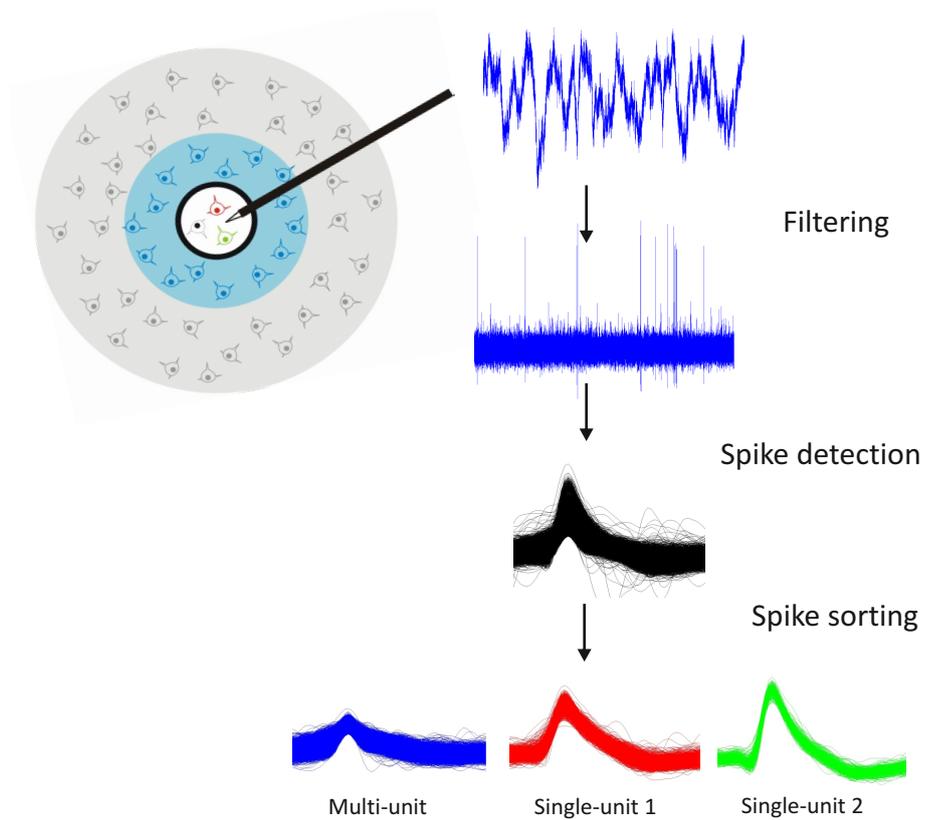


Figure 2.2: After amplification and filtering, the extracellular recording shows the spiking activity of single units on top of the background noise (right-top blue plot). After detection, the multi- and single-unit activity is identified using spike sorting (bottom).

waveforms recorded on each one of the sites of a multichannel electrode placed nearby a pyramidal cell (Harris et al., 2000).

Another frequent source of classification errors might occur when similar neurons at similar distances from the recording electrode generate spike shapes with very small differences. This might lead to grouping the activity of different neurons as if generated by only one. This contaminated activity, containing spikes from different neurons, might be solved by calculating the inter-spike-intervals of the times of this cluster of spikes and looking for well-defined refractory periods. However, such measurements might not be optimal if, for example, cells fire at low rates or in the presence of relatively high number of misses during detection (Ison and Quian Quiroga, 2007).

A solution to reduce errors in the classification of single neurons is to record the activity of the same cell from different locations simultaneously. The same spike is recorded from different channels with amplitudes varying as a function of distance from the cell. This relationship provides additional information about the spatial location of the neurons (Blanche et al., 2003; Chelaru and Jog, 2005) that helps spike sorting algorithms to obtain better identification results.

This multi-electrode approach was first adopted by McNaughton et al. (1983), with the use of an electrode including two adjacent recording sites made with two twisted wires called stereotrode. Following this approach, further improvements were made by Gray et al. (1995) introducing two more channels. The new design named with the term tetrode (O'Keefe and Recce, 1993; Wilson and McNaughton, 1993) reduced the classification errors to typically half those obtained with single-wires Harris et al. (2000). Another way to extract benefit from multi-electrodes, commonly known as polytrodes, is by building a *polyspike* concatenating the spikes of each recording site. This new *polyspike* waveform contains additional features that might be relevant for cluster separation when used by spike sorting algorithms.

2.5 The Spike Sorting Method

Spike sorting methods can be divided typically into four steps: filtering, spike detection, feature extraction and clustering (Figure 2.3). Band-pass filters within the range of 300-3000Hz are often used in order to eliminate low and high frequency components of the background noise. The detection is commonly done using amplitude thresholding, for example, estimating the standard deviation of the noise and setting the threshold to a certain number of times over this level.

After spike storage (the number of sample points for a spike must be decided) and alignment (every spike must have the point of maximum amplitude in the same sample), the third step is to extract relevant features of the detected spike shapes. The idea is to reduce the number of dimensions and take only those features that best separate the spikes. The most popular features are the peak amplitude, the spike width, the energy or a particular data point within the spike shape. Although very basic and not optimal for spike separation, these features are still widely used by several neurophysiology laboratories. A more advanced technique is principal component analysis (PCA). This method selects the directions of maximum variation of the data, but these are not necessarily the ones that give the best separation (Quiñan Quiroga, 2007). Other feature extraction method is wavelets, which do signal decomposition in time and frequency (Mallat, 1989). In this case, each spike shape is represented by a number of wavelet coefficients that capture different frequency features of the spike localizable in time. After wavelet analysis, one should eliminate those coefficients representing noisy features and keep only those ones that best separate the signal (i.e. those ones presenting multimodal distribution).

The last step is to group the features of the detected spikes into classes. Clustering methods can be classified into supervised, semi-supervised or fully unsupervised, depending on the degree of human intervention. In supervised clustering, the operator selects the clusters by, for example, drawing ellipses which enclose the data points on a multi-dimension feature map. In semi-supervised

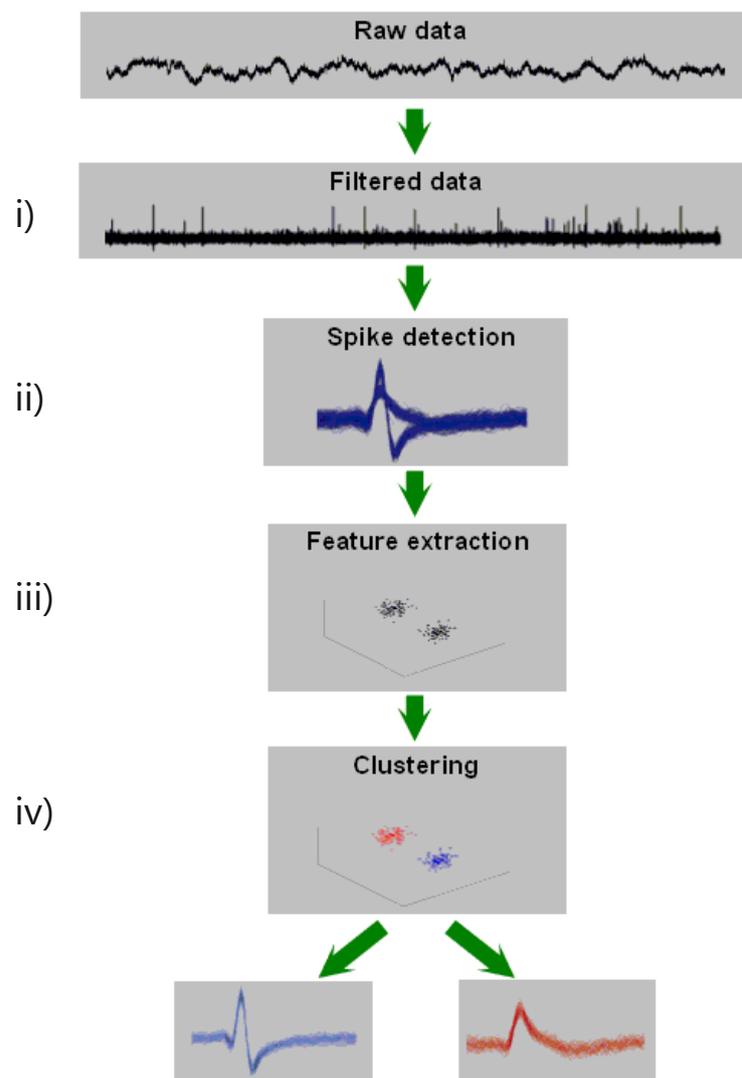


Figure 2.3: Description of the four steps of a spike sorting method: i) filtering; ii) spike detection; iii) feature extraction and iv) clustering. Adapted from Quian Quiroga (2007).

approaches, the algorithm typically provides automatically a preliminary selection of clusters. This solution is further optimised by the operator who makes the final decision by merging or rejecting clusters to obtain the optimal classes. Finally, the third type of approach is the unsupervised selection of clusters in which the algorithm does not assume any human intervention and provides a fully automatic identification of the optimal classes (Quiñan Quiroga et al., 2004).

In recent years, the automatic classification of single-units is becoming an essential part within the analysis of extracellular recordings. The new developments in multielectrode technology with an increasing number of recording sites per probe provide nowadays scientists with the opportunity to record from large-scale ensembles of neurons at a time (Nicolelis and Ribeiro, 2002; Csicsvari et al., 2003). In such high-density recordings, the manual supervision of each single channel might turn into a very time consuming task. Moreover, it is well known that the subjectivity introduced by the human intervention creates an additional source of sorting errors (Harris et al., 2000; Wood et al., 2004b) which increases with the number of recorded channels. The use of unsupervised spike sorting algorithms is therefore crucial to fully reach the potential of these massive parallel recordings. Moreover, these techniques will considerably improve the performance of on-line applications such as in the case of brain machine interfaces (Nicolelis, 2001; Serruya et al., 2002; Musallam et al., 2004), where manual intervention is not possible.

Many spike sorting algorithms have been developed (see Lewicki (1998) for a review). Some of these methods are based on Bayesian statistical frameworks, relying in some cases on a Gaussian model of the distribution of the spike waveforms (Lewicki, 1994; Harris et al., 2000; Wood et al., 2004a). Although the amplitude of the background noise is mainly dominated by a Gaussian distribution, the data clusters might not follow the same structure. Past studies, showing the non-Gaussian variability of the spike shapes and the non-stationarity nature of the extracellular recordings due for example to small electrode drifts (Fee et al., 1996b; Snider and Bonds, 1998) or the presence of bursting cells (McCormick

et al., 1985; Henze et al., 2000), have also motivated the development of different classification algorithms taking in consideration more realistic models of clusters (Shoham et al., 2003; Bar-Hillel et al., 2006; Wood and Black, 2008). Other solutions are based on nonparametric approaches using classification techniques based for example on nearest-neighbours interactions. One such a method is Wave_clus (Quián Quiroga et al., 2004), an unsupervised spike sorting algorithm using wavelets and superparamagnetic clustering (SPC) (Blatt et al., 1996, 1997; Domany, 1999).

2.6 Wave_clus

Wave_clus is a spike sorting method which includes the four main processing blocks described above: filtering, detection, feature extraction and clustering.

Filtering is done by using a bandpass zero-phase forward and reverse elliptic filter (second order) within the range of 300-3000 Hz. The detection is performed by estimating the standard deviation of the noise based on the median of the absolute value of the signal (Donoho and Johnstone, 1994), and setting the threshold detection to five times the value of the noise estimate. The median was used instead of the mean in order to reduce the influence of the spikes in the estimation of the noise level. The threshold Thr used in Wave_clus is as follows:

$$Thr = 5 \cdot \sigma_n, \quad (2.2)$$

$$\sigma_n = median \left\{ \frac{|x|}{0.6745} \right\} \quad (2.3)$$

where σ_n is an estimation of the standard deviation of the noise, and $|x|$ the absolute value of the bandpass filtered signal. After detection, 64 points (i.e. ~ 2.5 ms) are stored after alignment for further analysis. Due to limitations in the sampling frequency, it is possible to have misalignments between spikes, since the maximum point can be located in different parts of the spike, introducing distortions during the clustering process. In order to avoid this effect, the spike

shapes are realigned after interpolating by a factor of 2 using cubic splines.

Feature extraction is done using a four-level multiresolution decomposition with Haar wavelets, obtaining 64 wavelet coefficients for every spike shape. The selection of the coefficients that best separate the different spikes is done by identifying those ones with distributions differing most from a normal distribution. This quantification is done using a Kolmogorov-Smirnov (KS) test for normality which compares the cumulative distribution function (CDF) of the wavelet coefficients $F(x)$ with a CDF of a normal distribution of the same mean and variance $G(x)$. For each coefficient its deviation from normality is quantified as

$$\max(|F(x) - G(x)|) \quad (2.4)$$

taking in this case the 10 coefficients with the highest deviation. After the extraction and selection of the optimal features, an unsupervised selection of clusters is done using superparamagnetic clustering (SPC).

The SPC is a nonparametric clustering algorithm based on the physical properties of a magnetic system. In this method the data is modelled as a Potts system (Wu, 1982) in which a spin state $s = 1, 2, \dots, q$ is assigned to each data point x_i with $i = 1, 2, \dots, N$. The spin-spin correlations between the K -nearest-neighbours are measured using Monte Carlo simulations (Binder and Heermann, 1988; Wolf, 1989) and then used to group the points into clusters (Blatt et al., 1996). In the first stage, the strength of the interactions between all pairs of points $\langle i, j \rangle$ is computed as

$$J_{ij} = \begin{cases} \frac{1}{K} \exp\left(-\frac{\|x_i - x_j\|^2}{2a^2}\right) & \text{if } x_i \text{ is nearest neighbour of } x_j \\ 0 & \text{else} \end{cases} \quad (2.5)$$

where K is the number of nearest-neighbours and a represents the average distance across all the neighbouring points. After calculating all the pair interactions, an initial configuration is created by assigning random spin values from

1 to q to each data point x_i . Then, for each Monte Carlo iteration, a data point is randomly selected and its spin state s_i is changed into a new state. The effect of this new situation on the rest of data points is quantified by calculating the probability that a change on s_i will also change the state of its nearest-neighbour s_j , computed as

$$p_{ij} = 1 - \exp\left(-\frac{J_{ij}}{T}\delta_{s_i,s_j}\right) \quad (2.6)$$

where δ_{s_i,s_j} is the spin-spin correlation between neighbouring points in each iteration and T is the temperature. Note that only neighbouring points with the same spin state and a high interaction strength will be able to change their state together. Finally, after computing p_{ij} and δ_{s_i,s_j} for M Monte Carlo simulations, the neighbouring points x_i and x_j will be grouped in the same cluster if $\langle\delta_{s_i,s_j}\rangle \geq \theta$. In this case, Wave_clus uses $\theta = 0.5$ as grouping threshold, $N = 500$ Monte Carlo iterations, $q = 20$ spin states and $K = 11$ nearest-neighbours.

As it is possible to infer from Equation 2.6, the temperature T plays a crucial role in how points are grouped into clusters. In this case, low temperatures will make all the neighbouring points change their state together (ferromagnetic phase) grouping them into a single cluster regardless the strength of their interactions J_{ij} . For high temperatures, the pair probabilities will decrease even for points within high density regions. In this case, the data will be grouped into many clusters including a very few number of points (paramagnetic phase). In the middle, there is the *superparamagnetic* phase, in which only neighbouring points within high density regions will be grouped in the same cluster.

In this case, the SPC algorithm spans a wide range of temperatures varying from 0 to 0.2 in increments of 0.01. For each temperature, a different type of data partition is generated with new clusters merging or breaking. The idea is to find the optimal clusters appearing at the optimal temperature.

In Wave_clus, clusters are automatically identified according to a thresholding procedure based on their size. The algorithm looks for the highest tem-

perature at which a cluster containing more than a certain number of points is found. The rationale is that as the temperature increases one expects the existing clusters to break into smaller ones. The cluster size is relatively related to the firing rate of the corresponding neuron and therefore, this threshold should be selected based on the minimum firing rate that one wants to consider. In the cluster assignment, with $T_i = 0.01, 0.02, \dots, 0.2$ where $i = 1, 2, \dots, 20$ is the number of the partition, the selected temperature T_{clus} will be the highest temperature at which

$$A_{clus}^j - A_{clus-1}^j \geq a \quad (2.7)$$

where A_{clus}^j represents the size of the j^{th} cluster at T_{clus} and a is the design parameter representing the increment in the number of spikes allowed for cluster assignment.

In Figure 2.4 we show a snapshot of the Wave_clus graphical user interface (GUI). In the top plot, a segment of the filtered extracellular signal is depicted. The red line shows the detection threshold using the expression in equation 2.2. In the middle of the figure we see a series of plots representing from left to right the projection into 2 (out of 10) wavelet coefficients and the clusters identified by Wave_clus in this particular case. In this example, one multi-unit (spike shapes in blue) and two single-units (in red and green) were automatically detected. The temperature diagram generated by SPC is depicted at the bottom-left corner, which depicts the size of each cluster (A_i^j) in which the data has been partitioned (y axis) function of the temperature T_i (x axis). Setting the cluster size threshold to 60 spikes, Wave_clus selected automatically $T_{clus} = 0.14$, identifying the partition in which the three clusters were found. Each coloured dot represents the corresponding clusters depicted in the middle plots.

In some cases, the optimal clustering solution might be reached as a combination of clusters appearing at different temperatures. In this situation, the single-temperature approach of Wave_clus is not able to select the optimal solu-

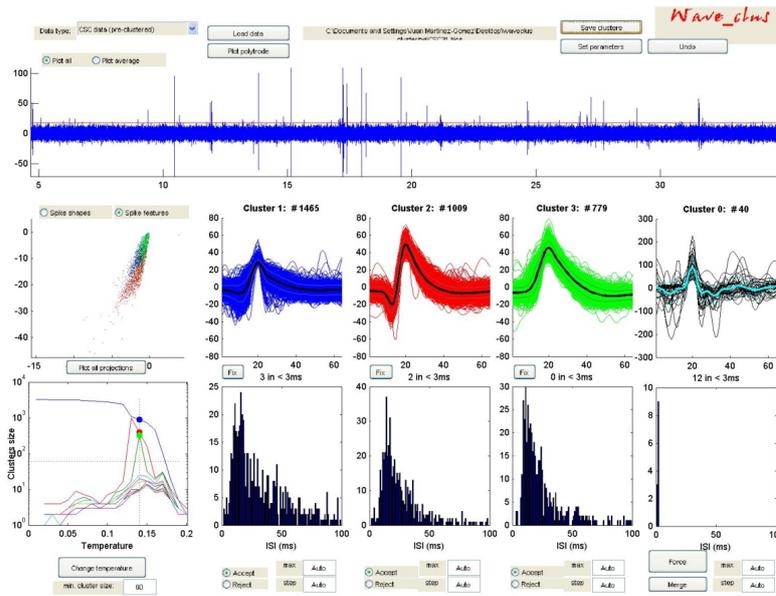


Figure 2.4: GUI of Wave_clus showing the outcome of the automatic spike sorting.

tion and a user intervention to manually select the right clusters is needed. In the GUI the operator can change the temperature, obtaining a different solution, selecting or rejecting the appropriate clusters. In Figure 2.5 we show an example of such scenario in which the user identified the single-units in red and green by selecting the temperatures $T = 0.03$ and $T = 0.12$, respectively.

2.7 Summary of Chapter 2

In this chapter we reviewed different techniques to measure the activity of the brain. Up to date, none of the current available methods is capable of providing a full spatiotemporal scale that allows the study of the brain at all possible resolutions.

Techniques such as EEG, MEG or fMRI, provide unprecedented information about the neural activity underlying a number of cognitive process. However, they are not able to give information on how this neural activity is created by the collective behaviour of individual neurons. On the other hand, promising methods such as optical imaging while having the best spatial resolution they do

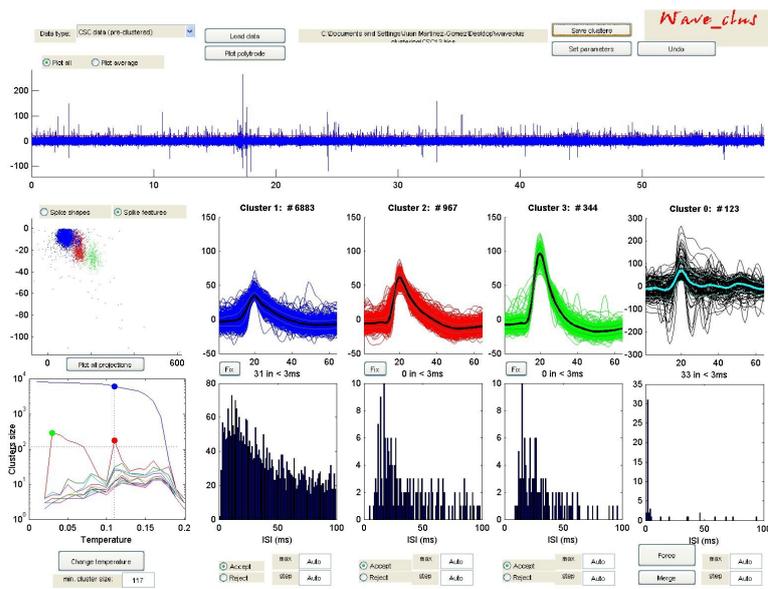


Figure 2.5: GUI of Wave_clus showing the outcome of the manual spike sorting.

not provide a direct measurement of the spiking activity of neurons.

Direct access to action potentials of groups of individual neurons *in vivo* is done using extracellular recordings. We described how extracellular activity is used for the study of single-cells and we reviewed some important aspects affecting the analysis of the extracellular signals such as the increasing advances in electrode technology and the identification of single-unit activity using spike sorting algorithms.

We emphasised the importance of using automatic methods for a fast and optimal classification of single-units as well as realistic simulation platform for testing these algorithms in order to take full advantage of the current powerful technology for high-density recordings.

In Part II, dedicated to expose the results of this work, we propose solutions that contribute to address the questions posed here. In Chapter 3 we present a new simulation platform to test spike sorting algorithms under realistic conditions. In Chapter 4 we make direct use of this synthetic data and we present a new solution for automatic classification of single-unit activity.

Part II

Results

Chapter 3

Realistic Simulation of Extracellular Recordings

3.1 Introduction

As described in Chapter 2 a well established methodology to study how mental processes are encoded by neurons in the brain is to analyse the extracellular activity of neurons using microelectrodes. In Section 2.5 we saw that this analysis is done by first detecting the spiking activity from the background noise and then grouping the recorded spikes into clusters based on the similarity of their shapes, a procedure known as spike sorting (see Section 2.5).

The quality of the resulting data depends crucially on the particular detection and sorting technique used, whose performance is usually hard to establish with real data. To quantify the performance of these methods, the use of synthetic data in which we know the original spike labels provides an adequate framework.

In principle, two different approaches might be taken to generate synthetic extracellular recordings. On the one hand, one could generate simple simulations by adding spikes to Gaussian noise (Lewicki, 1994). This approach is very fast and easy to implement but it misses important features of the real recordings that make spike sorting especially challenging. In particular, some characteristics

such as the presence of non-Gaussian distribution of clusters, multi-unit activity, and spectral similarity between noise and spikes might be critical issues to take into account in the analysis. On the other hand, detailed models of extracellular waveforms (Hines and Carnevale, 1997; Holt and Koch, 1999; Gold et al., 2006) could be more accurate but computationally too intensive when simulating extracellular recordings generated by a large number of neurons.

Here, we explored a compromise between both alternatives by developing relatively easy simulations that reproduce relevant real features of extracellular recordings for testing spike detection and sorting algorithms. In particular, our simulations will replicate the amplitude and spectral distributions of the noise, multi-unit and single-unit activity found in extracellular recordings. These simulations follow the approach presented by Quian Quiroga et al. (2004), where real spike shapes were used to create background noise and single-unit activity. Here, we introduced several improvements. First, we propose a biophysical strategy for generating the background noise. Second, we reproduced the multi-unit activity that makes the synthetic data very similar to real recordings. Multi-unit activity is commonly found in extracellular recordings and it represents the activity of several distant neurons whose spikes can be detected but are not large enough to be clustered because the difference in their shapes is masked by the background noise. This situation typically produces clusters with high variance and low signal to noise ratio. Third, we replicated actual values found in real recordings in terms of amplitude and power spectrum distributions. The resulting simulations also allowed the study of the spike sorting performance with sparse firing neurons (Hahnloser et al., 2002; Perez-Orive et al., 2002; Quian Quiroga et al., 2005, 2006), which, due to their very low baseline firing rates, might be masked by higher frequency firing neurons, making their identification particularly difficult. Our method also allowed an efficient generation of synthetic datasets, thus providing a simple strategy to generate large number of realistic experimental scenarios that are useful to quantify, for example, the performance and limits of different spike sorting algorithms.

3.2 Materials and Methods

We generated synthetic extracellular recordings that modelled the contribution of the background noise, multi-unit and single-unit activity. The simulations were created using a database with 594 different average spike shapes, taken from real recordings in monkey neocortex and basal ganglia. In order to reproduce conditions of real recordings -in which the peak of the spike might not occur during a recording sample, thus introducing misalignments- we used the following procedure: First, the data first simulated at a sampling frequency of 96 kHz. Then, using interpolated waveforms of the original spike shapes, the spike times were simulated to fall continuously between samples (Quiñan Quiroga et al., 2004). Finally, the data was downsampled to 24 kHz.

Figure 3.1 shows a summarised description of the method. Considering a recording electrode at the centre of the sphere, the extracellular recorded space was divided into three major areas: 1) Background noise; 2) Multi-unit activity and 3) Single-unit activity. The main idea of the method is to generate, in a computationally-efficient way, the relevant features -such as amplitude distribution and power spectrum- that allow reproducing realistic extracellular recording scenarios for the three zones described in the figure. The neurons in area 1 generate the background noise, whose activity was constructed based on the contributions of many individual point-source neurons. This strategy avoids complex models of single cells that might not have a relevant effect on the total signal at large distances from the recording electrode. The activity of neurons located at areas 2 and 3 was simulated using spike waveforms from the database and with amplitude distributions that follow real multi-unit and single-unit peak amplitude distributions, given that these values are accessible from real recordings. To consider the fact that different recording conditions can vary the levels of the background noise activity, we considered the amplitude values of the spikes of areas 2 and 3 relative to the detection threshold, which was in turn determined from an estimation of the standard deviation of the noise (see Equation 2.3).

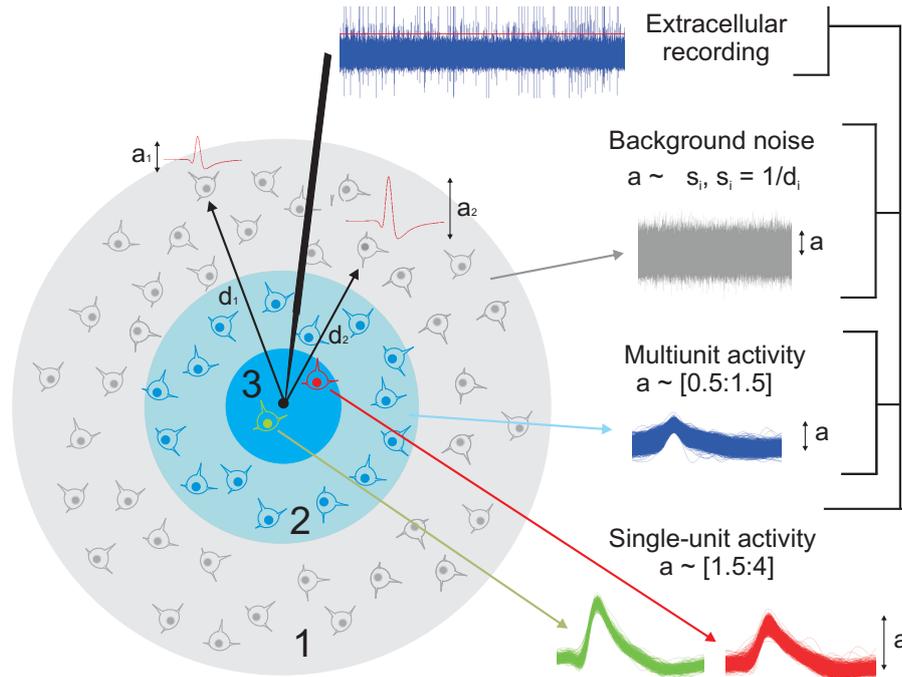


Figure 3.1: Overview of the simulation approach. Neurons in area 1 generate the noise activity, which is given by the superimposition of their spikes s_i . Each peak amplitude a_i is modulated by its distance d_i to the recording electrode. The activity of neurons in areas 2 and 3 (multi-unit and single-unit activity, respectively) is created by adding spikes on top of the background noise, following typical amplitude distributions a_i measured for multi-unit and single unit spikes.

Finally, the synthetic extracellular recordings were generated by superimposing the spikes corresponding to multi-unit and single-unit activity to the background noise, as detailed below. For this we used different single-unit amplitude levels, thus replicating the conditions of different signal-to-noise ratios.

3.2.1 Generation of Background Noise

The core of our approach to generate the noise activity was to simulate the contribution of distant neurons, as given in real extracellular recordings. This gives a hybrid simulation strategy, in which the characteristics of the noise activity, such as its amplitude and frequency distribution, arise naturally from the realistic biophysical process of its generation and it is not imposed beforehand. Moreover, this approach reproduces real noise characteristics without modelling the computationally expensive details of the contribution of each neuron. For

this, the far-apart neural activity was modelled by superimposing a large number of spike shapes (one spike per sample of generated data) placed at uniformly distributed random times. The amplitude distribution of all the waveforms was obtained by first assuming that each neuron generates an action potential of peak amplitude 1. These simulated neurons represent different point source charges distributed uniformly within a sphere of radius 1 (in adimensional units). Only far away spikes, at a distance larger than 0.5, were considered for generating the background noise.

The activity of close-by neurons was considered to contribute to the multi-unit and single unit activity (areas 2 and 3 in Figure 3.1), which was modelled following the typical amplitude ranges found in real data as described in the sections below. For a single spike shape s_i its peak amplitude a_i at the recording site (the centre of the sphere) was modelled to be equal to $1/d_i$ (d_i : distance of the cell i from the centre), following Coulomb's law (Rall, 1962). Alternatively, one could consider the neuron as a dipole with the soma and one dendrite represented by a negative and positive charge, respectively. In this case, the amplitude-distance relationship should be modelled as $1/d_i^2$. However, for large distances, we did not observe significant differences in the amplitude distribution generated by both approaches (data not published).

As mentioned above, this assumption leads to a large simplification in our simulation of the extracellular recordings, since we do not consider specific morphologies and physiological characteristics of distant neurons. The effect of other possible sources of electrical activity that might contribute to the noise power spectrum, such as electronic noise, axons, dendrites or synapse currents (Llinas, 1988; Farrant et al., 1994), was simulated by adding Gaussian noise to the noise generated by the spikes. Note that these simulations were generated to test spike detection and sorting algorithms, for which the extracellular signal is typically highpass filtered with a cutoff frequency of 300Hz or larger, to visualize the spiking activity. Therefore, we did not consider or try to reproduce sources of low frequency activity - i.e. below 300Hz as the ones considered for simulations of

local field potentials (Brunel and Wang, 2003; Bedard et al., 2006; Mazzone et al., 2008).

3.2.2 Generation of Multi-unit Activity

Multi-unit activity is produced by neurons located around 50 to 140 μm from the electrode tip (Buzsaki, 2004). The amplitude of the recorded multi-unit activity is close to the detection threshold. As shown by the blue spikes in Figure 3.1, the noise masks any difference between the spikes of different neurons, making the identification of single-units unfeasible. Spikes corresponding to multi-unit activity tend to produce large and typically non-Gaussian clusters. This is because clusters are created by the activity of more than one neuron and the amplitude distribution is truncated by the detection threshold.

In our simulations, multi-unit activity was created by mixing the activity of our whole database of 594 spike shapes using amplitudes uniformly distributed between 0.5 and 1.5 times the level of the detection threshold. Following previous criteria (Quiñones Quiroga et al., 2004) the detection threshold was set at 4 times the estimation of the standard deviation of the generated noise (see Equation 2.3). These amplitude ranges were chosen to follow the distribution of amplitude values for multi-unit activity found in real data (see Appendix A). For each unit contributing to the multi-unit activity the spike times were generated using a Poisson distribution with a mean firing rate of $20/N$ Hz (where N was the number of different units used to create the multi-unit activity), which summed up gives an overall multi-unit firing rate of 20 Hz.

3.2.3 Generation of Single-unit Activity

The single-unit activity (area 3 in Figure 3.1, spikes in green and red) was simulated using spikes of different shapes added to the background noise. The peak spike amplitudes were assigned with values between 1.5 and 4 times level of the detection threshold, that fall within the range of typical values found in real

recordings (see Appendix A). For each single unit, the variability of the spike amplitude was given by the background noise. The spike train of each single-unit followed a Poisson process, with a mean firing rate randomly selected between 0.5 and 5 Hz. Spikes that fell within the 2 ms following a previous one were removed -i.e. overlapping spikes were not considered.

3.2.4 Real Data

The real extracellular data presented here comes from a 30 minutes recording from the human medial temporal lobe. The subject was a pharmacologically intractable epileptic patient who was implanted with intracranial electrodes for clinical reasons. Besides the intracranial EEG contacts, each probe had a total of 9 micro-wires at its end, with 8 active recording channels and 1 reference to record single-neuron activity (Fried et al., 1997). The differential signal from the micro-wires was amplified and filtered between 300 and 3000 using a non-causal filter (Quiari Quiroga et al., 2004) to remove low frequency activity and high frequency artifacts. After filtering, the data was sampled at 32 kHz. The main features of this recording were similar to other 192 recordings from the human medial temporal lobe, in terms of the power spectrum and amplitude characteristics (see Figure A.1 in Appendix).

3.2.5 Spike Detection and Sorting

The real and synthetic extracellular recordings were analysed using Waveclus (see Section 2.6), an unsupervised spike detection and sorting algorithm (Quiari Quiroga et al., 2004). Spike detection was done using an amplitude threshold set to $Thr = 4 \cdot \sigma_n$ for this data (see Equation 2.3 for details).

Example Number	Amplitude	Firing rate (Hz)	α	r^2
1	4	1	1	0.991
2	4	5	1.18	0.994
3	2	5	0.9	0.984
4	2	5	0.93	0.984
5	3	0.5	0.97	0.988

Table 3.1: Simulation and spectral features for the 5 simulated data sets. The values of the amplitude are all relative to the detection threshold.

3.3 Results

We generated 5 sets of simulations of 2 minute length each with a sampling rate of 24 kHz. Each simulation contained the activity of one multi-unit and two single-units. The noise level σ_n was normalised to give a detection threshold of $28\mu V$, as typically used for real data (see Figure A.1 in Appendix). The peak amplitudes of the spikes generating the multi-unit clusters were uniformly distributed between 0.5 and 1.5 times the detection threshold for all the simulations. Each single-unit in examples 1 to 5 generated a spike train following a Poisson distribution. The values of peak amplitude and mean firing rates are described in Table 3.1. The different amplitudes (i.e. signal to noise ratios) and firing rates simulate different recording conditions to test spike sorting algorithms in varying regimes. The processing time to generate each example was around 9 minutes in an Intel Core 2 PC with a clock frequency of 2.4 GHz.

3.3.1 Real Data

Figure 3.2a shows 30 seconds of the continuous real data, which contains spikes with an amplitude between 30 and 130 μV . The estimation of the standard deviation of the noise (see Equation 2.3) was in this case equal to 6.5 μV and the threshold for spike detection was of 26 μV ; i.e. 4 times this value. For the same data, in Figure 3.2b we show the amplitude distribution of the signal, which followed a Gaussian distribution. Note that in the raw (band-pass filtered) data as well as in the amplitude distribution of Figure 3.2b there is no clear distinction

between the noise and the multi-unit activity (see inset), a typical challenging feature of real extracellular recordings. The black trace in Figure 3.2c shows the power density spectrum (PDS) of the data in Figure 3.2a (filtered between 300-3000Hz). For comparison, the curve in gray displays the non-filtered PDS. Both PDS traces were smoothed with a 3-datapoints moving average. The inset in Figure 3.2c depicts the log-log PDS within the red vertical lines. The plot in the inset of Figure 3.2b suggests that the PDS falls off following a power law distribution, with a $1/f^\alpha$, $\alpha = 1.1$ scaling dependence ($r^2 = 0.98$). Analyses performed on a large database of 192 channels of real recordings on the human medial temporal lobe, revealed similar PDS with r^2 equal to 0.98 ± 0.21 (mean \pm SD) and a mean r^2 coefficients equal to 0.992 ± 0.007 (see Appendix A).

From the full recording (of which the first 60 seconds are shown in Figure 3.2), three different clusters were automatically identified by using `Wave.clus` (see Section 3.2), as illustrated in Figure 3.3. The first cluster was classified as a multi-unit and the other two as single-units. Note the high variance and low amplitude observed in the multi-unit (left) as compared to the single unit clusters, a typical attribute of extracellular recordings.

3.3.2 Simulation of the Background Noise Activity

The noise activity was obtained by superimposing a large number of spikes, simulating the contribution of far-away neurons. For this, neurons were uniformly distributed inside a sphere and their spike amplitudes were inversely proportional to the distance to the centre of the sphere, which corresponds to the position of the electrode (see Section 3.2).

3.3.3 Matching of the Noise Amplitude Distribution

For the noise activity we did not consider point-source charges from relatively close distances d_i since the resulting amplitude of such spikes are in the range of the multi-unit and single unit activity. To have a rough approximation of the

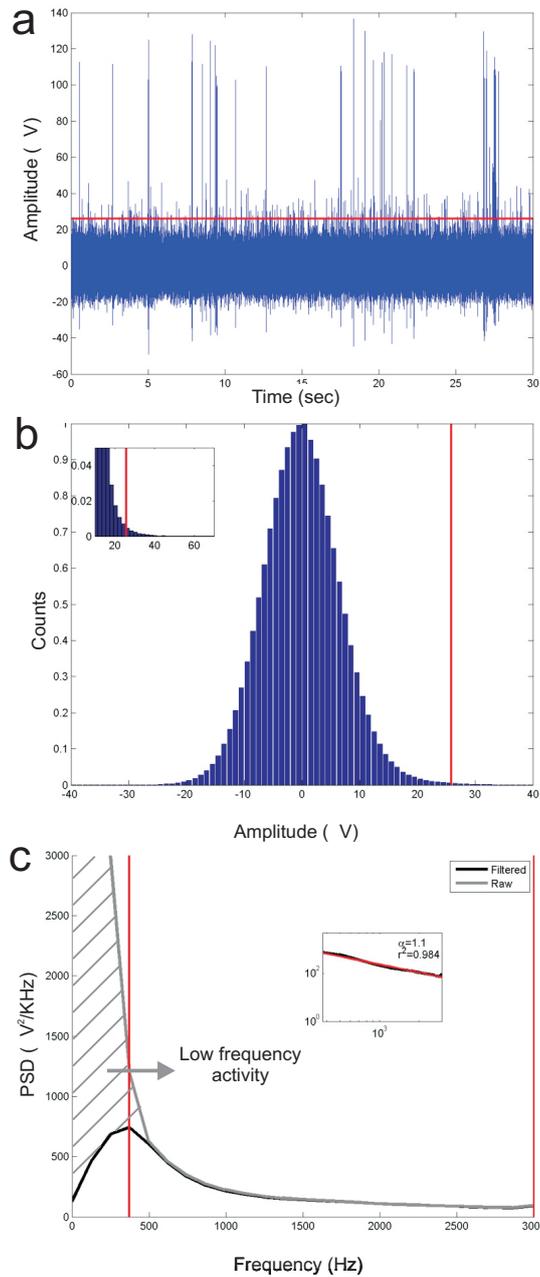


Figure 3.2: (a) 30 seconds of a continuous recording from the human medial temporal lobe. The voltage signal has been bandpass filtered between 300 and 3000 Hz. (b) Amplitude distribution of the voltage signal. The vertical red line represents the threshold for spike detection. The inset shows a zoom of the amplitude distribution around the detection threshold. (c) Power density spectrum (PDS) of the data in (a) with (black) and without (gray) bandpass filtering between 300-3000 Hz. The inset depicts the log-log PDS between the two red vertical lines. The red line in the inset indicates a linear fit to the data, which gave a slope $\alpha = 1.1$ and a correlation coefficient $r^2 = 0.984$.

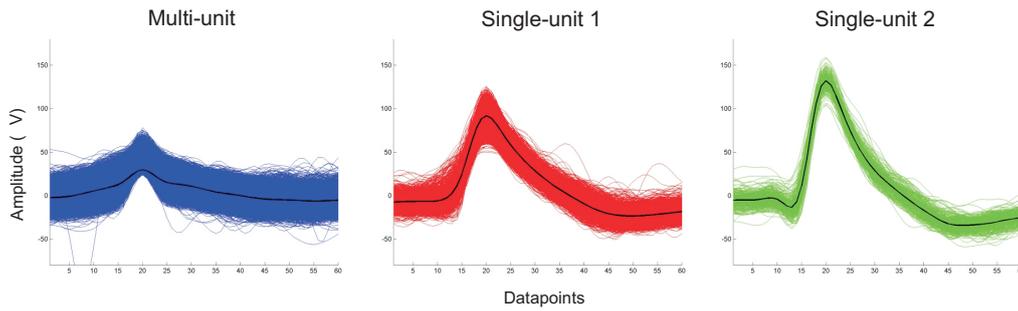


Figure 3.3: Clusters obtained after automatically spike sorting the data of Figure 3.2, using the whole duration of the recording (approx. 30 seconds). The cluster on the left corresponds to a multi-unit (in blue) and the other clusters correspond to two single-units (in red and green). For each cluster the thick black line marks the mean spike shape. Note the higher variance and non-symmetrical distribution across the maximum amplitude (datapoint 20) for the multi-unit compared to the single-units.

cutoff distance contributing to the noise, multi-unit, and single-unit activity, we evaluated the amplitude distribution of the simulated noise for 3 different cutoff distances ($d_i = 0.01$; 0.1 and 0.5, respectively). Figure 3.4 shows the simulated noise and the amplitude distributions generated using these 3 cutoff distances (Figure 3.4a - Figure 3.4c, respectively). We compared these simulations with a real continuous recording that did not have any visible single- and multi-unit activity (Figure 3.4d). As seen in Figure 3.4a, a low cutoff distance gives a noise distribution that clearly differs from the one of real recording, with too many large amplitude spikes. In fact, in this case there were $n = 285$ spikes crossing the detection threshold (see Equation 2.2), many more than the ones crossing the threshold for the real data ($n = 1$). On the other hand, considering only far apart action potentials (Figure 3.4b and Figure 3.4c) we obtained a more realistic representation of the continuous data and amplitude distribution of the synthetic background noise with few events above the detection threshold (16 and 4, respectively). Given these results, we chose a cutoff distance of 0.5 for simulating the noise activity. The processing time for creating these simulations was for all the examples approximately five minutes using a standard PC with a frequency of 2.4 GHz.

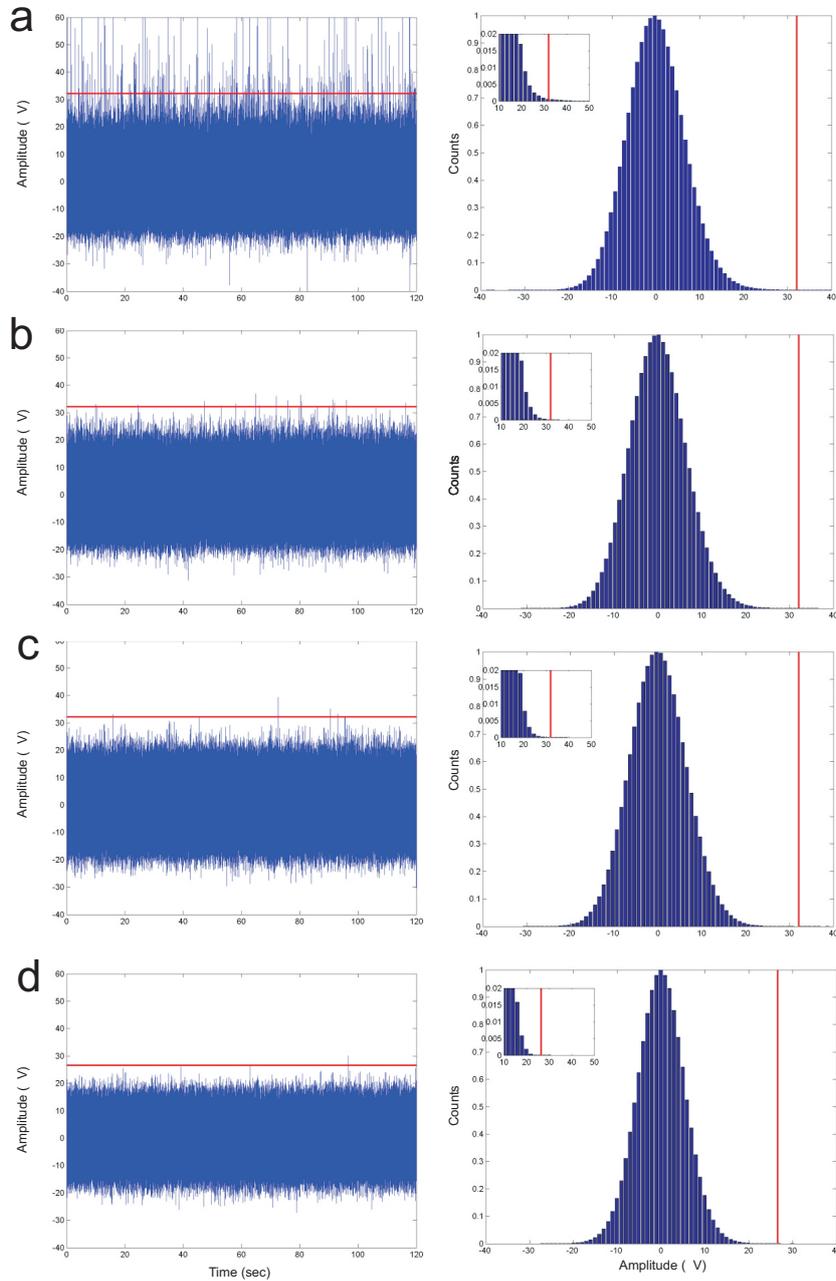


Figure 3.4: (a-c) Voltage signal (left) and amplitude distribution (right) of 3 synthetic noise signals for different values of d_i (from top to bottom, $d_i = 0.01$; 0.1 and 0.5, respectively; see text for details). The horizontal red line represents the detection threshold. (d) Noise from a real extracellular recording without spiking activity. For low values of d_i , (a) the amplitude distribution of the synthetic noise differs from the one observed in the real data.

3.3.4 Matching of the Noise Power Spectra

As mentioned before, a simple alternative to simulate the background noise of extracellular recordings is to use Gaussian noise, following the amplitude distribution discussed in the previous section. But, as shown in Figure 3.5a, after bandpass filtering between 300-3000Hz (as done for the real recordings), the power spectrum looks flat and clearly different to the spectrum found in real recordings (see Figure 3.2c). On the contrary, the noise simulation described in the previous section, based on a realistic simulation of the biophysical process contributing to the noise activity, gives a frequency distribution very similar to the one obtained in real recordings, as shown in Figure 3.2b. However, the value of the slope of a linear fit of the log-log plot of the power spectra gave a value of $\alpha = 1.8$, with $r^2 = 0.94$. The lower amounts of high frequency components as shown in the power spectrum as compared with real data (Figure 3.2c) might be because we are not taking in consideration other possible sources of electrical activity such as electronic noise or fast synaptic currents (Farrant et al., 1994), which could, in principle, be generated by an extra source of Gaussian noise, as the one of Figure 3.5a. In order to generate a more realistic synthetic noise that includes higher frequency components to reproduce the power law decay found in real data, we decided to combine the two approaches described below, adding to the noise of Figure 3.5b, in a proportion of 40% of Gaussian noise as the one of Figure 3.5a. Figure 3.5c shows the resulting noise characteristics, which lead to an increase in the high frequency activity, giving a PDS decay of $1/f^\alpha$, with $\alpha \sim 1$. Note that the voltage signal and amplitude distributions remained the same in all cases.

Another factor that influences the power spectrum and the amplitude distribution is the number of neurons contributing to the background noise. If this number is very low, the overall amplitude of the voltage signal generated by this spike activity will significantly decrease, increasing as a consequence the component of Gaussian noise. This, will also change the power spectrum, reducing the

slope α and make the PDS look flat and different to the one of the real data. In our model, we set the number of neurons contributing to the noise equal to the number of generated data samples (with each neuron firing only one spike during the simulation period). This level of neuronal activity gives a realistic noise signal in which very few spikes cross the detection threshold (Figure 3.5b). This, in combination with the gaussian noise, also generates a realistic PDS decay (Figure 3.5c). Note that a decrement in the number of superimposed spikes will modify the voltage signal, reducing the overall amplitude and standard deviation of the signal (i.e., a decrement of 50% in the number of neurons would result in a decrement of 25% in the standard deviation), increasing the spikes crossing the detection threshold. If this activity is very low (i.e., 0.1% of the generated samples) then, the power spectrum will be also seriously affected ($\alpha \sim 0.5$).

3.3.5 Addition of the Single-unit and Multi-unit Activity

Next, we added the single- and multi-unit spikes to the simulated noise. Note that in this case, rather than assuming a biophysical generation process, as for the noise, we experimentally have access to the amplitude and firing rate distribution of the single- and multi-unit data. In the following sections, we describe 5 simulated datasets reproducing different recording conditions that can be used as a platform to test spike sorting algorithms.

Figure 3.6 shows an example of one simulated data set with single-unit peak amplitude equal to 4 and a firing rate of 1 Hz. Note in the continuous recording showed in 3.6a the presence of small multi-unit spikes at the level of the detection threshold, as also observed in the real recordings (see Figure 3.2a). Moreover, in 3.6b there is a smooth transition between the background noise and the multi-unit activity around the detection threshold (see inset), as observed in the real recordings (see inset in Figure 3.2b). Figure 3.6c shows the PDS of this simulation, with a frequency dependence similar to the one of the real recording (3.2c). The values of $\alpha = 1$ and $r^2 = 0.991$ depicted in the figure (see

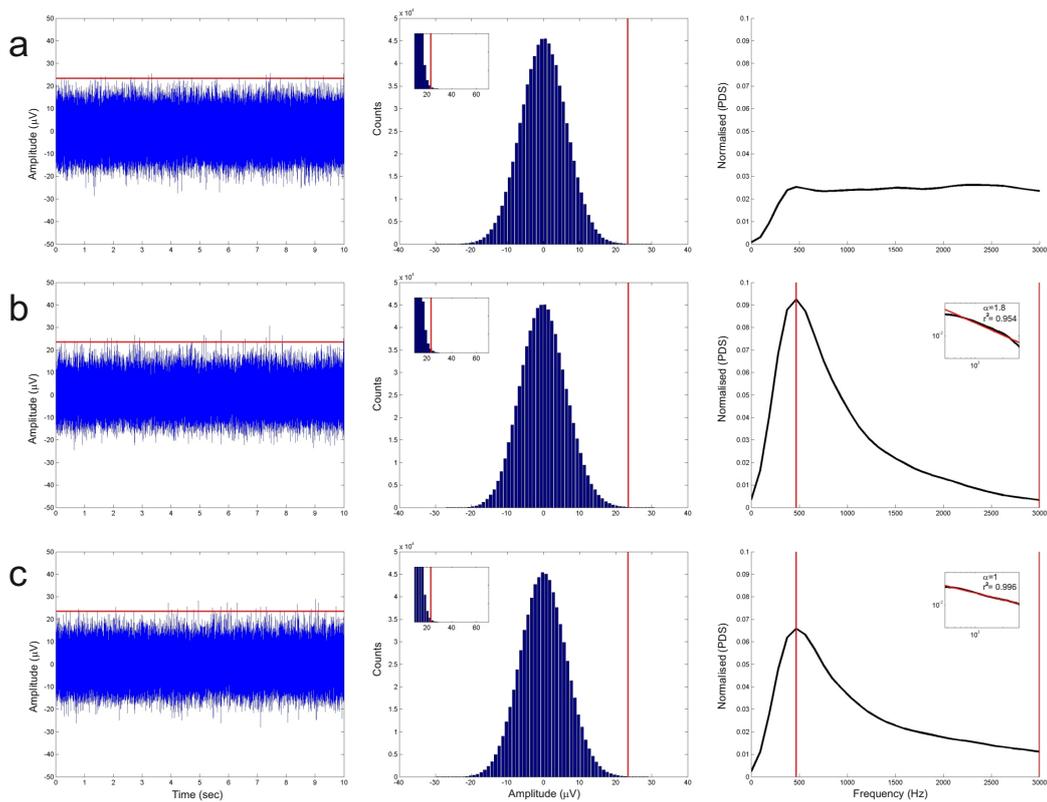


Figure 3.5: Continuous data (left), amplitude distribution (centre) and power density spectrum (PDS, right) of 3 synthetic noise realizations of 120 seconds each (only the first 10 seconds are displayed). (a) Gaussian noise. (b) Noise generated by the superimposition of spikes (see text for details). (c) Combination of (a) and (b). Note that the amplitude and frequency characteristics of the noise in (b) and (c) are very similar to the one of the real recording in Figure 3.2

inset) show that the simulation is able to reproduce major frequency features and also suggested a power-law distribution, as with the real extracellular recording situation (see 3.2c and Appendix A).

3.3.6 Testing of the Spike Sorting Algorithm

The results of the unsupervised spike sorting of the synthetic data of Figure 3.6 using `Wave_clus` are shown in Figure 3.7. As in the example with real data (Figure 3.3), the algorithm correctly identified all units generated in the simulation (one multi-unit and two single-units). The multi-unit cluster (left plot) has a wider spike shape due to the contribution of several units, as typically found in real recordings (see leftmost spike in Figure 3.3). Furthermore, for this cluster it is possible to see a non-uniform variance along the spike, which leads to non-Gaussian clusters found in real recordings. The presence of such non-Gaussian clusters are very challenging for many spike sorting algorithms that assume any particular model for the clusters. The middle and right plots depict the single-unit clusters, in which the amplitude variability is solely given by the background noise activity.

Table 3.1 and Table 3.2 summarize the characteristics of the 5 generated datasets and the performance of the unsupervised spike sorting on this data. The number of misses as well as the amplitude, firing rate and number of spikes, are referred to the simulated single-unit clusters. In Table 3.2, the number of spikes refers to the number of detections by the algorithm. In examples 1, 2 and 5, all spikes were correctly detected and classified. For simulations using single units of lower amplitudes, as in the cases of examples 3 and 4, the single-unit spikes tended to overlap with the multi-unit cluster, leading to false positives and misses. Note that in the case of example 3, where the spike shapes of the two single-units are similar, the number of misses and false positives increased due to some misclassifications. In Table 3.1, the last two columns represent the scaling exponent α and the correlation factor r^2 of the power spectrum for each

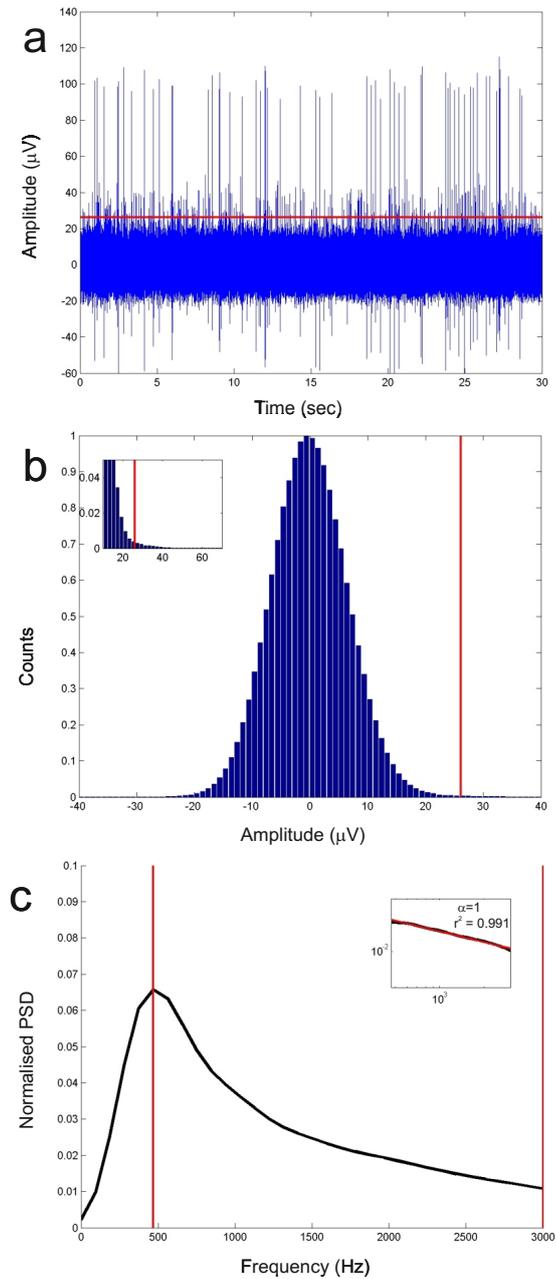


Figure 3.6: Example of a 2 minute simulation of an extracellular recording including background noise, multi-unit activity and two single-units. (a) First 30s of continuous data filtered between 300 and 3000 Hz. (b) Amplitude distribution of the voltage signal. The vertical red line represents the amplitude detection threshold. The inset shows a zoom around the detection threshold. (c) Power density spectrum (PDS) of the continuous data from 300 to 3000 Hz. The inset depicts the log-log PDS between the two red vertical lines. The red line in the inset shows a linear fit to the data, which gave a slope $\alpha = 1$ and a correlation coefficient $r^2 = 0.991$.

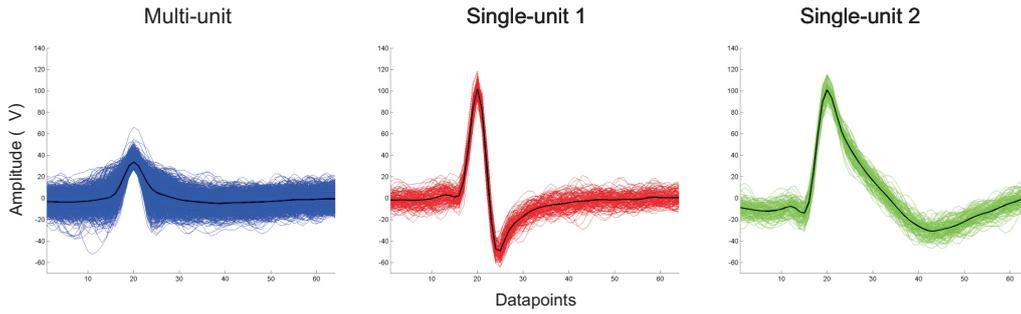


Figure 3.7: Result of the automatic clustering for thirty minutes of the recording in the previous figure. The algorithm identified one multi-unit (left) and two single-unit clusters (right). Note that the multi-unit cluster shares similar critical features, such as spike shape variability and non-symmetry of peak amplitude distribution, compared to the real recording shown in Figure 3.2.

Example number	Number of detected spikes	Spike detection		Spike sorting	
		Misses	False positives	Misses	False positives
1	236	0	0	0	0
2	1058	0	0	0	0
3	1163	5	80	31	106
4	1127	10	87	10	87
5	118	0	0	0	0

Table 3.2: Performance of the unsupervised spike sorting algorithm. All values are referred to single-unit spikes. Values in the spike sorting section of the table include cluster classification errors.

example.

3.4 Discussion

3.4.1 Hybrid Simulation Approach

In this chapter we showed the use of a hybrid two-fold strategy to characterize most salient features of real data to test spike sorting algorithms. The first component of our approach was to simulate the background noise by adding the contribution of many different far-away action potential point-sources, uniformly distributed around the electrode, where the amplitude of each of these

action potentials was determined by their distance from the electrode following Coulomb's law. This simple and physiologically plausible model reproduced the main characteristics of the noise activity in real recordings, avoiding the need of simulating computationally expensive details of each of the far-away neurons. Note that this strategy is not plausible for close-by neurons -which generate the multi-unit and single unit activity- since in this case the particular physiological details of the neurons and their relative positions and orientations with respect to the recording site become relevant (Rall, 1962; Holt and Koch, 1999; Gold et al., 2006; Harris et al., 2000). Therefore, the second component of our approach was to generate the single-unit and multi-unit activity by simply using template spike shapes and the peak amplitude distributions available for a large number of real recordings.

This simple two-fold strategy to generate the noise and the spiking activity allowed the performance of relatively fast simulations -about 9 minutes for a simulation of a 2 minutes recording using an Intel Core 2 PC with a clock frequency of 2.4 GHz with Matlab-, which reproduced the main features of extracellular recordings, as detailed below.

3.4.2 Multi-unit Activity

The presence of multi-unit activity in extracellular recordings makes the processes of spike detection and sorting particularly challenging, an issue that has not been dealt in previous simulation studies (Lewicki, 1994; Quian Quiroga et al., 2004; Smith and Mtetwa, 2007). This is mainly for two reasons. First, the low signal-to-noise ratio and high amplitude variance of multi-unit clusters complicates the detection of spikes because there is no clear separation between background noise and spikes. Second, multi-unit clusters are usually non-Gaussian, which compromises the use of spike sorting algorithms assuming Gaussian clusters or introducing a pre-whitening of the data (Pouzat et al., 2002; Lewicki, 1998). Third, multi-unit clusters can be relatively large compared to those of

single-units, given that they are composed by the activity of several units (in spite of the fact that a large proportion of spikes of these units may be not detected). This can complicate the accurate sorting of single-units and even their identification, since these can be easily merged within the large multi-unit clusters. This issue becomes more problematic for the identification of sparse firing neurons (Quiari Quiroga et al., 2005; Quiari Quiroga, 2007).

3.4.3 Background Noise

The characteristics of the background noise recorded from in-vivo extracellular recordings play a fundamental role in the identification and classification of single- and multi-unit activity. In this respect our noise simulations, generated by superimposing a large number of spikes from far-away neurons, reproduced the amplitude and power spectra distribution of the real recordings. In this study we did not systematically study up to which distance neurons contribute to the single- and multi-unit activity and beyond that to the background noise. Rather we used a heuristic minimum distance for the noise contribution that could reproduce the amplitude distribution observed in real recordings and used typical amplitude values observed in real data for the multi- and single-units. For the noise generation, decreasing this heuristic distance gave a large number of spikes crossing the detection threshold, in contrast to what we found in the real data (see Figure 3.4). A more precise and systematic investigation of the distance ranges giving rise to the noise, multi- and single-unit activity is ripe of future investigation, but it would have also to consider other factors, such as the impedance profile of the electrodes and more detailed characteristics of the neurons nearby the recording site. These results could be then compared to those obtained from real recordings, by retracting acutely implanted electrodes once a neuron has been identified and observing at what distance the spikes of a single neuron gets mixed with those of other neurons -in between 50 to 140 microns- and at what distance are the spikes of the first neuron are no longer detected

-at about more than 140 microns- (Gerstein and Clark, 1964; Henze et al., 2000; Buzsaki, 2004).

It has to be noted that with this approach we did not try to replicate the lower frequency activity of LFPs and the high frequency components generated by single-cell currents (Llinas, 1988; Farrant et al., 1994) or by relatively lower electronic noise from the recording system (Fee et al., 1996a). To replicate the power law decay in the frequency spectrum observed in real data, high frequency noise was introduced by superimposing Gaussian noise to the background activity generated by the contribution of the spikes from far-away neurons. For this study it was not necessary to model the contribution of LFPs, since their activity is outside the frequency range considered for spike sorting, which is in fact typically filtered to perform the spike detection.

Alternatively, it is in principle possible to obtain noise signals from real recordings by eliminating the spiking activity (Pouzat et al., 2002). Then, the identified spikes corresponding to different neurons can be added to the noise to evaluate the performance of spike sorting algorithms. However, the difficulty of such an approach is to eliminate the spiking activity because this depends on the setting of the detection threshold, which leaves some of the (low amplitude) spikes and removes part of the noise activity due to random crossings of the threshold. To overcome these problems, one could still use real recordings without spiking activity, as we used in Figure 3.4d to compare the amplitude distributions of our noise simulations. However, the availability of such data is not guaranteed, especially if one requires a large number of simulations.

The noise model presented here did not assume any particular distribution of the data, as in the case of generation of white or coloured Gaussian noise in which clusters have hyper-spheres or ellipse boundaries (Chandra and Optican, 1997). It has to be noted that several experimental conditions, such as electrode drifts, bursting activity, overlapping spikes, multi-unit activity or mis-alignments during detection can produce non-Gaussian clusters (Quiñero, 2007), which compromise the use of algorithms assuming clusters with a multi-

dimensional Gaussian shape (Lewicki, 1994; Pouzat et al., 2002) and indicates that the use of such algorithms might not be plausible in general.

In conclusion, we presented a method to generate simulations that reproduced the most salient characteristics of extracellular recordings, like the power spectrum and amplitude distribution, which can be useful for testing spike detection and sorting algorithms. Moreover, these simulations might also be used to quantify filter distortions in the spike shapes (Quiari Quiroga, 2009), to simulate recording conditions with different electrode designs (Robinson, 1968) or as a platform to test effects of spike detection and sorting for inferring information about, for example, time patterns and synchronised activity in a neural population (Quiari Quiroga and Panzeri, 2009). Our hybrid approach to generate the synthetic data was able to provide several simplifications making this algorithm simple and fast.

3.5 Summary of Chapter 3

In this chapter we presented a new approach to produce simulations of extracellular recordings. The generation of realistic noise and multi-unit activity together with single-unit activity offers new opportunities for the development of more efficient analytic tool. The possibility to generate realistic scenarios with for example, neurons with high and low firing rates present during the same recording, makes these results an optimal platform in which to compare the performance between different spike sorting algorithms. In Chapter 4, we apply the method presented here to develop a new solution for automatic spike sorting under realistic experimental conditions such as the ones mentioned earlier.

Chapter 4

Automatic Spike Sorting

4.1 Introduction

As described in Chapter 2, `Wave_clus` provides a first unsupervised classification of units based on the automatic selection of a single temperature. However, in some cases is likely to find clusters appearing at different temperatures in which further intervention of an expert operator is necessary to reach the optimal solution (see Figure 2.5 in Chapter 2). This condition turns the method into a semi-automatic approach, slowing the sorting process and introducing sometimes changes in the solution depending on which person performed the analysis. Here, we show a new approach to the automatic selection of clusters using SPC. Our proposed solution introduces two major improvements that successfully overcome the difficulty of identifying clusters of different densities and sizes by analysing the data at different temperatures. First, we construct a threshold variable with the temperature that successfully detects natural clusters and rejects overclustering. Second, we simultaneously select clusters at different temperatures. We quantified the performance of our algorithm using the simulation platform described in the previous chapter. As a result, we developed a robust and fully automatic spike sorting algorithm which significantly increases the number of identified units without manual intervention.

4.2 Materials and Methods

Our method follows the clustering approach of `Wave_clus` in which SPC is used to provide an unsupervised selection of clusters. As described in Section 2.6, SPC is a nearest-neighbour algorithm in which the grouping of points into clusters strongly depends on a parameter called the temperature T which controls the level of correlation between neighbouring points. For low temperatures, the correlation between neighbouring points increases and all the data points, even located at far distances in the feature space, will be grouped as a single cluster. For high temperatures, this correlation becomes very weak thus grouping the data points into several clusters of small size. The optimal solution, in which natural clusters appear, is expected to be found within a range of temperatures between the lower and upper extremes.

Limitations of `Wave_clus`

In extracellular recordings, the activity of neurons with different firing rates and spike amplitudes might generate clusters that differ in their density and location within the feature space. This condition increases the difficulty of the problem in which clusters of different densities might appear at several temperatures. Another complication is the fact that at the highest temperatures, high density clusters are likely to split, resulting in an undesired partition of the data, called overclustering. We show this scenario with an example of real data, in particular an extracellular recording from the human medial temporal lobe of an epileptic patient.

The purpose of this example is to illustrate two main complications in the partition of the data arising from a typical recording: i) clusters appearing at different temperatures and ii) overclustering. Figure 4.1a contains the temperature diagram generated by SPC and four different clusters identified at several temperatures. At the lowest temperature $T = 0$ only the biggest cluster (BC) appears (see the blue curve in the temperature diagram). At $T = 0.02$ from

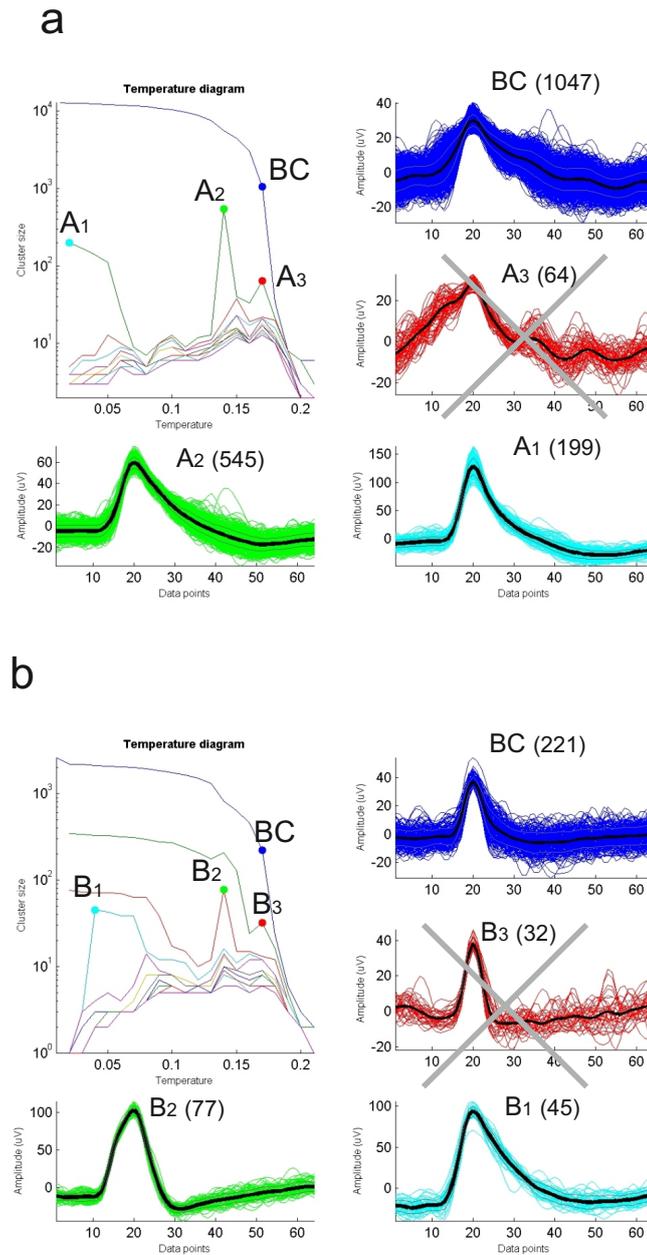


Figure 4.1: (a) Example of a real recording from the human medial temporal lobe (approximately 30 min). Clusters were obtained after manual spike sorting using the temperature diagram depicted in the figure. Clusters in green and cyan represent the spike shapes of two single-units, appearing at temperatures $T = 0.02$ and $T = 0.14$, respectively. The multi-unit cluster (in blue) is represented by the biggest cluster (BC) at $T = 0.17$. Cluster in red, also at $T = 0.17$, represents an overclustering case of the multi-unit and should not be selected. (b) Similar case reproduced with a simulation of 5 minutes duration, including one multi-unit and two single-units. Note also here, 3 salient peaks in the temperature diagram, with B_1 and B_2 representing good clusters and B_3 appearing at a spurious partition which should be avoided.

BC we can separate cluster A_1 (single-unit in cyan). As the temperature further increases, the BC gives at $T = 0.14$ a new cluster A_2 (single-unit in green) but at this temperature the cluster A_1 no longer exists. At a higher temperature $T = 0.17$ a new partition is generated in which we obtain a new cluster A_3 (in red) of smaller size (64 spikes). However, cluster A_3 seems to be an overclustering of the BC at the clustering temperature $T = 0.17$ where their shapes overlap. In neurophysiology, the new cluster A_3 would be normally labeled as noise or a small part of the BC (multi-unit activity). A simple way to solve this problem would be for instance, to set a hard threshold to reject clusters with less than ~ 60 points. However, this solution would not detect small clusters that might represent in some cases the sparse activity of neurons that should be considered.

In Figure 4.1b we reproduced the same problem using synthetic data. We included one multi-unit and the activity of two single-units with different spike amplitudes and firing rates. As in the example depicted in Figure 4.1a, the data broke into two simulated single-unit clusters in two different stages. In the first one, the cluster B_1 appeared at a low temperature ($T = 0.04$). At $T = 0.14$ the second partition occurs where from BC we extract B_2 . The simulation also reproduced the overclustering case typically found in real data where in this case, a further breaking of the BC at $T = 0.17$ generated a spurious cluster (B_3).

The automatic selection of clusters used by Wave_clus, based on a hard-threshold approach applied on a single temperature, limits the method to deal with challenging situations commonly found in extracellular recordings, as demonstrated earlier on both real and synthetic data. The main limitations of Wave_clus are i) identification of clusters based on a single threshold for all temperatures and ii) selection of clusters at a single temperature. When these complications arise, the optimal clustering is reached by means of a supervised approach. In this situation, the operator has the possibility of choosing the temperature (e.g. $T = 0.14$ in Figure 4.1a and Figure 4.1b) and also selecting or rejecting clusters at different temperatures. In the following section we describe a new criterion for doing this automatically.

4.2.1 Improvements

The aim of the new algorithm is to obtain an optimal unsupervised classification of single-units using the several clustering partitions provided by SPC. Our method overcomes the difficulties found in extracellular recordings arising from clusters of different densities and sizes. In order to deal with this complicated situation, our strategy introduces two major improvements: i) Selection of optimal clusters and rejection of overclustering cases of applying a variable threshold to identify the optimal cluster size at each temperature and ii) simultaneous identification of clusters appearing of different temperatures by analysing the data partitions at several temperatures.

Criterion 1

To identify optimal clusters of different sizes at all the data partitions, we define a threshold as a function of the temperature, in contrast with the hard-threshold approach based on a fix number of spikes used by the old algorithm (see equation 2.7). We know that in SPC, the probability of breaking clusters within high-density regions increases with the temperature. At higher temperatures, this phenomenon may lead to spurious partitions of the data (overclustering). As a result, artificial clusters, may appear at high temperatures, leading to false positives if a threshold based on a fixed cluster size is adopted. On the other hand, in partitions occurring at lower temperatures, the size of the natural clusters might be smaller. Thus, an increment in the threshold level to avoid overclustering at high temperatures might lead to misses at lower ones. Moreover, in systems with phase transitions, the data partitions as temperature increases are related to the decay of the size of the BC (Ison et al., 2002). Our strategy deals with this complication by means of creating an adaptive threshold using the dynamics of the BC across all partitions. We define the value of the threshold θ_i as

$$\theta_i = b \cdot \frac{BC_0}{BC_i} \quad (4.1)$$

where b is an optimization parameter related to the size of the cluster to be identified, equivalent to a used in the old algorithm (see equation 2.7 in Chapter 2). BC_0 and BC_i represent the size of the biggest cluster at $T = 0$ (i.e. the total number of spikes) and $T_i = 0.01, 0.02, \dots, 0.2$, respectively. The algorithm selects a potential clustering at temperature T_i when $A_i^j - A_{i-1}^j \geq \theta_i$, where A_i^j represents the size of the j^{th} cluster at T_i . Once the desired temperature is located we select the fragments containing A_i^j spikes or more. With this criterion, the algorithm becomes more conservative in the selection of clusters as temperature increases.

Criterion 2

In order to fully reach the potential of the SPC, we should analyse the data at all partitions. For example, one cluster identified at a certain temperature T_1 , might no longer appear at some other temperature T_2 in which a new cluster might be identified. Our second criterion deals with this complication, selecting clusters at different temperatures simultaneously. This fact introduces the problem of identifying whether two clusters appearing at different temperatures are the same or not. To overcome this difficulty, we added a second condition to this criterion, quantifying the overlapping of clusters.

Assuming that, in SPC, clusters break down as T increases, in the simple case of having two data partitions at temperatures T_1 and T_2 ($T_1 < T_2$), three possible scenarios might arise: a) points in the cluster A at T_1 correspond to the same cluster at T_2 ; b) points in one single cluster A at T_1 correspond to different clusters, i.e. B and C, at T_2 and c) points in different clusters, i.e. A and B, at T_1 correspond to different clusters, i.e. C and D, at T_2 . For all these possible conditions, we want to know if the cluster A appearing at T_1 has broken down into different clusters B and C at T_2 , discarding cluster A if this situation occurs. Then, we need to quantify if clusters at T_1 and T_2 are included in each other. The *Jaccard* coefficient (Jaccard, 1901) measures similarity between clusters with sizes A and B coming from different partitions. It is given by $C_{jaq} = \frac{|A \cap B|}{|A \cup B|}$. The problem with this quantification is that if B is a small subset but completely

included in A, C_{jaa} will be small. Instead, we proposed an *overlapping coefficient* C_{ov} defined as

$$C_{ov} = \frac{|A \cap B|}{\min(A, B)} \quad (4.2)$$

with C_{ov} ranging between 0, if they are completely different and 1 when they are both the same (i.e they do not share any spike), or if one is a subset of the other.

4.2.2 Simulated Data

In order to quantify the performance of our spike sorting algorithm we used realistic simulations of extracellular recordings. We generated these synthetic data following the approach described in Chapter 3, a model that reproduces the contribution of the background noise, multi- and single-unit activity. All the simulations had a duration of 5 minutes with a sampling rate of 24 kHz and a noise level based on its standard deviation (see equation 2.3 in Chapter 2) and equal to $7\mu V$. We included the activity of 1 multi-unit per simulation and 1 to 5 single-neurons. The amplitudes of the single-units were uniformly distributed between 70 to $120\mu V$ and the firing rates between 0.1 to 2 Hz. This variability is expected to generate clusters with different sizes and densities, creating difficult problems for the spike sorting algorithm.

4.2.3 Algorithm Optimization and Validation

We trained and tested both the former and new clustering algorithms using two datasets of 100 simulations each. As an optimization parameter for each method, we varied the threshold for choosing a cluster, represented by a and b in the the former and new algorithm, respectively (see equations 2.7 and 4.1). We evaluated the outcome of the algorithms with values of a equal to 10, 30, 50 and 60 spikes and values of b equal to 5, 10, 15 and 20 spikes. The performance of each algorithm was quantified by its number of hits, false positives and misses. For single-units, an identified cluster was considered as hit when more than

Optimization parameter		Misses	False positives	Errors
a (Former)	10	165	158	323
	30	58	14	70
	50	51	5	56
	70	57	2	59
b (New)	5	15	75	90
	10	18	6	24
	15	21	1	22
	20	23	0	23

Table 4.1: Number of simulated units misses, false positives and errors for the different optimization parameters a and b used in the former and new algorithm, respectively. In this analysis a total number of 410 units were included

50% of its spikes come from more than the 50% of the spikes of the simulated cluster. A selected cluster was defined as false alarm when it did not meet the hit criterion. Missed clusters were calculated as the number of simulated units minus the number of hits. For multi-units, a hit was considered when more than 50% of its spikes came from the simulated multi-unit. Note that in this case, we do not apply the same rule as in single-units; since, by design, a significant number of the spikes included in the simulated multi-unit will not cross the detection threshold (see Martinez et al. (2009) for details). We applied a paired T-Test to quantify the significance in the differences resulting in the sorting errors of each clustering algorithms for each one of the 100 simulations in the training and test datasets.

4.3 Results

4.3.1 Optimization

We optimised the performance of the former and new clustering algorithms using a total of 100 simulations containing 410 clusters: 100 multi-units and 310 single-units. Table 4.1 shows the classification results in terms of misses, false positives and errors, corresponding to the two methods for the different optimization parameters a and b , respectively.

Based on these analyses, we selected a and b at which each algorithm provides the smaller number of classification errors. For the training dataset used in this study, the optimal thresholds for the former and new methods were 50 and 15 with sorting errors of 13.7% (56/410) and 5.4% (22/410), respectively. For both algorithms all the multi-unit clusters were correctly detected.

4.3.2 Example of Use of Criterion 1

In Figure 4.2 we show the benefits of using criterion 1. This example, includes one multi-unit and two single-units with high firing rates and spike amplitudes.

Figure 4.2a depicts the cluster assignment using the former algorithm. In this case, this simulation presents a significant increment in the number of spikes at a high temperature due to overclustering. The automatic approach of the former algorithm, based on a fixed threshold, selected this temperature, thus giving a wrong partition. At this temperature ($T_{clus} = 0.16$), the partition presents an overclustering of the data, imposing the breaking of the multi-unit (blue cluster) into small spurious clusters (red and green clusters). In Figure 4.2b we show how the variable threshold (see equation 4.1), becomes more conservative at high temperatures, avoiding the overclustering made by the former algorithm. Instead, this criterion selects the correct partition ($T_{clus} = 0.02$) at which three natural clusters appear, corresponding to one multi-unit (blue) and two simulated single-units (red and green).

4.3.3 Example of Use of Criterion 2

In the example of Figure 4.3 we show how the criterion 2 improves the clustering results by analysing the data at different partitions. In Figure 4.3a, the former algorithm, based on a single-temperature selection approach, looks for the highest temperature at which the increment in the number of spikes crosses the threshold a . In this case the cluster assignment is done at $T_{clus} = 0.14$, identifying one multi-unit (blue) and one single-unit (red). In Figure 4.3b, we show how the

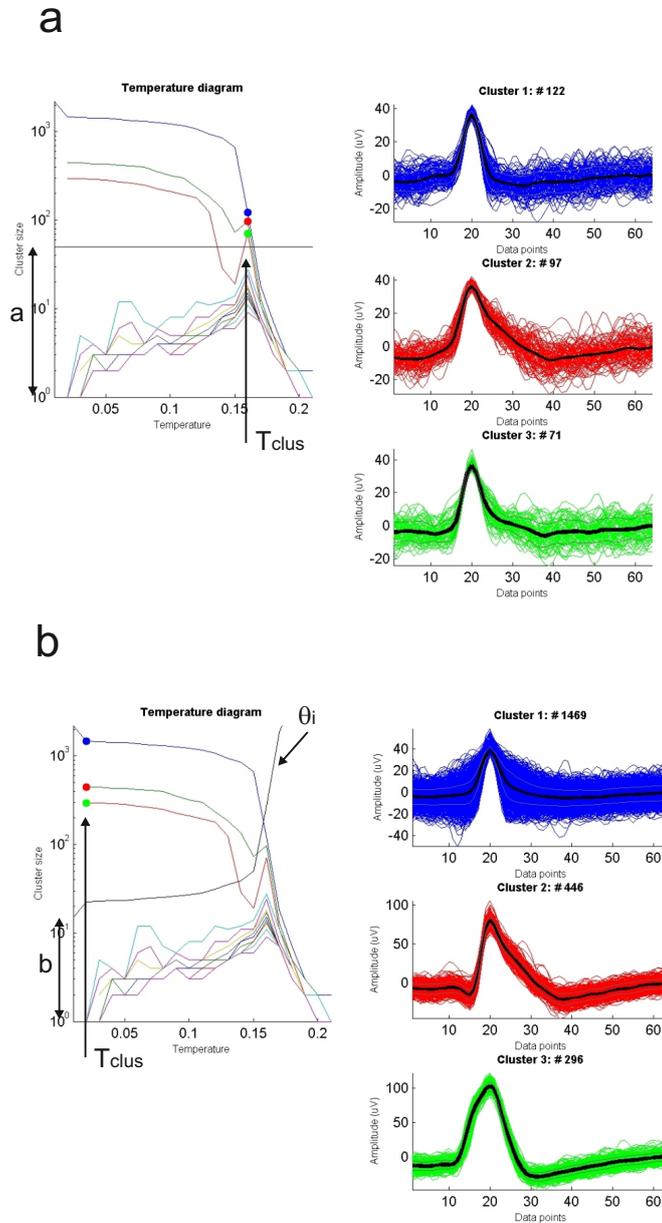


Figure 4.2: Example of use of criterion 1. (a) Former criterion for the selection of clusters based on a fixed threshold a at the highest temperature (see equation 2.7). Note how the partition at $T_{clus} = 0.16$ imposes the clusters red and green to appear as an overclustering of the biggest fragment (in blue). (b) Criterion for the selection of clusters used by the new algorithm. Note here how a threshold variable with the temperature as in θ_i , dramatically improves the selection of clusters, becoming more conservative as the temperature increases (see equation 4.1). In this case, the new solution selects the optimal partition at $T_{clus} = 0.02$ and at the same time avoids the potential overclustering at $T = 0.16$.

new criterion, based on a multi-temperature selection of clusters, identifies two more clusters appearing at $T_{clus1} = 0.03$, corresponding to the activity of two correctly classified single units (green and cyan). The new algorithm also allows the selection of the clusters appearing at $T_{clus2} = 0.14$. Note also the role of the overlapping criterion, when selecting clusters from different partitions. The overclustering coefficient C_{ov} (see equation 4.2) between the BC at T_{clus1} (black dot) and the same cluster at T_{clus2} (blue dot), prevented the first one from being selected which further broke into a new one at T_{clus2} (red dot).

4.3.4 General Results

Once both algorithms were optimised with the training dataset, we performed the quantification of both the old and the new implementations, using the testing set of 100 simulations with a total of 430 simulated clusters (100 multi-units and 330 single-units). In Figure 4.4 we present the outcome of both methods in terms of number of misses, false positives and classification errors. In comparison with the old algorithm, the new implementation significantly improves the spike sorting results. The classification errors in the new algorithm (35/430), compared to the old method (86/430) were significantly lower ($p < 0.05$). In terms of false positives, the new implementation also outperformed the old one (4/430 and 8/430, respectively) although, these differences were not statistically significant ($p = 0.158$). Both algorithms correctly identified all the multi-unit clusters.

In terms of computational speed, both algorithms present similar performances. Note that all the improvements are implemented after spike detection, feature extraction and Monte Carlo simulations for the calculation of the temperature diagram by the SPC algorithm, taking all these stages together an overall processing time of approximately 60 seconds (for our test set). This makes the implementation of the new automatic solution very efficient, since it provides a significant improvement in terms of classification performance at a low computational cost as compared to the former solution. In this case, the

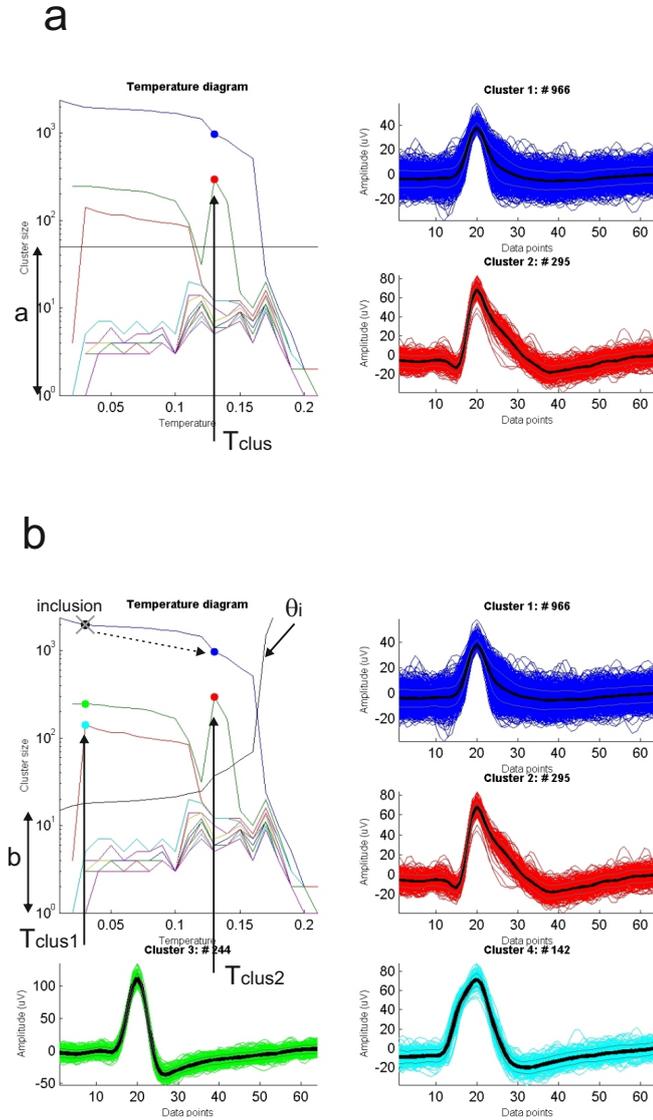


Figure 4.3: Example of use of criterion 2. (a) Former criterion for the selection of clusters based on single-temperature selection of clusters. Note that looking only at one partition, the former algorithm selected the highest temperature meeting its threshold criterion, in this case $T_{clus} = 0.14$. (b) Identification of clusters based on a multi-temperature analysis approach. Looking at several partitions, the new algorithm is capable of selecting clusters appearing at different temperatures ($T_{clus1} = 0.03$ and $T_{clus2} = 0.14$). Also, the *overlapping coefficient* prevents the new clustering algorithm from selecting clusters containing the same group of points (i.e. spikes in the BC at T_{clus1} (black dot) appear also in the read and blue cluster at T_{clus2} as quantified by C_{ov} (see equation 4.2) and therefore the BC at T_{clus1} should be discarded).

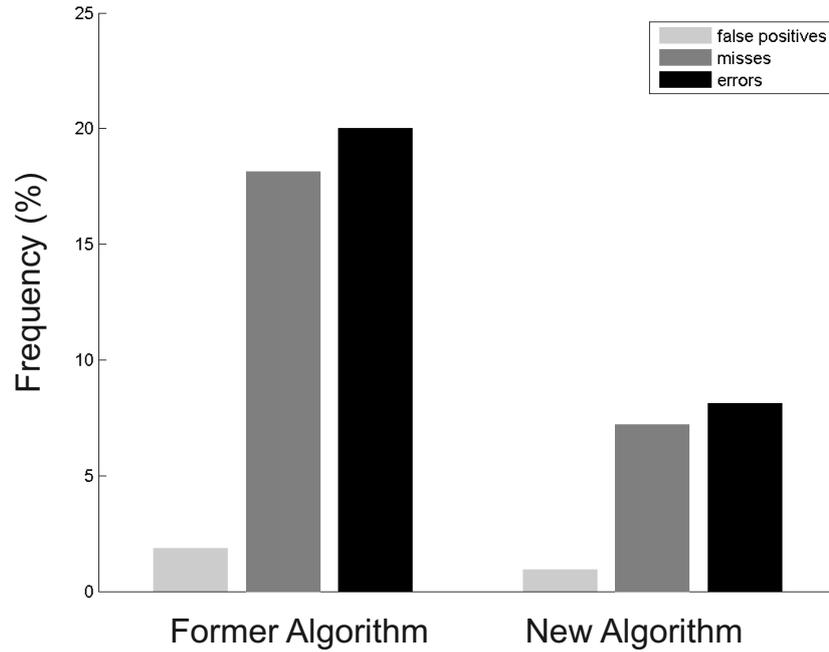


Figure 4.4: Outcome of both algorithms in percentage of misses, false positives and classification errors for a total of 430 clusters (100 multi-units and 330 single-units) across 100 simulations.

new algorithm calculated the automatic solution in 0.22 ± 0.087 seconds (average \pm SD across the test set) while the former algorithm obtained the results in 0.15 ± 0.038 seconds.

Figure 4.5 depicts a scatter plot showing the uniform distribution of amplitudes and firing rates of the simulated single-units (grey dots). The misses, corresponding to the former and new implementations, are represented by black circles and dots, respectively. Note that neurons with low amplitudes (i.e. $< 80\mu V$) and low firing rates (i.e. $< 0.5\text{Hz}$) are less likely to be identified with both methods. However, a number of units with both high levels of signal-to-noise ratio and firing rates, easier in principle to be identified, were also missed when using the former algorithm.

Figure 4.6a depicts the number of misses for each algorithm for different firing rate bands (0.1-0.5, 0.5-1, 1-1.5 and 1.5-2 Hz). Note that with both algorithms the performance was worst for neurons with very low firing rates (band

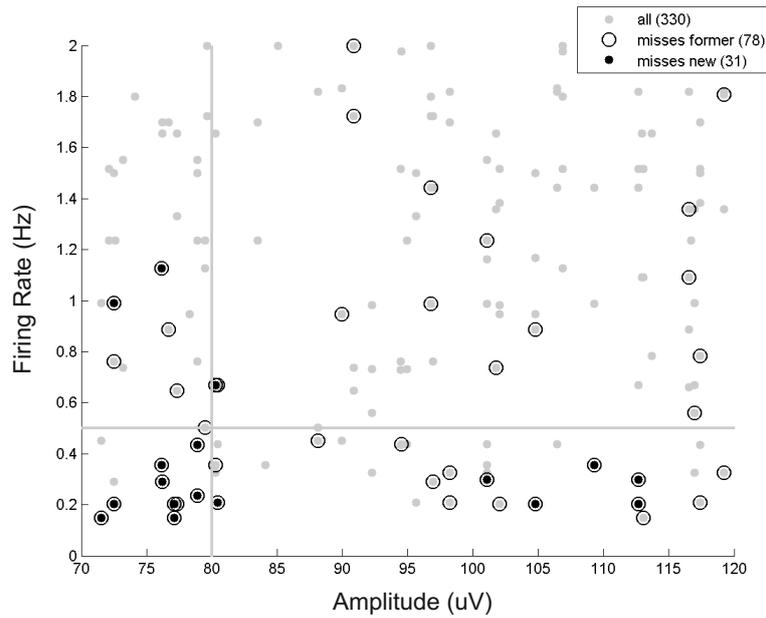


Figure 4.5: Scatter plot of the uniformly distributed firing rates and peak amplitudes across all the single-unit clusters (grey dots). Black circles represent the missed clusters with the former algorithm (78) and black dots the missed clusters using the new algorithm (31).

of 0.1-0.5Hz). However, under this challenging condition, the new implementation significantly decreased the number of misses (25/92) as compared with the former one (49/92). In Figure 4.6b, we see the number of missed clusters for different ranges of spike amplitudes (70-80, 80-90, 90-100, 100-110 and 110-120 μV). In this case, a relative large number of units within the lowest amplitude range were not detected with the former method (20/72), while with the new implementation the performance was significantly better (3/72). In all these cases the differences were statistically significant ($p < 0.05$).

4.4 Discussion

In the recent years many spike sorting techniques have been proposed (Lewicki, 1994; Fee et al., 1996a; Harris et al., 2000; Letelier and Weber, 2000; Hulata et al., 2002; Pouzat et al., 2002; Quian Quiroga et al., 2004). However, most of these methods have been designed to provide a semi-automatic solution assuming

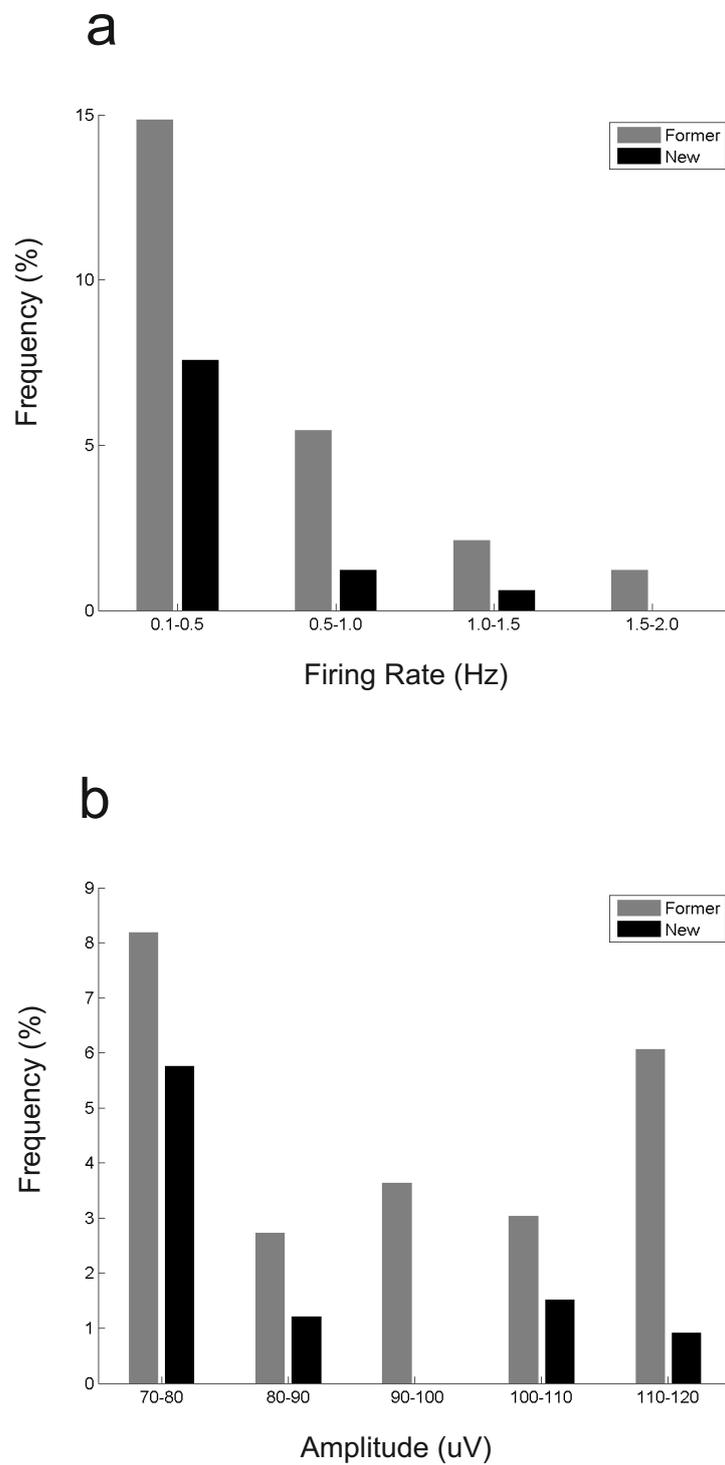


Figure 4.6: Percentage of misses with the former (grey) and new algorithm (black) for different firing rates (a) and amplitude bands (b).

in more or less proportion the intervention of a human operator. One of the most widely used methods for spike sorting is Wave.clus (Quiari Quiroga et al., 2004) and one of the few algorithms providing a fully unsupervised clustering solution. In this case, Wave.clus selects one optimal temperature identifying automatically the corresponding optimal clusters. In a relatively high number of recordings, this solution is optimal. However, the manner in which the system changes depends on the structure in the data. Recordings including the activity of neurons with different firing rates, amplitudes and spike shapes are likely generate complex scenarios in which several clusters of many sizes might appear at different temperatures.

In this chapter we presented a fully automatic clustering that significantly improves the unsupervised performance of a previously developed and widely used spike sorting algorithm. Our method successfully overcomes the problem of identifying clusters of different densities and sizes appearing at different temperatures while avoiding spurious clusters. To obtain these results we proposed two criteria: we defined a variable clustering threshold following the dynamic of the biggest cluster and we considered clusters at different temperatures. We also included a measure that quantifies the degree of overlapping when clusters appear at different partitions. As with any automatic algorithm, its outcome might result in some cases not correct for an expert operator, being apparently better the use of a semi-supervised approach. However, the possibility of gathering large data sets by recording from multiple channels and the need of tools for a fast analysis of this data makes the use of unsupervised spike sorting algorithms a very attractive method for the study of large populations of neurons in the years to come.

4.5 Summary of Chapter 4

In Chapter 4 we presented a new solution for automatic spike sorting. The method described here significantly improves the unsupervised outcome of a well

known spike sorting algorithm by providing a more efficient solution when dealing with clusters of different sizes and densities, typically found in real extracellular recordings. The excellent results obtained in the automatic isolation of single-neuron activity will to take full advantage of the increasing advances in multi-channel electrode technology by reducing de errors and processing time during manual intervention.

Part III

Conclusions

Chapter 5

General Conclusions

The main goal of this thesis was the development of technology for the automatic analysis of neural activity. We addressed this general problem with two specific contributions: the development of realistic simulations of extracellular recordings and improvements to an automatic spike sorting algorithm.

The quality of neural data depends crucially on the particular spike sorting technique that is used to isolate the activity of each of the recorded neurons. Furthermore, advances in multichannel electrodes provide the appropriate technology for recording from large-scale neuronal ensembles. However, this offers a challenge for the development of efficient spike sorting methods, since manual sorting becomes a very subjective and time consuming task and automatic algorithms -which are in turn more difficult to develop- should be preferred.

For an optimal quantification of the performance of the different spike sorting techniques we created a robust simulation framework in which the challenging conditions of real extracellular recordings were reproduced.

The results obtained in both these areas contribute to the development of automated and efficient quantitative platforms for the analysis of large-scale population codes, thus contributing to a better understanding of how our brain functions.

5.1 Simulation of Extracellular Recordings

In chapter 3 we presented a method to generate simulations that reproduced the most salient characteristics of extracellular recordings. Such theoretical platforms play a crucial role in the development of optimal spike detection and sorting algorithms.

A number of different algorithms for spike sorting have been proposed in the past years (Lewicki, 1994; Fee et al., 1996a; Harris et al., 2000; Letelier and Weber, 2000; Hulata et al., 2002; Pouzat et al., 2002; Quiñero et al., 2004). However, it is still very difficult to compare and evaluate the outcome of these methods due to the lack of shared datasets that allow to unequivocally quantify their performance. On the one hand, it is in general not possible to perform such quantifications with real datasets because we do not have access to 'ground truth', i.e.: the original labels marking to which neuron belong each spike (but see Wehr et al. (1999); Harris et al. (2000); Henze et al. (2000) for notable exceptions, in which simultaneous intracellular and extracellular recordings provided labels to each spike).

We presented a new platform based on a hybrid and computationally simple approach to generate synthetic datasets. Our method accurately reproduced salient features of the background noise such as power spectrum and amplitude distribution. The single-unit activity was simulated by adding spikes trains following a Poisson process on top of the background noise. We also generated multi-unit activity as the contribution of many spike shapes with peak amplitude distributions following those corresponding to real data.

In summary, we developed realistic simulations reproducing those features of real recordings that are particularly challenging for any spike detection and sorting algorithm. These results provide an adequate testing platform to develop robust algorithms as well as a benchmark for quantifying the performance between different spike detection and sorting techniques.

5.2 Automatic Spike Sorting

In chapter 4 we presented a new implementation for automatic clustering that significantly improved the unsupervised solution of a widely used spike sorting algorithm.

Making use of the results presented in Chapter 3, we tested our method using a set of simulations including multi-unit and single-unit activity with different levels of signal-to-noise ratio and firing rate as well as different number of neurons. These synthetic data sets, reproduced several realistic situations in which the correct classification of neural activity becomes very challenging for an automatic clustering algorithm. For example, the activity of sparse neurons (cells with low levels of activity) is very difficult to identify since is usually masked by other clusters with high firing rates such as the multi-unit. Also, the activity of neurons with low signal-to-noise ratio activity is usually mixed with the multi-unit.

Under this challenging but also typical condition in extracellular recordings, the previous solution for automatic sorting did not provide the expected results in a considerable number of cases, being necessary the intervention of a human operator to reach the optimal solution.

Our new solution for automatic clustering successfully dealt with the problem of detecting neurons with different levels of activity and signal-to-noise ratio in the same recording, increasing the number of isolated units without manual intervention. Moreover, the new method also improved the identification of sparse neurons, commonly very difficult to identify using automatic spike sorting algorithms.

The results presented in Chapter 4 will also contribute to the analysis of large-scale recordings. New developments in the area of intracortical electrodes (Drake et al., 1988; Hetke et al., 1994; Rousche and Normann, 1998) allow us to record from multiple extracellular signals at a time by using a high number of channels. This, together with the use of better invasive technology for optimal

chronic implants, provides a new scenario in which it is increasingly common to generate large datasets during the course of an experiment. The analysis of this vast amount of data usually becomes a very time consuming process since, for the reasons mentioned above, the operator usually needs to supervise dozens of channels and manually classify the neural activity.

In summary, the results presented in this work contribute to the efficient and fast automatic classification of single-unit activity, taking full advantage of the new advances in the electrode technology and helping to understand how cognitive process are encoded by individual neurons.

Chapter 6

Suggestions for Future Work

In this chapter we present suggestions for future work based on the results described in Part II. The work presented here has made novel contributions on the automatic analysis of neural data. In Chapter 4, we presented improvements for an efficient unsupervised classification of single-units, that we tested using a realistic simulation framework described in Chapter 3. These results opened new questions that we address here.

6.1 Improvements on Simulation of Extracellular Recordings

In chapter 3 we presented a new method to generate efficient simulations of extracellular recordings. This method uses a hybrid and computationally simple approach, where the features of the background noise arise naturally from its biophysical process of generation. We simulated the single-unit activity by adding spikes on top of the background noise. Moreover, we simulated the multi-unit activity as a contribution of many spikes of several shapes with amplitude distributions matching the ones found in real recordings. This approach resulted in a powerful method in which we tested the automatic spike sorting algorithm presented in chapter 4. However, the nature of the synthetic multi- and single-units in which the amplitudes and shapes are predefined might limit the generation of

more complex scenarios.

Simulations including phenomena such as the spike shape variability due to neuronal morphology or to the orientation of the neuron with respect to the recording electrode (see Chapter 2), might be used for example to create more challenging conditions in which to test spike sorting algorithms. This new approach would be addressed by developing simulations with a higher level of complexity. Here, the spiking activity of both multi- and single-unit would be generated by detailed computational models of neurons of different morphologies, reproducing realistic spike shapes at different locations within the extracellular space (Gold et al., 2006). This would also open the possibility to develop simulations of extracellular recordings in polytrodes for testing new spike sorting algorithms. In this case, the variation on the spike amplitude and shape across the different channels needs to be considered for the development of new analysis techniques using polytrode data.

6.2 Spike Sorting on Multi-Electrode Data

In Chapter 2 we briefly described the powerful capabilities of using multi-electrode probes for spike sorting. Recording from the same spike shape in different channels opens new ways of analysing single-neuron activity. One of the key aspects when using multi-electrode probes is how to combine the information provided by different channels for an optimal spike sorting. For example, we could concatenate the spikes shapes recorded from different channels to build a new *polyspike* waveform that would be used further for sorting purposes.

Here, we suggest as a future line of work the study of analytical tools to deal with multi-electrode data. For example, the extraction of optimal features from the *polyspike* will be of great relevance for a better cluster isolation. As in single-channel data, one could use a particular set of wavelet coefficients to represent the most salient features of the spike shape. In this case, we suggest the study of optimal wavelet coefficients as well as new features for separation of

polyspike waveforms. Moreover, in recordings using a high number of electrodes, we propose further studies for automatic selection of the best channels (i.e. the ones with higher signal-to-noise ratio) for spike sorting. We also consider a relevant question to study the advantages of using multi-electrodes over single-channel probes (using the simulations of multi-electrode data proposed in the previous section). It is of great interest to know the gain in the total number of isolated units when using i) the combination of several channels or ii) each channel individually. Finally, to create a fast and robust analytical framework using multi-electrode probes we propose to study and quantify the performance of the automatic clustering method described in Chapter 4 on *polyspike* data.

Part IV

Appendices

Appendix A

Simulations

A.1 Extracellular Recordings

In our simulations of extracellular recordings we tried to replicate the main characteristics observed in real data, as measured in the human medial temporal lobe. These recordings were done in epileptic patients, refractory to medication to localize the focus of the seizures and evaluate the possibility of a surgical resection. The basic features of a typical recording were described in Figure 3.2 of Chapter 3.

Figure A.1 shows the histograms of the scaling exponents α corresponding to the linear fits of the power density spectrum (PDS) and the relative peak amplitude for multi- and single-unit spike shapes for 192 channels. The identification of all the clusters was done using `Wave_clus` as a spike sorting algorithm (see Chapter 2). The labeling of the clusters as multi- and single-unit was done by an expert operator, according to criteria defined elsewhere (Quian Quiroga et al., 2005; Quian Quiroga, 2007). The scaling exponent α had a mean equal to 0.98 ± 0.21 (mean \pm SD) and the mean correlation coefficient r^2 was 0.99 ± 0.007 . Figure A.1b and Figure A.1c show the differences in peak amplitude for multi- and single-unit activity. Note the lower peak amplitudes for multi-unit spikes (Figure A.1b), about 1.5 times the detection threshold level (see Chapter 3), compared to the single-unit normalised amplitudes, between 1.5 and 4

(Figure A.1c).

A.2 Simulations for Testing the Spike Sorting Algorithm

Five simulations of extracellular recordings of one multi-unit and two single-units were generated using different firing rates and peak amplitudes (see Table 3.1 in Chapter 3 for a detailed description). Spike detection and sorting was done using `Wave_clus`. Figure A.2 shows the spike shapes automatically detected by `Wave_clus` for examples 2-5. The results corresponding to example 1 are shown in Chapter 3.

Figure A.3 shows the PDS plots for the simulations of the previous figure, after band-pass filtering between 300-3000 Hz. The amplitude of the power spectrum was normalised by the total energy within the 300 - 3000 Hz frequency range. The plots on the left side show the PDS and the ones on the right show the log-log plots with their corresponding linear fits, the scaling exponent α , and the correlation factor r^2 .

Figure A.4 discloses the amplitude distribution for the four simulations of Figure A.2. The vertical red line represents the amplitude threshold used for spike detection, set at 4 times the estimation of the standard deviation of the noise (see Section 3.2). The plots on the right side show a detail of the distribution around the detection threshold. Note the smooth transition between the noise and the detected spikes, as typically seen in real recordings (see inset in middle plot of Figure 3.2 in Section 3.3).

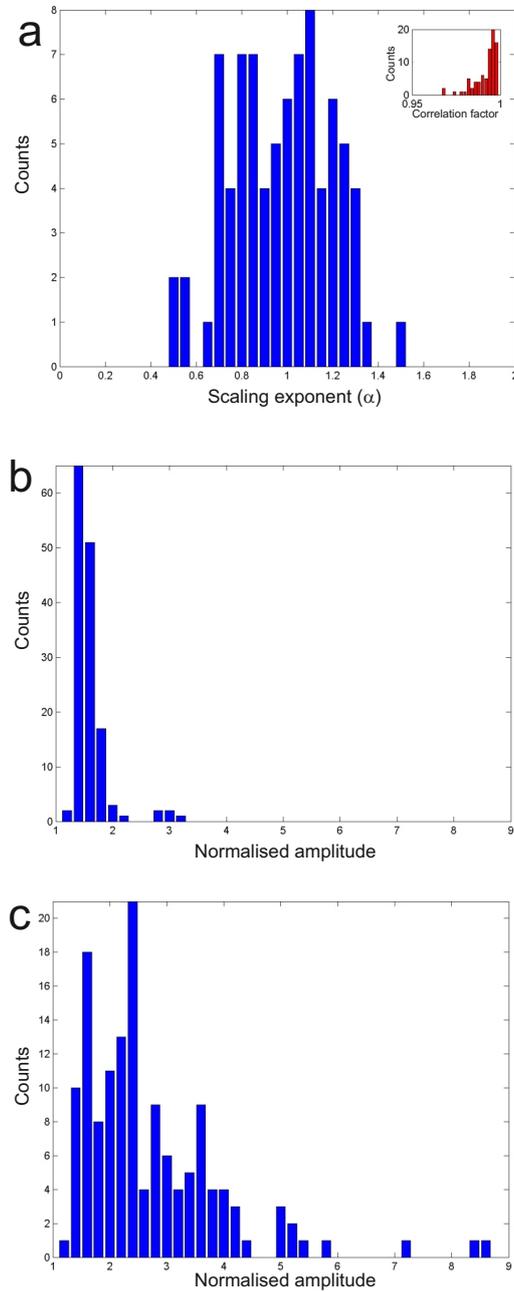


Figure A.1: Scaling exponents α (a), normalised peak amplitude for multi-unit (b) and single-unit activity (c) for 192 channels.

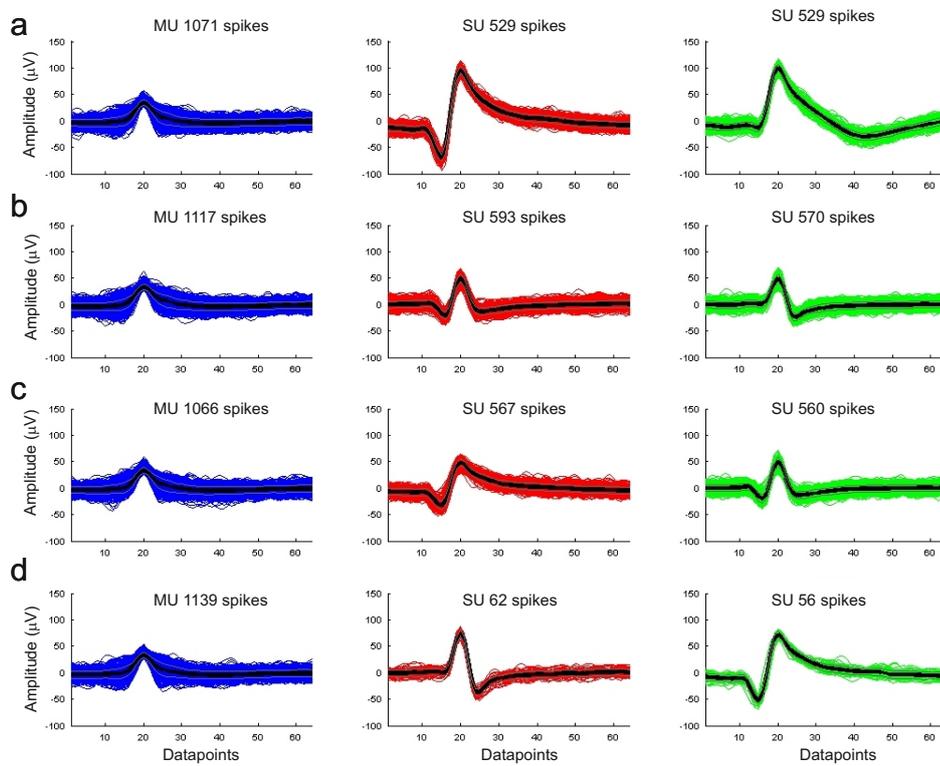


Figure A.2: Clusters of spike shapes as identified by automatic spike sorting for four examples of synthetic extracellular recordings. Each row in the plot represents the clusters of multi-unit (MU) and single-units (SU) for a different simulation example.

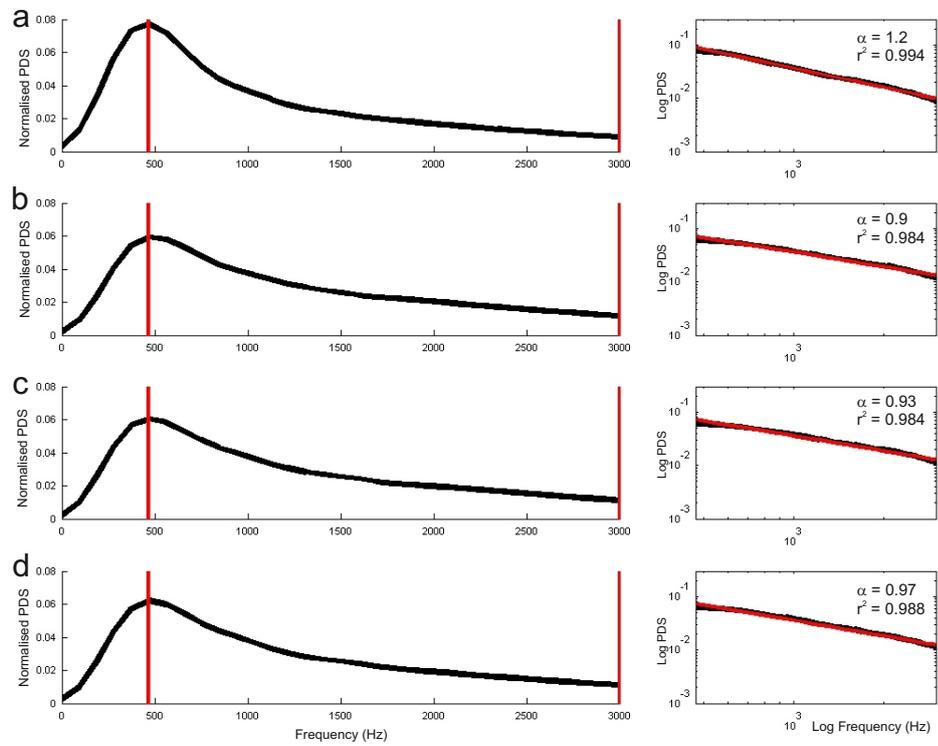


Figure A.3: Power density spectrum for the four simulations of Figure A.2.

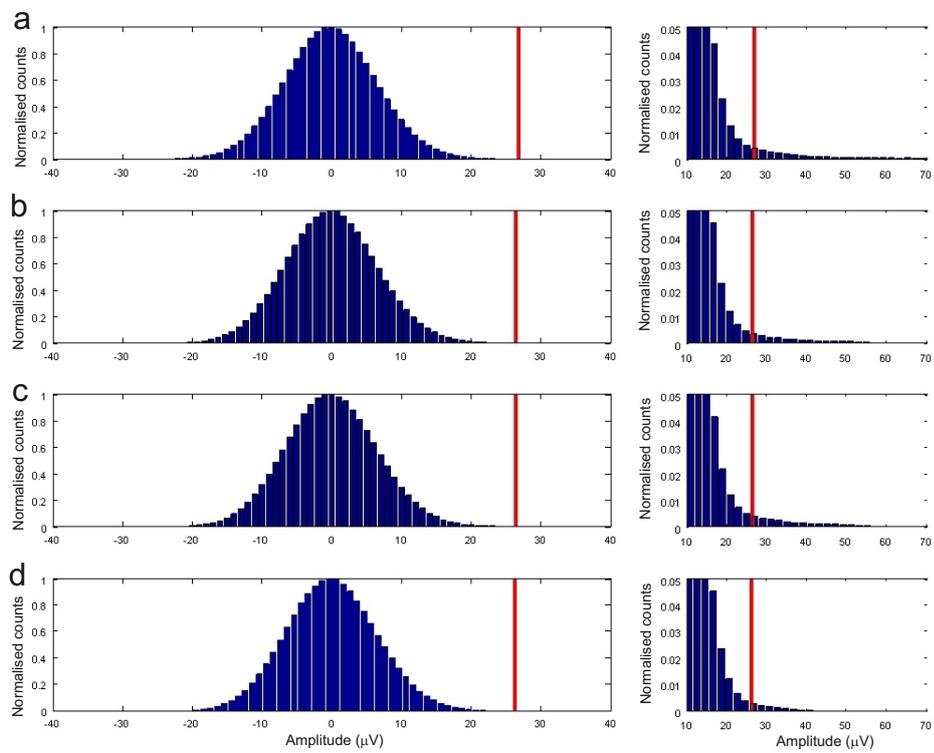


Figure A.4: Amplitude distribution for the four simulations of Figure A.2. The plots on the left side show the distribution of the signal around the detection threshold in more detail.

Bibliography

- Adrian, E. (1928). The basis of sensations. *Norton, New York*.
- Andersen, R., Budrick, J., Musallam, S., Pesaran, B., and Cham, J. (2004). Cognitive neural prosthetics. *Trends Cogn. Sci.*, 8:486–493.
- Averbeck, B., Latham, P., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.*, 7:358–366.
- Bar-Hillel, A., Spiro, A., and Stark, E. (2006). Spike sorting: Bayesian clustering of non-stationary data. *J. Neurosci. Methods*, 157:303–316.
- Bedard, C., Kroger, H., and Destexhe, A. (2006). Model of low-pass filtering of local field potentials in brain tissue. *Phys. Rev. E*, 73:051911.
- Binder, K. and Heermann, D. (1988). Monte carlo simulations in statistical physics: An introduction. *Berlin: Springer-Verlag*.
- Blanche, T., Hetherington, P., Rennie, C., Spacek, M., and Swindale, M. (2003). Model-based 3d cortical neuron localization and classification with silicon electrode arrays. *Soc Neurosci Abstr*.
- Blatt, M., Wiseman, S., and Domany, E. (1996). Super-paramagnetic clustering of data. *Phys. Rev. Lett*, 76:3251–3254.
- Blatt, M., Wiseman, S., and Domany, E. (1997). Data clustering using a model granular magnet. *Neural Computation*, 9:1805–1842.

- Brun, V., Otnæss, M., Molden, S., Steffenach, H., Witter, M., Moser, M., and Moser, E. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science*, 296:2243–2246.
- Brunel, N. and Wang, X. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? *J. Neurophysiol.*, 90:415–430.
- Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neurosc.*, 7:446–451.
- Buzsaki, G. (2006). Rhythms of the brain. *Oxford University Press*.
- Chandra, R. and Optican, L. (1997). Detection, classification, and superposition resolution of action potentials in multi-unit single-channel recordings by an on-line real-time neural network. *IEEE Trans. Biomed. Eng.*, 44:403–412.
- Chelaru, M. and Jog, M. (2005). Spike source localization with tetrodes. *J. Neurosci. Methods*, 142:305–315.
- Claverol-Tinture, E. and Nadasdy, Z. (2004). Intersection of microwire electrodes with proximal cal stratum-pyramidale neurons at insertion for multiunit recordings predicted by a 3-d computer model. *J. Neuroscience Methods*, 51:2211–2216.
- Csicsvari, J., Henze, D., Jamieson, B., Harris, K., Sirota, A., Bartho, P., Wise, K., and Buzsaki, G. (2003). Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J. Neurophysiology*, 90:1314–1323.
- Denk, W., Strickler, J., and Webb, W. (1990). Two-photon laser scanning fluorescence microscopy. *Science*, 6:372–378.
- Desimone, R., Albright, T., Gross, C., and Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *J. Neuroscience*, 4:2051–2062.

- Domany, E. (1999). Super-paramagnetic clustering of data: the definitive solution of an ill-posed problem. *Physica A*, 263:158–169.
- Donoho, D. and Johnstone, I. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455.
- Drake, K., Wise, K., Farraye, J., Anderson, D., and BeMent, S. (1988). Performance of planar multisite microprobes in recording extracellular single-unit intracortical activity. *IEEE Trans. Biomed. Eng.*, 35:719–732.
- Farrant, M., Feldemeyer, D., Takahashi, T., and Cull-Candy, S. (1994). Nmda-receptor channel diversity in the developing cerebellum. *Nat. Lond.*, 368:335–339.
- Fee, M., Mitra, P., and Kleinfeld, D. (1996a). Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-gaussian variability. *J. Neurosci. Methods*, 69:175–188.
- Fee, M., Mitra, P., and Kleinfeld, D. (1996b). Variability of extracellular spike waveforms of cortical neurons. *J. Neurophysiology*, 76:3823–3833.
- Fried, I., MacDonald, K., and Wilson, C. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron*, 18:753–765.
- Gallant, J., Braun, J., and Van Essen, D. (1993). Selectivity for polar, hyperbolic, and cartesian gratings in macaque visual cortex. *Science*, 259:100–103.
- Georgopoulos, A., Schwartz, A., and Kettner, R. (1986). Neuronal population coding on movement direction. *Science*, 233:1416–1419.
- Gerstein, G. and Clark, W. (1964). Simultaneous studies of firing patterns in several neurons. *Science*, 143:1325–1327.
- Goense, J., Zappe, A., and Logothetis, N. (2007). Highresolution fmri of macaque v1. *Magn Reson Med*, 25:740–747.

- Gold, C., Henze, D., Koch, C., and Buzsaki, G. (2006). On the origin of the extracellular action potential waveform: a modeling study. *J. Neurophysiology*, 95:3113–3128.
- Gray, C., Maldonado, P., Wilson, M., and McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J. Neuros. Methods*, 63:43–54.
- Grinvald, A. and Hildesheim, R. (2004). Vsdi: A new era in functional imaging of cortical dynamics. *Nat. Rev. Neurosci*, 5:874–885.
- Gross, C., Bender, D., and Rocha-Miranda, C. (1969). Visual receptive fields of neurons in inferotemporal cortex of the monkey. *Science*, 166:1303–1306.
- Gross, C., Rocha-Miranda, C., and Brender, D. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *J. Neurophysiology*, 35:96–111.
- Hahnloser, R., Kozhevnikov, A., and Fee, M. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419:65–70.
- Halgren, E., TL, B., and Candall, P. (1978). Activity of human hippocampal formation and amygdala neurons during memory testing. *EEG Clini. Neuropsychiol.*, 45:585–601.
- Hamalainen, M., Hari, R., Ilmoniemi, R., Kuutila, J., and Lounasma, O. (1993). Magnetoencephalography -theory, instrumentation, and applications to non-invasive studies of working human brain. *Rev. Mod. Phys*, 65:413–493.
- Harris, K., Henze, D., Csicsvari, J., Hirase, H., and Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiology*, 84:401–414.
- Henze, D., Borhegy, Z., Csicsvari, J., Mamiya, A., Harris, K., and Buzsaki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J. Neurophysiology*, 84:390–400.

- Hetke, J., Lund, J., Najafi, K., Wise, K., and Anderson, D. (1994). Silicon ribbon cables for chronically implantable microelectrode arrays. *IEEE Trans. Biomed. Eng.*, 41:314–321.
- Hines, M. and Carnevale, N. (1997). The neuron simulation environment. *Neural Comp.*, 9:1179–1209.
- Holt, G. and Koch, C. (1999). Electrical interactions via extracellular potential near cell bodies. *J. Comput. Neurosci.*, 6:169–184.
- Hubel, D. (1957). Tungsten microelectrodes for recording single units. *Science*, 125:549–550.
- Hubel, D. and Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol.*, 160:106–154.
- Hulata, E., Segev, R., and Ben-Jacob, E. (2002). A method for spike sorting and detection based on wavelet packets and shannon’s mutual information. *J. Neurosci. Methods*, 117:1–12.
- Ison, M., Balenzuela, P., Bonasera, A., and Dorso, C. (2002). Dynamical properties of constrained drops. *Eur. Phys. J. A*, 14:451–457.
- Ison, M. and Quian Quiroga, R. (2007). Interspike interval distributions, spike sorting, and real experimental conditions. *The British Neuroscience Association Meeting, Harrogate (UK)*.
- Jaccard, P. (1901). Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin de la Société Vaudoise des Sciences Naturelles*, 37:547–579.
- Kreiman, G., Koch, C., and Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nature Neuroscience*, 3:946–953.

- Letelier, J. and Weber, P. (2000). Spike sorting based on discrete wavelet transform coefficients. *J. Neurosci. Methods*, 101:93–106.
- Lewicki, M. (1994). Bayesian modeling and classification of neural signals. *Neural Computation*, 6:1005–1030.
- Lewicki, M. (1998). A review of methods for spike sorting: The detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9:R53–R78.
- Llinas, R. (1988). The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Sci. Wash*, 242:1654–1664.
- Logothetis, N. (2008). What can we do and what cannot do with fmri. *Nature*, 453:869–878.
- Logothetis, N., Pauls, J., and Poggio, T. (1994). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5:552–563.
- Mallat, S. (1989). A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Analysis and Machine Intell.*, 2:674–693.
- Martinez, J., Pedreira, C., Ison, M., and Quiñan Quiroga, R. (2009). Realistic simulation of extracellular recordings. *J. Neurosci. Methods*, 184:285–293.
- Mazzoni, A., Panzeri, S., Logothetis, N., and Brunel, N. (2008). Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PloS Comput. Biol*, 4:1–20.
- McCormick, D., Connors, B., Lighthall, J., and Prince, D. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophysiology*, 54:782–806.

- McNaughton, B., O'Keefe, J., and Barnes, C. (1983). The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J. Neurosci. Methods*, 8:391–397.
- Musallam, S., Corneil, B., Greger, B., Scherberger, H., and Andersen, R. (2004). Cognitive control signals for neural prosthetics. *Science*, 305:258–262.
- Nicolelis, M. (2001). Actions from thoughts. *Nature*, 409:403–407.
- Nicolelis, M. and Ribeiro, S. (2002). Multielectrode recordings: the next steps. *Current Opinion in Neurobiology*, 12:602–606.
- Ogawa, S., Lee, T., Nayak, A., and Glynn, P. (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn Reson Med*, 14:68–78.
- O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp. Neurology*, 51:78–109.
- O'Keefe, J. (1979). A review of the hippocampal place cells. *Prog. Neurology*, 13:419–439.
- O'Keefe, J. and Recce, M. (1993). Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3:317–330.
- Optican, L. and Richmond, B. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. iii. information theoretic analysis. *J. Neurophysiology*, 57:162–178.
- Perez-Orive, J., Mazor, O., Turner, G., Cassenaer, S., Wilson, R., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science*, 297:359–365.
- Plonsey, R. (1969). Bioelectric phenomena. *New York: McGraw-Hill*.

- Pouzat, C., Mazor, O., and Laurent, G. (2002). Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J. Neurosci. Methods*, 122:43–57.
- Quian Quiroga, R. (2007). Spike sorting. *Scholarpedia*, 2:3583.
- Quian Quiroga, R. (2009). What is the real shape of extracellular spikes? *J. Neurosci. Methods*, 177:194–198.
- Quian Quiroga, R., Mukamel, R., Isham, E., Malach, R., and Fried, I. (2008). Human single-neuron responses at the threshold of conscious recognition. *PNAS*, 105:3599–3604.
- Quian Quiroga, R., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Computation*, 16:1661–1687.
- Quian Quiroga, R. and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci*, 10:173–185.
- Quian Quiroga, R., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single-neurons in the human brain. *Nature*, 435:1102–1107.
- Quian Quiroga, R., Snyder, L., Batista, A., Cui, H., and Andersen, R. (2006). Movement intention is better predicted than attention in the posterior parietal cortex. *J. Neuroscience*, 26:3615–3620.
- Rall, W. (1962). Electrophysiology of a dendritic neuron model. *Biophys. J*, 2:145–167.
- Ramon y Cajal, S. (1909). Histologie du systeme nerveux de l’homme et des vertebres. *Paris, Maloine*.

- Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997). Spikes: Exploring the neural code. *MIT Press. Cambridge, Massachusetts.*
- Robinson, D. (1968). The electrical properties of metal microelectrodes. *Proc. of the IEEE*, 56:1065–1071.
- Rousche, P. and Normann, R. (1998). Chronic recording capability of the utah intracortical electrode array in cat sensory cortex. *J. Neuroscience Methods*, 82:1–15.
- Serruya, M., Hatsopoulos, N., Paninski, L., Fellows, M., and Donoghue, J. (2002). Brain-machine interface: instant neural control of a movement signal. *Nature*, 416:141–142.
- Shadlen, M. and Newsome, W. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neuroscience*, 18:3870–3896.
- Shoham, S., Fellows, M., and Normann, R. (2003). Robust, automatic spike sorting using mixtures of multivariate t-distributions. *J. Neurosci. Methods*, 127:111–122.
- Shoham, S., O’Connor, D., and Segev, R. (2006). How silent is the brain: Is there a dark matter problem in neuroscience? *J. Comp. Physiology*, 192:777–784.
- Smith, L. and Mtetwa, N. (2007). A tool for synthesizing spike trains with realistic interference. *J. Neurosci. Methods*, 159:170–180.
- Snider, R. and Bonds, A. (1998). Classification of non-stationary neural signal. *J. Neurosci. Methods*, 84:155–166.
- Snyder, L., Batista, A., and Andersen, R. (1997). Coding of intention in the posterior parietal cortex. *Nature*, 386:167–170.
- Softky, W. and Koch, C. (1992). Cortical cells should fire regularly, but do not. *Neural Computation*, 4:643–646.

- Stark, E. and Abeles, M. (2007). Predicting movement from multiunit activity. *J. Neuroscience*, 27:8387–8394.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.*, 19:109–139.
- Victor, J. and Purpura, K. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiology*, 76:1310–1326.
- Wehr, M., Pezaris, J., and Sahani, M. (1999). Simultaneous paired intracellular and tetrode recordings for evaluating the performance of spike sorting algorithms. *Neurocomputing*, 26:1061–1068.
- Wilson, M. and McNaughton, B. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058.
- Wolf, U. (1989). Comparison between cluster monte carlo algorithms in the ising spin model. *Phys. Lett. B*, 228:379–382.
- Wood, F., , Fellows, M., Donoghue, J., and Black, M. (2004a). Automatic spike sorting for neural decoding. *Proceedings of the 26th Annual International Conference of the IEEE EMBS. San Francisco, CA, USA*.
- Wood, F. and Black, J. (2008). A nonparametric bayesian alternative to spike sorting. *J. Neurosci. Methods*, 173:1–12.
- Wood, F., Black, J., Vargas-Irwin, C., Fellows, M., and Donoghue, J. (2004b). On the variability of manual spike sorting. *IEEE Trans. Biomed. Eng.*, 51:912–918.
- Wu, F. (1982). The potts model. *Reviews of Modern Physics*, 54:235–268.