

Experimental Evidence for Cortical Phase Transitions in Natural Sleep

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

Master of Science

in Physics

University of Waikato

by

Iain-Patrick Gillies



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2006

Dedication

I dedicate this thesis–

To my wife Sasha

To my son Sean and daughter Isla

To my other children who are not yet with us

To my mother Catherine

To my father Daren

To my grandparents Iain and Flora

Thank you for all your love, support and inspiration.

Abstract

Sleep can be considered to consist of two very different phases: slow-wave sleep (SWS) and rapid-eye-movement (REM) sleep. Recently, Steyn-Ross *et al.* (2005) have developed a mean-field model describing the average behaviour of populations of neurons in the cerebral cortex, and applied this model to sleep. The model predicts that the transition between SWS and REM is a first-order phase transition, analogous to ice melting into water. As a system crosses the critical point of the phase transition, its fluctuations show discontinuous changes in power, correlation and spectral distribution. These are hallmarks of all first-order phase transitions.

In this thesis, the scalp recorded voltage fluctuations, the electroencephalogram (EEG), of sleeping humans and rats are examined to provide evidence of the phase transitions predicted by the Steyn-Ross *et al.* (2005) model. Using extensive, original MATLAB code, the previously mentioned statistics are calculated on the EEG time-series and the SWS to REM transitions are identified. The results show high values for the total power, low-frequency power and correlation time of the EEG during SWS, on approach to the transition. At the transition into REM, these high values suddenly drop, giving strong evidence that it is indeed a first-order phase transition. While identifying abrupt changes in DC-coupled EEG, it is observed that one-way, low-pass filtering a DC-shift produces waveforms that appear very similar to another phenomenon: the K-complex.

Steyn-Ross *et al.* [*Journal of Biological Physics*, **31**, 3/4 ,547–569 (2005)]

Acknowledgements

I would like to thank all those who have helped me with my research. Firstly, a big thanks to Dr Alistair and Dr Moira Steyn-Ross for originally encouraging me to do this MSc project. Thanks again Alistair, for your supervision and helpful advice. Thank you Dr Marcus Wilson for supervising in the last few months of the write-up. Thank you Prof Jamie Sleigh for your ideas and Dr Logan Voss for providing me with data. Thank you Dr Michael Cree for assistance with \LaTeX and Unix whenever it was required. A final, last but not least, thank you to Fisher & Paykel Healthcare, for providing data and, with the Tertiary Education Commission, supplying the funding of the Bright Futures Scholarship.

Contents

Dedication	iii
Abstract	v
Acknowledgements	vii
List of Figures	xv
List of Tables	xvii
Acronyms and Abbreviations	xix
1 Introduction	1
1.1 Thesis Overview	3
1.2 Thesis Structure	3
1.3 Original Work	4
2 Cortical Modelling Theory	7
2.1 Basic Neurophysiology	7
2.1.1 Neuronal Operation	7
2.1.2 Voltages of Neurons	9
2.2 The Macrocolumn and the EEG	10
2.3 Cortico-Dynamic Equations	12
2.3.1 Integral Equations	12
2.3.2 Converting the Integrals to Derivatives.	15
2.3.3 Approximations and Stability of the Steady States	18
2.4 Expanding from Anaesthesia to Natural Sleep	20
2.4.1 The Anaesthesia Model	20
2.4.2 Modelling Natural Sleep	21
2.4.3 Predictions of the Steyn-Ross Model	24
3 EEG Recording and Filtering	29
3.1 The Raw Sleep Recordings	29
3.1.1 Electrode Configuration	29
3.1.2 Sleep Stages	29
3.2 Human Sleep Recordings	33
3.3 The Binary Data Format	34

3.4	Filtering Electrical Interference	34
3.4.1	Filtering Human EEG	34
3.4.2	Filtering Rat EEG	38
4	Time Series Signal Processing	39
4.1	Total Power	39
4.2	Fractional Band Power	40
4.3	Correlation Time	41
4.4	Spectral Entropy	44
4.5	Epoch Overlap	45
4.5.1	A New Overlapping Algorithm	46
4.5.2	Overlap Time Trials	49
5	Human Data Results	51
5.1	Sleep Staging	52
5.2	Pre-processing – Artefacts Detection	54
5.3	Full Sleep Record Results	55
5.3.1	Patient One	58
5.3.2	Patient Two	63
5.3.3	Patient Three	66
5.3.4	Patients Four, Five and Six	70
5.4	How Does ECG Affect the Results?	77
5.5	Summary of the Human Results	79
6	Rat Data Results	81
6.1	Full Rat Time-Series	82
6.2	Time-Aligned Transitions	87
6.2.1	Locating Transitions	87
6.2.2	Differences in ECoG Amplitude	88
6.2.3	Results of the Time-Aligned Transitions	90
6.3	Looking for the DC Shifts	90
6.3.1	Removing the DC Drift	92
6.3.2	How Many Transitions Show a DC Shift?	94
6.3.3	Should Transition 11 be Included?	94
6.4	A K-complex or a DC-shift?	99
6.5	Summary of Rat Results	101
7	DC-Shift Detection Algorithm	103
7.1	Developing the New Algorithm	103
7.1.1	The Two Least-Squares Lines	104
7.1.2	Varying Δt to Make a Surface	104
7.2	Applying the Algorithm to the DC Transitions	107
7.3	How Well do the Lines Fit the Data?	112
7.4	DC-shift Detection – A Work in Progress	114

8	Conclusions and Future Work	119
8.1	Unfinished Work	120
8.2	Future Work	121
A	Matlab Codes Written	123
A.1	Time-Series Statistics Functions	123
A.1.1	Function: sleeppower	123
A.1.2	Function: fract_pwr	124
A.1.3	Function: correlation_time	124
A.1.4	Function: entropy	125
A.2	Utility Software Tools	126
A.2.1	Matlab Graphics	126
A.2.2	Data Display Software	127
A.2.3	Function: showline	129
A.2.4	Function: zoomall	131
A.2.5	Function: extreme	133
A.3	Data Smoothing	136
A.3.1	Convolution Moving Average	136
A.3.2	Whittaker Difference Smoothing	141
A.3.3	Comparison of Convolution and Whittaker Smoothing . . .	143
B	Extra Rat Transition Results	147
	References	155

List of Figures

2.1	A schematic diagram of a neuron	8
2.2	The macrocolumn	11
2.3	The weights ψ_{jk}	13
2.4	The impulse response of the soma and dendrites	15
2.5	Sigmoid function to relate soma voltage to firing rate	17
2.6	Simulated EEG time-series of an anaesthetised cortex	21
2.7	The EEG of a sleeping cat (Destexhe <i>et al.</i> , 1999)	22
2.8	The steady states manifold of the natural sleep model	24
2.9	Bird's-eye view of the sleep manifold	25
2.10	The predicted changes in EEG during a sleep cycle	26
2.11	A SWS to REM transition from numerically simulated EEG	27
3.1	The International 10-20 System of Electrode Placement.	30
3.2	Electrode placements for sleep studies	32
3.3	Sample EEG from the different stages of sleep	33
3.4	Magnitude and Phase response for 49 Hz low pass filter	36
3.5	The PSD of filtered and unfiltered EEG	37
3.6	The PSD of Rat EEG	38
4.1	The fraction band power of test data	41
4.2	Evaluating the correlation time	43
4.3	The spectral entropy of test data	45
4.4	Overlapping time epochs by 3/4	47
4.5	Overlapping time epochs by 1/3	48
4.6	Overlapping epochs can fill in the blanks	49
4.7	Computation time for the overlap algorithm	50
5.1	Possible REM classified as stage 2	53
5.2	Possible stage 2 classified as REM	53
5.3	A movement artefact during stage 3 sleep	55
5.4	Total power of EEG containing artefacts	56
5.5	Total power of EEG that does not contain artefacts	56
5.6	Filtering EEG to remove REM artefacts	57
5.7	The time-series analysis of the sleep recordings of patient one	59
5.8	The fractional band power of patient one	61
5.9	Spindles are detected by their short correlation time	62
5.10	The time-series analysis of the sleep recordings of patient two	64

5.11	The fractional band power of patient two	65
5.12	The time-series analysis of the sleep recordings of patient three . . .	67
5.13	The fractional band power of patient three	68
5.14	A very quick SWS to REM transition from the records of patient 3 .	69
5.15	The time-series analysis of the sleep recordings of patient four . . .	71
5.16	The fractional band power of patient four	72
5.17	The time-series analysis of the sleep recordings of patient five . . .	73
5.18	The fractional band power of patient five	74
5.19	The time-series analysis of the sleep recordings of patient six	75
5.20	The fractional band power of patient six	76
5.21	ECG added to EEG to test how the statistics are affected	77
5.22	ECG present in EEG changes the time-series statistics	78
6.1	A lab rat wearing the surgically attached electrode headset	81
6.2	Time-series statistics for rat 6 ₃	84
6.3	Time-series statistics for rat 12 ₂	85
6.4	Time-series statistics for rat 12 ₃	86
6.5	A function used for time-aligning transitions	88
6.6	The rms voltage levels of the rat ECoG	89
6.7	The time-aligned SWS to REM transitions of rat recording 6 ₁	91
6.8	Removing drift from DC ECoG	92
6.9	The DC ECoG before and after drift removal	93
6.10	DC ECoG that show DC-shifts at the SWS to REM transition	95
6.11	Average DC ECoG including transition 11	97
6.12	Average DC ECoG excluding transition 11	97
6.13	The data surrounding transition 11 contains many artefacts	98
6.14	Comparison between a K-complex, and high-pass filtered DC-shift .	100
6.15	Magnitude and phase response of a high-pass filter	101
7.1	Fitting two least-squares lines to data to detect a DC-shift	105
7.2	The DC-shift corresponds to the maximum in $\Delta V(t)$	106
7.3	The length of the least-squares lines varies sensitivity	106
7.4	A surface plot of $\Delta V(t, \Delta t)$	108
7.5	The average across Δt of the surface ΔV	109
7.6	DC-shifts detected in the transitions 1–7	110
7.7	DC-shifts detected in the transitions 8–15	111
7.8	Norm of residuals to measure the quality of the fit	113
7.9	Clear DC-shifts are better detected by dividing ΔV by R_{norm}	114
7.10	DC-shifts detected in the transitions 1–7 using the norms	115
7.11	DC-shifts detected in the transitions 8–15 using the norms	116
7.12	Histogram showing the positions of all detected DC-shifts	117
A.1	The arrangement of the array that is returned by <code>fract_pwr</code>	125
A.2	A screen shot of the program <code>sleepwork</code>	128
A.3	Autocorrelation and PSD opened by <code>sleepwork</code>	129
A.4	An example of the operation of <code>showline</code>	130
A.5	An example of <code>showline</code> in action	132

A.6	A diagram illustrating how extreme works	135
A.7	Using extreme to subsample raw EEG	137
A.8	Data smoothed with a convolution moving average	140
A.9	Edge effects of the convolution moving average	141
A.10	Data smoothed with a Whittaker smoother	144
A.11	Derivatives of the Whittaker and convolution-smoothed data	145
A.12	Convolution smoothing can fluctuate in anti-phase with the data . .	146
B.1	The time-aligned SWS to REM transitions of rat recording 6 ₃	148
B.2	The time-aligned SWS to REM transitions of rat recording 6 ₄	149
B.3	The time-aligned SWS to REM transitions of rat recording 11 ₂	150
B.4	The time-aligned SWS to REM transitions of rat recording 12 ₁	151
B.5	The time-aligned SWS to REM transitions of rat recording 12 ₂	152
B.6	The time-aligned SWS to REM transitions of rat recording 12 ₃	153

List of Tables

2.1	Ion concentrations and Nernst potentials for a resting neuron	9
2.2	Typical model constants for the neural macrocolumn	14
3.1	Cortical lobes labelled in the 10-20 system	30
4.1	Computation time for the overlap algorithm	50
5.1	Frequency bands commonly used in EEG analysis.	58

Acronyms and Abbreviations

AC	alternating current
ACh	acetylcholine (an neuromodulator)
DC	direct current
ECG	electrocardiogram (heart beats)
ECoG	electrocorticogram (brain activity recorded via electrodes inserted through the skull, into the cortex)
EEG	electroencephalogram (brain activity recorded via scalp electrodes)
EMG	electromyogram (muscle tone)
EOG	electro-oculogram (eye movement activity)
EPSP	excitatory postsynaptic potential
IPSP	inhibitory postsynaptic potential
LOC	left ocular (EOG from left eye)
PSD	power spectral density
REM	Rapid-eye-movement
ROC	right ocular (EOG from right eye)
rms	root-mean-square
SWS	slow-wave sleep

Chapter 1

Introduction

Sleep is essential to everyday life, but there is still much about it that remains unknown. Psychologists initially assumed that the brain, particularly the cerebral cortex, was idle during sleep and active during wakefulness. Dement and Kleitman (1957) recorded the electroencephalogram (EEG) — the electrical activity of the cortex measured with electrodes attached to the scalp — of a multitude of sleeping adult humans. They reported that “discrete periods of rapid eye movement potentials” coincided with high-activity EEG that is similar to waking EEG. They found that sleep went through cycles between two phases: low activity, with no eye-movements and high activity, with rapid eye-movements. They then divided sleep into two main categories: rapid-eye-movement (REM) sleep, and non-REM sleep. Non-REM was further sub-categorised, from lightest sleep to deepest sleep, as stages 1, 2, 3 and 4.

When people fall asleep they begin with stage 1 or 2, and usually proceed up to stage 3 and stage 4. The stages generally progress with 1–4 Hz oscillations increasing in amplitude and occurrence. In stages 3 and 4, the 1–4 Hz slow-waves are the dominating feature. Non-REM sleep has also come to be known as EEG-synchronised sleep, or slow-wave sleep (SWS). REM sleep, because its EEG looks similar to the waking EEG, is also known as paradoxical sleep. (EEG examples for the different sleep states are shown in Fig. 3.3 on Pg. 33)

While many were trying to mark the similarities between paradoxical and EEG-synchronised sleep, Dement commented that the two “are as far as night and day”. He went on to state, “I would even go so far as to suggest that there may be some validity in questioning whether they should be subsumed under the general heading of sleep.” (Steriade and McCarley, 2005, Pg. 9).

With advances in technology, there has been a big leap in the understanding of the differences between SWS and REM, particularly at the neurophysiological

level. By surgically implanting the electrodes into the brains of mammals such as rats and cats, voltage signals can be recorded from the cortex and other structures such as the thalamus. Research has also been performed on microscopic slices of the brain tissue of animals. Individual neurons are chemically or electrically stimulated and their responses are measured. During SWS there is a decreased amount of information transfer from the outside world, but the cortico-cortico and cortico-thalamic messages are still preserved (Steriade and McCarley, 2005, Pg. 10). In REM sleep the EEG is in an activated state, where the neurons in the cortical network are brought closer to the threshold where they can respond quickly to messages from the outside world. The paradox of REM sleep is that although the cortex is in this activated state, the motoric arousal threshold is much greater than it is for SWS (Steriade and McCarley, 2005, Pg. 19).

Recently, research into the sleeping states have used models of the brain to describe the waveforms observed in the EEG. The different approaches of modelling can be classified into two groups. One is the *neural network*, where the individual neurons and their connections are modelled. This approach has been used by Hill and Tononi (2005) and Bavhenov *et al.* (2002). The other type of models, first used by Nunez (1974) and by Freeman (1992), are continuum, or *mean-field* models, where populations of tens of thousands of neurons are grouped together and averages are computed. Taking averages of this many neurons may seem like a huge approximation, but the EEG itself is an average of the voltages of the neurons directly beneath the electrode. Other contributors to mean-field models include Wright and Liley (1996), Robinson *et al.* (1997), Liley *et al.* (1999) and Rennie *et al.* (2000).

Steyn-Ross *et al.* (1999, 2001a,b) have taken the mean-field model of the cortex by Liley *et al.* (1999), a system of coupled differential equations describing the average voltages of the neurons, and adapted it to model the effects of anaesthetic agents such as propofol. They solved their system of equations as a function of anaesthetic concentration and found that for low levels and high levels of anaesthetic, there was only one solution, but at intermediate concentrations, there were three solutions. Solutions of this form lead to phase transitions, which are commonly seen in thermodynamics. A stability analysis revealed that where three solutions existed, only two of them were stable. The model showed that if the concentration of the anaesthetic was increased enough, the cortex would make a discontinuous jump from the conscious state to the unconscious state. Once the cortex was anaesthetised, the cortex would not make the jump back to consciousness until the anaesthetic concentration was reduced to a much lower level. The jumps in and out of consciousness seen in the Steyn-Ross *et al.* (1999, 2001a,b) model are first-order phase transitions,

analogous to ice melting into water.

More recently, Moira and Alistair Steyn-Ross have become interested in natural sleep, and whether the change from SWS to REM may also be a phase transition. In their paper, Steyn-Ross *et al.* (2005) design a new model to account for the neuro-modulator effects that have been observed during sleep using the same philosophy as their anaesthesia model. The natural sleep model does indeed predict the existence of a phase transition between SWS and REM.

Dement claimed that SWS and REM were as far apart as night and day. In this thesis, I look for phase transitions between SWS and REM and try to see if they are as far apart as ice and water.

1.1 Thesis Overview

The Steyn-Ross *et al.* (2005) model for natural sleep makes a number of predictions describing how the characteristics of the EEG fluctuations will change as the cortex undergoes the SWS to REM phase transition. The primary aim of my research was to analyse real EEG obtained from both human sleep subjects and laboratory rats, and attempt to find evidence of the cortical phase transitions. At the SWS to REM transition, the model predicts sharp discontinuous changes in statistical values calculated from the EEG fluctuations. The same statistics are calculated on real EEG and the SWS to REM transitions are identified and compared with phase transitions. The model also predicts that there will be a quick upwards jump in the DC voltage level of the EEG at the phase transition. To test for this, EEG of the rats had been recorded down to zero frequency. A new algorithm was developed for detecting these DC-shifts.

1.2 Thesis Structure

In Chapter 2 I present the relevant theory necessary for describing the Steyn-Ross *et al.* (2005) model. This includes: elementary neurophysiology, the Liley *et al.* (1999) equations on which the Steyn-Ross model is based, stability analysis, a previous anaesthesia model and its predictions (Steyn-Ross *et al.*, 2001a,b), and finally the natural sleep model and its predictions.

Chapter 3 contains a summary of common practices involved with sleep studies, particularly recording EEG and other bio-signals and categorising sleep into its stages. I also give details specific to my EEG data, such as how, where and from whom the recordings were obtained. The chapter is concluded with a small section

describing the filters that were used in pre-processing the EEG.

Chapter 4 describes the time-series statistics that were calculated for the analysis of the EEG data. The analytical equations given are not my own work, but I did write the functions that were used to implement them myself. The four time-series statistics are: total power, frequency band power fractions, correlation time and spectral entropy. During the process of writing the functions I devised a method of overlapping the time epochs very efficiently. This is explained at the end of the chapter along with results of a quick experiment measuring its computation time.

In Chapter 5 the measures taken for detecting and removing artefacts from the signal are described. The figures showing the results of the time-series analysis of the human EEG are then presented and the SWS to REM transitions are described. The chapter also gives a small discussion on how the heart beats of the electrocardiogram (ECG) can be found in the EEG, and how these artefacts affect the time-series statistics.

Chapter 6 presents the results of the time-series statistics that were calculated on the fluctuations of the rat EEG. The data is presented as a time-series of the full recording, and also as averages of small data subsets that have been time-aligned so the SWS to REM transitions occur simultaneously. The focus of the chapter then moves on to the detection of DC-shifts. The transitions that do and do not show DC-shifts at the SWS to REM transition are discussed. Towards the end of the chapter is a section describing a certain waveform known as the K-complex. I make an observation and give results indicating the possibility that another phenomenon could be misclassified as a K-complex.

Chapter 7 gives the details of a DC-shift detection algorithm that I have developed. I give the basic idea behind the algorithm and the results it produces in its current form. I conclude the chapter by listing the limitations of the algorithm and giving suggestions on how it could be improved.

Chapter 8 is a summary of the conclusions made from the research presented in this thesis. It also contains things that did not work so well and that I would change if more time were available. It then concludes with future research topics to provide answers to questions that have arisen during work on the current topic.

1.3 Original Work

For the sake of completeness, this thesis briefly presents work that is not my own. However, the majority of this thesis is my own original work. Specifically, my original contributions are:

- All the necessary code that has been used to assist with data display and analysis with the exception of a few toolboxes (a summary of the time-series signal processing functions and other utility codes that have been written is given in Appendix A).
- All figures presented in this thesis, unless otherwise stated. Most of the reprinted figures are found in Chapter 2.
- The overlapping algorithm that is implemented in the time-series statistics functions (Chapter 4).
- Artefact detection and removal from the EEG time-series (Chapter 5).
- The calculation and analysis of the time-series statistics (Chapters 5 and 6).
- The final alignment of the time-aligned transitions (Chapter 6). Logan Voss extracted the data subsets containing the SWS to REM transitions from the full time-series.
- Analysis of the rat DC data (Chapter 6).
- The observation of the similarities between low-pass filtered DC-shifts and K-complexes (Chapter 6).
- The development of the DC-shift detection algorithm, with initial inspiration from Jamie Sleight and Alistair Steyn-Ross (Chapter 7).

Chapter 2

Cortical Modelling Theory

This chapter summarises the theory that has been adopted and developed in the Steyn-Ross model of the cortex. The equations and concepts described here are explained in greater detail in previous papers (Steyn-Ross *et al.*, 1999, 2001a,b, 2003, 2004, 2005; Liley *et al.*, 1999; Wilson *et al.*, 2005) and theses (Steyn-Ross, 2002; Whiting, 2003).

2.1 Basic Neurophysiology

The human brain consists of many different parts such as the cerebellum, brain stem, thalamus and cerebrum. The largest of these is the cerebrum. Its outer layer of grey matter, about 2–3 mm thick, is called the cerebral cortex. It is twisted and folded to compact its large surface area into a small volume. The cerebral cortex is responsible for sensing, thinking, learning, emotion, voluntary movements and consciousness (Weiten, 1992).

2.1.1 Neuronal Operation

The brain cells that make up the cortex are called neurons. A typical neuron is shown in Fig. 2.1. Many neurons connect together, sending and receiving electronic signals. The main body of the neuron is the *soma*. Signals leave the soma through the *axon*, then at the *synapses*, they jump to the *dendrites* of other neurons. A single neuron may have as many as 15,000 synaptic connections (Weiten, 1992). The dendrites of the neuron collect the synaptic input voltages and bring the signals back to the soma where they are integrated together. When the voltage of the soma increases above a threshold, a large voltage spike (action potential) is generated and it travels down the axon. At the end of the axon, it branches out to many




Figure 2.1: A schematic diagram of a neuron. The dendrites of a neuron collect (1) excitatory or (2) inhibitory voltage pulses that have come from a presynaptic neuron. The pulses propagate down the dendritic tree and are combined at the cell body (soma). If the soma voltage is raised above the threshold voltage a large voltage spike, called an action potential (3a), is released and it travels down the axon. At each of the nodes of Ranvier (3b, 3c), the action potential is boosted again. At the end of the axon the action potential propagates down the to the axon terminals where neurotransmitters are released to stimulate more neurons (4). [image source: Hammond (2001)]

terminals which connect to the dendrites of more neurons via more synapses. At the synapse, neurotransmitter chemicals are released from the *presynaptic* neuron, and are received by the *postsynaptic* neuron. There will be either a positive (excitatory) or negative (inhibitory) voltage pulse, which is called a *postsynaptic potential* (PSP), in the dendrites after they have received the neurotransmitter. A positive PSP (EPSP) arises from stimulation from an excitatory neuron, while a negative PSP (IPSP) comes from an inhibitory neuron. In the cortex, roughly 85% of the neurons are excitatory and 15% are inhibitory (Braitenberg and Schüz, 1991).

2.1.2 Voltages of Neurons

If the potential difference between the intracellular and extracellular fluids of a cortical neuron were measured, one would record a voltage of approximately $V = -70$ mV. This is the resting potential of the neuron. The voltage difference is due to the fact that the intracellular and extracellular fluids have different ionic concentrations. Inside the cell there is an abundance of K^+ , while outside the cell the concentrations of Na^+ and Cl^- are high.

The Nernst potential V_X is the voltage across the membrane due to the difference in concentrations of ion X. It is defined as:

$$V_X = V_i - V_o = \frac{RT}{zF} \ln \frac{[X]_o}{[X]_i}, \quad (2.1)$$

where $R = 8.314 \text{ J K}^{-1} \text{ mol}^{-1}$ is the ideal gas constant, T is the temperature, z is the ion charge (i.e., Na^+ has a charge of 1), and $F = 9.648 \times 10^4 \text{ C mol}^{-1}$ is the Faraday constant. The square brackets $[X]$ mean the concentration of X, and the subscripts i and o denote inside and outside the cell respectively. Table 2.1 lists the concentrations and Nernst potentials for the common ions.

Table 2.1: Ion concentrations (in mmol/L) and Nernst potentials (mV) for a resting human neuron. The Nernst potentials were calculated using Eq. (2.1) with $T = 37^\circ\text{C}$, which gives $RT/F = 26.7$ mV. [Values from Steyn-Ross (2002)]

Ion	Concentration		Nernst Potential
	Inside	Outside	
Na^+	10	140	+70
K^+	140	4	−95
Cl^-	4	103	−87

The differences in concentrations of the ions are needed to form action poten-

tials. When the potential difference across the soma membrane (soma voltage) exceeds a certain threshold, the sodium ion channels open. Na^+ rush into the neuron to balance out the concentrations, consequently raising the soma voltage. At the peak of the action potential, the potassium ion channels open and allow K^+ to escape, which lowers the soma voltage again. The concentration gradients are maintained by *ion pumps*, which bring two K^+ ions into the cell and expel three Na^+ ions. The Nernst potential is also called the *reversal* potential because the direction of ion flow will reverse if the soma voltage is greater than the Nernst potential.

The Goldman–Hodgkin–Katz formula has a similar form as Eq. (2.1), but it combines the concentrations of the ions $[X]$ and their permeabilities P_X together.

$$V = \frac{RT}{F} \ln \left[\frac{P_K[\text{K}]_o + P_{\text{Na}}[\text{Na}]_o + P_{\text{Cl}}[\text{Cl}]_i}{P_K[\text{K}]_i + P_{\text{Na}}[\text{Na}]_i + P_{\text{Cl}}[\text{Cl}]_o} \right] \quad (2.2)$$

The exact values of the permeabilities need not be known, just their ratios. The permeabilities $P_K : P_{\text{Na}} : P_{\text{Cl}}$ with respect to P_K are 1 : 0.05 : 0.25. At a temperature of 37°C, $V = -70.2$ mV, which is the resting potential.

2.2 The Macrocolumn and the EEG

Because the cortex contains billions of neurons and $\sim 10^{14}$ synapses, modelling each connection would be impossible. Rather than attempting this unfeasible task, the Steyn-Ross model opts for the mean-field approach where neurons occupying the same volume are grouped together to give averages. The Steyn-Ross model is an adaption of the Liley model (Liley *et al.*, 1999). The basic building block of both models is the *macrocolumn*. The macrocolumn is a cylindrical volume, ~ 0.3 – 1 mm in diameter and as deep as the cortical tissue ~ 2 – 3 mm. A schematic diagram is shown in Fig. 2.2.

The voltages of the excitatory and inhibitory neurons are averaged and labelled as V_e and V_i respectively. When an action potential is fired down the axon, a dipole is created. As seen in Fig. 2.2, the excitatory neurons are aligned so that their dendrites and axons are roughly perpendicular to the scalp, so the dipoles sum together. The inhibitory neurons are randomly orientated so their dipoles approximately cancel. Therefore, the EEG measured from the scalp can be broadly considered to arise from the average behaviour of the excitatory neurons over a small area immediately below the electrode.




Figure 2.2: The macrocolumn is a cylindrical volume of the cortex containing $\sim 100,000$ neurons (only four are shown). $\sim 85\%$ are (pyramidal) excitatory neurons which are actually triangular in shape and line up with their axons pointing away from the scalp. The other $\sim 15\%$ are (stellate or basket) inhibitory neurons which have a more circular shape and are aligned arbitrarily. Outputs connections from an excitatory neuron ($e \rightarrow e$, $e \rightarrow i$) can be short-range, long-range and subcortical, but inhibitory neuron output connections ($i \rightarrow e$, $i \rightarrow i$) can only be short-range.

[image source: Steyn-Ross (2002)]

2.3 Cortico-Dynamic Equations

The original equations describing the dynamics of the macrocolumn used in the Steyn-Ross *et al.* (1999) model, were the same as were used by Liley *et al.* (1999). Since then there have been some small modifications. This section follows the theory presented in one of their more recent papers (Steyn-Ross *et al.*, 2005).

The model consists of a set of eight differential equations, in two-dimensional space and time. In these equations, the subscripts e and i denote excitatory and inhibitory neurons. In some of the terms, there are two subscripts (e.g., in ψ_{ei}). The first subscript is the type of the presynaptic neuron, while the second is the postsynaptic neuron.

2.3.1 Integral Equations

The first two of the eight equations describe how the average voltages of the excitatory and inhibitory macrocolumn, V_e and V_i , react in time to dendritic inputs. Before stating them in differential form, it helps to understand them by first writing them as integrals. Here, all the excitatory and inhibitory PSPs are convolved with the soma impulse response function L and then integrated together to give the soma voltage. The two equations are identical apart from the subscripts e and i of the postsynaptic neuron, so only one equation has been written and k can take the value e or i .

$$V_k = V_k^{\text{rest}} + \int_{-\infty}^t L_k(t-t') [\rho_e \psi_{ek}(t) \Phi_{ek}(t) + \rho_i \psi_{ik}(t) \Phi_{ik}(t)] dt' \quad (2.3)$$

The impulse response L_k is defined as:

$$\frac{1}{\tau_k} \exp[-t/\tau_k], \quad t \geq 0 \quad \text{soma response} \quad (2.4)$$

where τ_k is time constant of the decay curve.

The terms $\rho_j \psi_{jk} \Phi_{jk}$ (where $j, k \in \{e, i\}$), in Eq. (2.3), are input perturbations from other neurons. The product $\rho_e \Phi_{ek}(t)$ is the excitatory voltage input to a type k neuron. The strength of the synapse is modelled by ρ , which has units mV.s and Φ is the input spike rate flux, with the units spikes.s⁻¹. The ψ_{jk} are dimensionless and they weight the effectiveness of the input $\rho_j \Phi_{jk}$. The equation governing this weighting is:

$$\psi_{jk} = \frac{V_j^{\text{rev}} - V_k}{V_j^{\text{rev}} - V_k^{\text{rest}}}. \quad (2.5)$$

To explain the weighting of the inputs, consider this example of an excitatory neuron giving input $\rho_e \Phi_{ee}(t)$ to another excitatory neuron. If the potential of the postsynaptic neuron V_e is at its resting potential V_e^{rest} , we have $(V_e^{\text{rev}} - V_e^{\text{rest}})/(V_e^{\text{rev}} - V_e^{\text{rest}}) = 1$ and the full voltage input is received. In another scenario, if V_e were already almost as high as V_e^{rev} , the numerator of $\psi_{ee}(t)$ would be close to zero, and V_e would increase by only a small fraction of the input. The ψ_{jk} weights constrain V_e and V_i to lie between the reversal potentials, V_e^{rev} and V_i^{rev} . Figure 2.3 shows how the weights $\psi_{ek}(t)$ and $\psi_{ik}(t)$ change with varying soma voltage V_k using the parameters in the current model. (A full list of typical parameters is given in table 2.2.) The soma potential of the postsynaptic neuron is *increased* with excitatory inputs, and *decreased* with inhibitory inputs.




Figure 2.3: The value of ψ_{jk} weights the effectiveness of the voltage inputs $\rho_j \Phi_{jk}$ on the postsynaptic neuron. In the current model (Steyn-Ross *et al.*, 2005), the value of $V_e^{\text{rest}} = V_i^{\text{rest}} \approx -64$ mV, $V_e^{\text{rev}} = 0$ mV and $V_i^{\text{rev}} = -70$ mV. [image source: Steyn-Ross *et al.* (2005)]

The total input spike rate flux Φ is an integration of all local (β : within the same macrocolumn), distant (α : from other macrocolumns) and subcortical (sc: from structures such as the thalamus and brainstem) voltage spikes. These are convolved with the dendrite impulse response curve H . The total input spike rate is given by the per-synapse spike rate ϕ_{ek} , ϕ_{ik} multiplied by the number of $e \rightarrow k, i \rightarrow k$

Table 2.2: Typical model constants for the neural macrocolumn. Subscript label k means destination cell can be either of type e (excitatory) or i (inhibitory). Note that there is considerable uncertainty in many of these values.

Symbol	Description	Value	Unit
$\tau_{e,i}$	membrane time constant	0.050, 0.050	s
$V_{e,i}^{\text{rev}}$	cell reversal potential	0, -70	mV
$V_{e,i}^{\text{rest}}$	cell resting potential	-64, -64	mV
$\rho_{e,i}$	synaptic gain	$(1.00, -1.05) \times 10^{-3}$	mV·s
$N_{ee,ei}^{\alpha}$	long-range $e \rightarrow k$ synaptic connections	3710, 3710	—
$N_{ee,ei}^{\beta}$	local $e \rightarrow k$ synaptic connections	410, 410	—
$N_{ie,ii}^{\beta}$	local $i \rightarrow k$ synaptic connections	800, 800	—
$N_{ee,ei}^{\text{sc}}$	subcortical $e \rightarrow k$ synaptic connections	50, 50	—
$N_{ie,ii}^{\text{sc}}$	subcortical $i \rightarrow k$ synaptic connections	50, 50	—
v	axonal conduction speed	9	m s^{-1}
$\Lambda_{ee,ei}$	axonal inverse-length scale	40, 40	m^{-1}
$\gamma_{ee,ei}$	rate constant for $e \rightarrow k$ synaptic input	70, 70	s^{-1}
$\gamma_{ie,ii}$	rate constant for $i \rightarrow k$ synaptic input	58.6, 58.6	s^{-1}
$Q_{e,i}^{\text{max}}$	maximum firing rate for sigmoid	30, 60	s^{-1}
$\theta_{e,i}$	inflexion-point voltage for sigmoid	-58.5, -58.5	mV
$\sigma_{e,i}$	sigmoid width	4, 6	mV

connections N_{ek}, N_{ik} .

$$\Phi_{ek}(t) = \int_{-\infty}^t H_{ek}(t-t') \left[\underbrace{N_{ek}^{\beta} \phi_{ek}^{\beta}(t')}_{\text{local}} + \underbrace{\phi_{ek}^{\text{sc}}(t')}_{\text{subcortical}} + \underbrace{N_{ek}^{\alpha} \phi_{ek}^{\alpha}(t')}_{\text{distant}} \right] dt' \quad (2.6)$$

$$\Phi_{ik}(t) = \int_{-\infty}^t H_{ik}(t-t') \left[N_{ik}^{\beta} \phi_{ik}^{\beta}(t') + \phi_{ik}^{\text{sc}}(t') \right] dt' \quad (2.7)$$

The ϕ_{ek}^{β} and ϕ_{ik}^{β} are the spike rates from other neurons within the macrocolumn, ϕ_{ek}^{sc} and ϕ_{ik}^{sc} are spike rates from subcortical structures, and ϕ_{ek}^{α} is the spike rate from neurons in other macrocolumns. Notice that there is no $N_{ik}^{\alpha} \phi_{ik}^{\alpha}$ in Eq. (2.7). This is because inhibitory neurons do not have long-range output connections (i.e., to other macrocolumns). The ϕ_{jk}^{α} and ϕ_{jk}^{β} are dependent on the the soma voltages V_e and V_i . They are explained in the next section.

The H_{jk} term is the dendrite impulse response curve. It models how the ion channels open and let the excitatory or inhibitory voltage pulses enter the cell. H_{jk}

is defined as:

$$H_{jk}(t) = t \gamma_{jk}^2 \exp[-t \gamma_{jk}], \quad t \geq 0 \quad \text{dendrite response} \quad (2.8)$$

The soma impulse response function, L , and the dendrite impulse response function, H , are both zero when $t < 0$. The functions are also normalised so that their area is exactly 1. They are plotted in Fig. 2.4. Notice that H decays much faster than L . This becomes important in the “slow membrane” approximation made in the Steyn-Ross *et al.* (2005) model.




Figure 2.4: The impulse response of the soma and dendrites. The potential of the dendrite response equalises much faster than the soma after an input spike. [image source: Steyn-Ross *et al.* (2005)]

2.3.2 Converting the Integrals to Derivatives.

As integrals the equations are easier to understand, but they are difficult to use in analysis. With the use of the calculus theorem for differentiation of integrals with variable limits (see page 83 of Spiegel (1963)), Eq. (2.3) is converted in to a pair of first-order differential equations (one each for $k = e, i$).

$$\tau_e \frac{dV_e}{dt} = V_e^{\text{rest}} - V_e + \rho_e \psi_{ee}(t) \Phi_{ee}(t) + \rho_i \psi_{ie}(t) \Phi_{ie}(t) \quad (2.9)$$

$$\tau_i \frac{dV_i}{dt} = V_i^{\text{rest}} - V_i + \rho_e \psi_{ei}(t) \Phi_{ei}(t) + \rho_i \psi_{ii}(t) \Phi_{ii}(t) \quad (2.10)$$

If one ignores all ρ 's, ψ 's and Φ 's from Eq. (2.9), it should be recognisable as an exponential decay of V_e from some initial level to its equilibrium position V_e^{rest} , with τ_e as the time constant. The situation is exactly the same for the inhibitory neuron in Eq. (2.10).

By using the same calculus theorem again, *twice*, Eqs. (2.6, 2.7) are converted into four second-order differential equations. The first pair of equations describes the excitatory postsynaptic neurons,

$$\left(\frac{d}{dt} + \gamma_{ee} \right)^2 \Phi_{ee} = [N_{ee}^\alpha \phi_{ee}^\alpha(t) + N_{ee}^\beta Q_e(t) + \phi_{ee}^{\text{sc}}(t)] \gamma_{ee}^2 \quad (2.11)$$

$$\left(\frac{d}{dt} + \gamma_{ie} \right)^2 \Phi_{ie} = [N_{ie}^\beta Q_i(t) + \phi_{ie}^{\text{sc}}(t)] \gamma_{ie}^2 \quad (2.12)$$

and the second pair describes the inhibitory neurons,

$$\left(\frac{d}{dt} + \gamma_{ii} \right)^2 \Phi_{ii} = [N_{ii}^\beta Q_i(t) + \phi_{ii}^{\text{sc}}(t)] \gamma_{ii}^2 \quad (2.13)$$

$$\left(\frac{d}{dt} + \gamma_{ei} \right)^2 \Phi_{ei} = [N_{ei}^\alpha \phi_{ei}^\alpha(t) + N_{ei}^\beta Q_e(t) + \phi_{ei}^{\text{sc}}(t)] \gamma_{ei}^2 \quad (2.14)$$

In these last four equations, new variables $Q_{e,i}$ have appeared and replaced ϕ_{jk}^β . These were the internal macrocolumn spike rates, which are assumed to be equal to the mean rate of spiking of the macrocolumn. The new variable Q is an average spike rate that is determined by the average soma voltage; the higher the soma voltage, the higher the firing rate. All the neurons have a threshold soma voltage that must be exceeded before they will fire an action potential. Among the population of neurons, the thresholds are assumed to have a Gaussian distribution with mean $\theta_{e,i}$ and standard deviation $\sigma_{e,i}$. $Q_{e,i}$ is then the cumulative sum of the distribution and is described by a sigmoid function.

$$Q_{e,i} = \frac{Q_{e,i}^{\text{max}}}{1 + \exp \left[\frac{-\pi (V_{e,i}(t) - \theta_{e,i})}{\sqrt{3} \sigma_{e,i}} \right]} \quad (2.15)$$

The last two of the set of eight differential equations describe the long-range excitatory inputs from distant macrocolumns. They are modelled as damped waves




Figure 2.5: The sigmoid functions that give the firing rates Q_e and Q_i plotted against the soma voltage. The firing rates $Q_{e,i}$ have been normalised by their maximums firing rates: $Q_e^{\max} = 30 \text{ s}^{-1}$, $Q_i^{\max} = 60 \text{ s}^{-1}$.

that are generated by other sources $Q_e(\vec{r}, t)$, propagating along the 2-D cortex.

$$\left[\left(\frac{\partial}{\partial t} + v\Lambda_{ee} \right)^2 - v^2 \nabla^2 \right] \phi_{ee}^\alpha(\vec{r}, t) = (v\Lambda_{ee})^2 Q_e(\vec{r}, t) \quad (2.16)$$

$$\left[\left(\frac{\partial}{\partial t} + v\Lambda_{ei} \right)^2 - v^2 \nabla^2 \right] \phi_{ei}^\alpha(\vec{r}, t) = (v\Lambda_{ei})^2 Q_e(\vec{r}, t) \quad (2.17)$$

The v in the above equations is the average speed that an action potential travels down the axon. The other new parameter Λ_{ek} is the inverse-length scale for long-range connections. It governs the average distance that action potentials travel (analogous to the time-constants for the decays of the impulse responses earlier).

Equations (2.9–2.17) make up the (slightly modified) Liley *et al.* (1999) model as used by Steyn-Ross *et al.* (2005).

To keep their model of the cortex consistent with the real cortex Steyn-Ross *et al.* (2005) feed the system with white noise $\xi_m(t)$ through the subcortical terms ϕ_{jk}^{sc} seen in Eqs. (2.11–2.14).

$$\phi_{jk}^{\text{sc}}(t) = s_{jk} Q_j^{\max} + \sqrt{s_{jk} Q_j^{\max}} \cdot \xi_m(t). \quad (2.18)$$

Here m can take the values $m = 1$ for $j = e, k = e$, $m = 2$ for $j = i, k = e$, $m = 3$

for $j = e, k = i$ and $m = 4$ for $j = i, k = i$. Each ξ_m is completely independent of all the others, and they all have a mean of zero: $\langle \xi_m(t) \rangle = 0$. The other parameters s_{jk} are just scalars to allow variability to the level of the noise.

2.3.3 Approximations and Stability of the Steady States

The Steyn-Ross *et al.* (2005) model makes an approximation that was made in previous work (Steyn-Ross *et al.*, 2001a,b, 2003, 2004), namely, the “slow neuron” adiabatic limit. The assumption is that the relaxation time of the impulse response of the dendrites is much quicker than the relaxation time impulse response of the soma, i.e., $\gamma_{jk}^{-1} \ll \tau_k$. To implement their approximations the temporal derivatives d/dt in Eqs. (2.11–2.14) and Eqs. (2.16–2.17) are set to 0. The cortex is also assumed to be homogeneous, meaning that all macrocolumns are identical so only one macrocolumn needs to be modelled. In effect this means that the Laplacian ∇^2 in Eqs. (2.16–2.17) can be set to zero. The Eqs. (2.9–2.17) now reduce to only two equations:

$$\begin{aligned} \frac{dV_e}{dt} = & \frac{1}{\tau_e} \left\{ V_e^{\text{rest}} - V_e(t) + \rho_e \psi_{ee}(t) \left[(N_{ee}^\alpha + N_{ee}^\beta) Q_e + s_{ee} Q_e^{\text{max}} \right] \right. \\ & \left. + \rho_i \psi_{ie}(t) \left[N_{ie}^\beta Q_i + s_{ie} Q_i^{\text{max}} \right] + \Gamma_e(t) \right\} \end{aligned} \quad (2.19)$$

$$\begin{aligned} \frac{dV_i}{dt} = & \frac{1}{\tau_i} \left\{ V_i^{\text{rest}} - V_i(t) + \rho_e \psi_{ei}(t) \left[(N_{ei}^\alpha + N_{ei}^\beta) Q_e + s_{ei} Q_e^{\text{max}} \right] \right. \\ & \left. + \rho_i \psi_{ii}(t) \left[N_{ii}^\beta Q_i + s_{ii} Q_i^{\text{max}} \right] + \Gamma_i(t) \right\} \end{aligned} \quad (2.20)$$

The white-noise (stochastic) terms are written as $\Gamma_{e,i}$:

$$\Gamma_e(t) = \rho_e \psi_{ee}(t) \sqrt{s_{ee} Q_e^{\text{max}}} \cdot \xi_1(t) + \rho_i \psi_{ie}(t) \sqrt{s_{ie} Q_i^{\text{max}}} \cdot \xi_2(t) \quad (2.21)$$

$$\Gamma_i(t) = \rho_e \psi_{ei}(t) \sqrt{s_{ei} Q_e^{\text{max}}} \cdot \xi_3(t) + \rho_i \psi_{ii}(t) \sqrt{s_{ii} Q_i^{\text{max}}} \cdot \xi_4(t) \quad (2.22)$$

Another member of the cortical modelling group at Waikato, Marcus Wilson, has run numerical simulations on a two-dimensional grid of macrocolumns without using the adiabatic approximation. His results on the dynamics of the system are reported in Wilson *et al.* (2005).

By also setting to zero the white noise terms $\Gamma_{e,i}$, the steady states of V_e and V_i can be found — i.e., solutions where $V_{e,i}(\vec{r}, t)$ is constant in space and time. In most cases there is only one solution, but for a particular parameter set more than solution may exist. The region in the parameter space where solutions cease to

exist is usually of most interest. If there are multiple solutions the the stability of the steady states must be checked. Some states can be stable and others can be unstable.

A stable steady state is analogous to a ball sitting in the valley of two hills (two dimensions, height and width). If the ball is pushed part way up the hill, it will roll back down. It may oscillate back and fourth, but it will eventually settle back to its equilibrium position. An unstable state, conversely, is analogous to a ball sitting on the top of a hill. The ball is not moving, so the top of the hill is steady. It is not stable though, because if the ball is slightly pushed one way or another, it will roll down and not return to the top.

To analyse the stability of the steady states and to consider the size of fluctuations about these states Steyn-Ross (2002) first wrote the equations for soma voltages as the sum of deterministic and stochastic parts:

$$\frac{dV_e}{dt} = F_1 + \Gamma_e(t)/\tau_e \quad (2.23)$$

$$\frac{dV_i}{dt} = F_2 + \Gamma_i(t)/\tau_i \quad (2.24)$$

where $F_{1,2}$ are all the terms on the right hand side of Eq. (2.19, 2.20) except $\Gamma_{1,2}$. The stability is determined by perturbing the soma voltage V_k away from its steady state V_k^0 by a small amount δV_k .

$$V_e = V_e^0 + \delta V_e \quad (2.25)$$

$$V_i = V_i^0 + \delta V_i \quad (2.26)$$

The linear approximations to the deterministic part of Eqs. (2.23, 2.24) are combined in the matrix equation,

$$\frac{d}{dt} \begin{bmatrix} \delta V_e \\ \delta V_i \end{bmatrix} = \mathbf{A} \begin{bmatrix} \delta V_e \\ \delta V_i \end{bmatrix} \quad (2.27)$$

where \mathbf{A} is the 2×2 Jacobian matrix. \mathbf{A} needs to be computed at every steady state position (V_e^0, V_i^0) .

$$\mathbf{A} \equiv \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \equiv \begin{bmatrix} \frac{\partial F_1}{\partial V_e} & \frac{\partial F_1}{\partial V_i} \\ \frac{\partial F_2}{\partial V_e} & \frac{\partial F_2}{\partial V_i} \end{bmatrix} \quad (2.28)$$

The eigenvalues $\lambda_{1,2}$ of \mathbf{A} are then calculated via the formula

$$\lambda_{1,2} = \frac{T \pm \sqrt{T^2 - 4D}}{2} \quad (2.29)$$

where T is the trace and D the determinant of \mathbf{A} . The solutions of the differential equations (2.19, 2.20) are then written:

$$\delta V_e = a_1 e^{\lambda_{1t}} + a_2 e^{\lambda_{2t}} \quad (2.30)$$

$$\delta V_i = b_1 e^{\lambda_{1t}} + b_2 e^{\lambda_{2t}} \quad (2.31)$$

Here $a_{1,2}$ and $b_{1,2}$ are constants that depend on the initial conditions. Both eigenvalues must be negative for the steady state to be stable. If one of the eigenvalues is not negative, the fluctuations will grow exponentially. The eigenvalues can be complex, which will lead to oscillatory components, but their real part must still be negative.

2.4 Expanding from Anaesthesia to Natural Sleep

2.4.1 The Anaesthesia Model

In earlier work (Steyn-Ross *et al.*, 1999, 2001a,b, 2003, 2004) the model was used to predict changes in the cortex when it was under the effect of an anaesthetic such as propofol. It was easier to model anaesthesia than natural sleep, because there is only one major control variable: the concentration of the anaesthetic drug.

The drug propofol works by enhancing the inhibitory effect of the neurotransmitter GABA (gamma-amino butyric acid) by holding the chloride ion channel open longer and allowing more Cl^- to enter the postsynaptic neuron. This causes the soma voltage to be lowered further from its firing threshold. To incorporate the effect of the anaesthetic into their model, the parameters γ_{ie} and γ_{ii} were divided by a dimensionless factor λ .

$$\gamma_{ik} \rightarrow \frac{\gamma_{ik}}{\lambda} \quad (2.32)$$

In Fig. 2.6 the thick curve in (a) marks the steady states of Eqs. (2.9–2.17) as a function of anaesthetic effect λ . There are three branches of the “S-bend” corresponding to three solutions for a particular range of λ . The upper and lower branches are stable, but the middle branch is unstable. As the concentration of anaesthetic concentration increases and the average soma voltage decreases along the steady state curve, the amplitudes of its fluctuations get larger. When the corner

of the “S-bend” is reached (marked A_3), the soma voltage must suddenly drop from the upper branch to the lower branch. Following this, the fluctuation amplitude dramatically decreases. This behaviour for the fluctuations is a hallmark of a first-order phase transition.

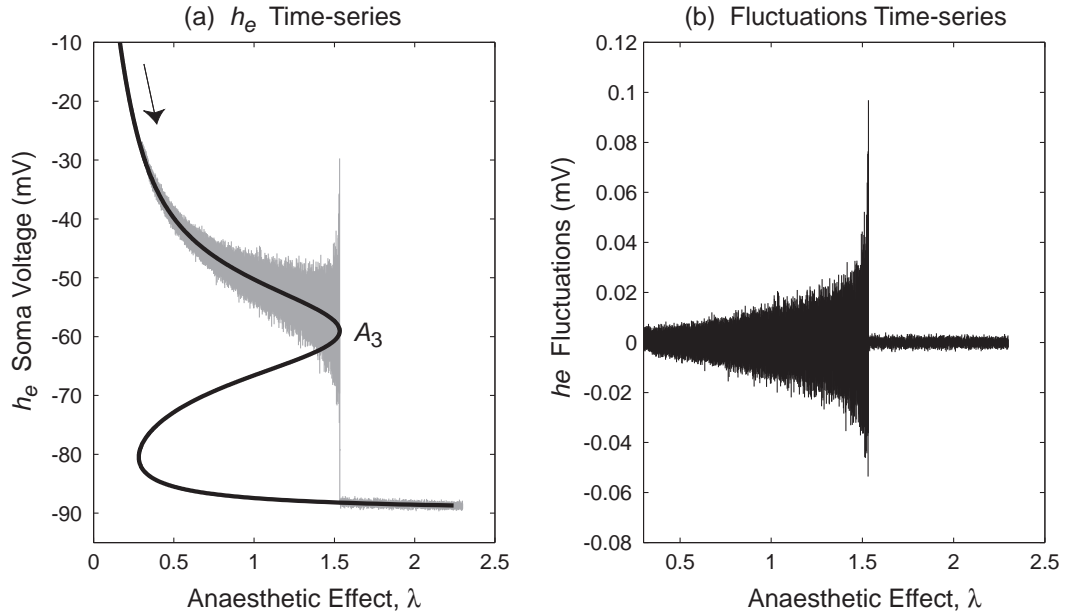


Figure 2.6: An EEG time-series simulation of an anaesthetised cortex from an earlier model (Steyn-Ross *et al.*, 2001a). In their older notation h_e was the symbol assigned to the excitatory soma voltage (now V_e). The change from one branch to another makes this a first-order phase transition similar to those commonly seen in thermodynamics. In (a) the amplitude of the fluctuations was scaled up by 300 to make them visible; they are at their true size in (b). [image source: Steyn-Ross *et al.* (2001a)]

The large increase in soma voltage fluctuations as a patient approaches anaesthesia is well known by anaesthetists and is referred to as a biphasic response. This experimental behaviour is consistent with the predictions of the model.

2.4.2 Modelling Natural Sleep

A figure in a paper by Destexhe *et al.* (1999) (reprinted here in Fig. 2.7) is reminiscent of the results of the anaesthetic modelling. This prompted the question as to whether phase transitions also occur in natural sleep — specifically the transition from SWS to REM sleep.

Sleep becomes more complex to model than anaesthesia because there are many neuromodulators that are involved in the sleep cycle. To try to keep the model as

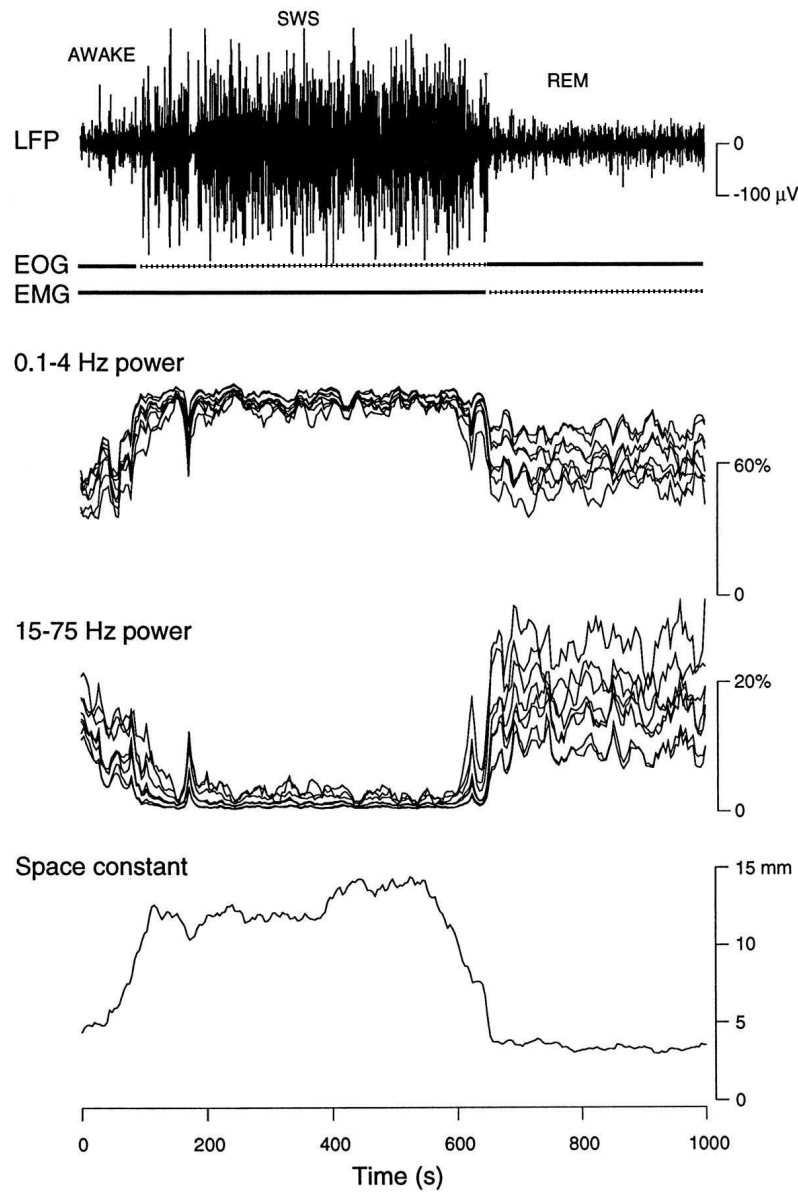


Figure 2.7: The EEG, fractional band power, and space constant of a sleeping cat studied by Destexhe *et al.* (1999). The sharp drop in the EEG fluctuations is similar to what was observed in Fig. 2.6

simple as possible Steyn-Ross *et al.* (2005) model the effect of only two neuromodulators: *adenosine* and *acetylcholine*

Adenosine seems to control tiredness, or fatigue. Its concentration steadily increases while one is awake, and then decreases again during sleep (España and Scammell, 2004). Adenosine works by increasing K^+ leak current. This lowers the resting voltage V_e^{rest} and makes the cells less likely to fire (Hasselmo, 1995).

The second neuromodulator, acetylcholine (ACh), has a high concentration during REM sleep, but during non-REM sleep, the concentration drops to almost zero (España and Scammell, 2004). There are two major effects of ACh. The first is opposite to that of adenosine. ACh reduces the K^+ leak current, which raises V_e^{rest} and makes the cells more likely to fire. The second effect of ACh is somewhat contrary to the first. While it increases V_e^{rest} , it also reduces the amplitude of the excitatory PSP (Hasselmo, 1995). This corresponds to a decrease in the excitatory gain ρ_e .

The two neuromodulator effects are included into the differential equations by changing the parameters V_e^{rest} and ρ_e from Eqs (2.9) and (2.10):

$$V_e^{\text{rest}} \rightarrow V_e^{\text{rest}} + \Delta V_e^{\text{rest}} \quad (2.33)$$

$$\rho_e \rightarrow \lambda_{\text{ACh}} \rho_e \quad (2.34)$$

where λ_{ACh} is a dimensionless scalar to model the concentration of ACh and ΔV_e^{rest} is an additive constant. A high ACh concentration corresponds to a low value of λ_{ACh} . Steyn-Ross *et al.* (2005) have chosen to not include the $+\Delta V_e^{\text{rest}}$ into the ψ_{jk} weights of Eq. (2.5) to avoid unnecessary complexity. The soma voltage equations now become:

$$\begin{aligned} \frac{dV_e}{dt} = \frac{1}{\tau_e} \Big\{ & V_e^{\text{rest}} + \Delta V_e^{\text{rest}} - V_e(t) + \lambda_{\text{ACh}} \rho_e \psi_{ee}(t) \left[(N_{ee}^\alpha + N_{ee}^\beta) Q_e + s_{ee} Q_e^{\text{max}} \right] \\ & + \rho_i \psi_{ie}(t) \left[N_{ie}^\beta Q_i + s_{ie} Q_i^{\text{max}} \right] + \Gamma_e(t) \Big\} \end{aligned} \quad (2.35)$$

$$\begin{aligned} \frac{dV_i}{dt} = \frac{1}{\tau_i} \Big\{ & V_i^{\text{rest}} - V_i(t) + \lambda_{\text{ACh}} \rho_e \psi_{ei}(t) \left[(N_{ei}^\alpha + N_{ei}^\beta) Q_e + s_{ei} Q_e^{\text{max}} \right] \\ & + \rho_i \psi_{ii}(t) \left[N_{ii}^\beta Q_i + s_{ii} Q_i^{\text{max}} \right] + \Gamma_i(t) \Big\} \end{aligned} \quad (2.36)$$

The steady states of Eqs. (2.35) and (2.36) were found on the domain defined by the new parameters ΔV_e^{rest} and λ_{ACh} . The steady states for V_e in the model also contains an ‘‘S-bend’’, but now there is a 3-D manifold (shown in Fig. 2.8). The amount of overlap in the three steady states reduces as ΔV_e^{rest} rises until there is only one steady state.




Figure 2.8: The steady states of V_e of the macrocolumn for the Steyn-Ross *et al.* (2005) model of natural sleep are plotted as a function of ΔV_e^{rest} and λ_{ACh} . For most of the surface there is only one steady state, but on the left side of the graph, inside the green \supset shape, the upper and lower surfaces overlap and three steady states exist. The thick black line traces out an elliptical tour. This will be explained later in section 2.4.3. [image source: Steyn-Ross *et al.* (2005)]

2.4.3 Predictions of the Steyn-Ross Model

The Steyn-Ross *et al.* (2005) model makes predictions of how the EEG will change as a sleeper progresses through his sleep cycle. The ellipse drawn on to the manifold of steady states marks out a hypothetical sleep cycle on the ΔV_e^{rest} , λ_{ACh} domain. A bird's-eye view of the manifold is shown in Fig. 2.9. The starting position, marked as (+), is on the lower branch on what is presumed to be SWS. The trajectory of the ellipse takes it towards the edge of the cusp (marked by the green curve). From here there is nowhere else to go except to the upper branch which is assumed to be REM sleep. This jump, occurring at the position marked (o), is a first-order phase transition. The ellipse trajectory continues around the manifold where only one steady state exists. It eventually passes under the upper branch again and returns to the start. This cycle was chosen to agree with the evidence that the SWS to REM transition is a sudden change, but the transition back from REM to SWS is not (see chapter 5).

Steyn-Ross *et al.* (2005) also calculated the theoretical changes in statistics that can be measured from real EEG. Figure 2.10 shows the predicted changes in EEG time-series statistics for one complete sleep cycle that follows the elliptical tour.




Figure 2.9: A bird's-eye view of the sleep manifold from Fig. 2.8. The elliptical tour starts at (+) and undergoes a first-order phase transition as it passes from SWS to REM sleep at (o). [image source: Steyn-Ross *et al.* (2005)]

The model predicts that there will be a large surge in both the total power and correlation time of the fluctuations of the EEG at the SWS to REM transition (These quantities are discussed in more detail in chapter 4). This is followed by a sharp drop to a lower level. The power fractions show how the energy is spectrally distributed. Before the transition more than 50% of the energy will be contained below 4 Hz, with all power concentrated at zero frequency the instant before transition. After REM begins, there is much more energy located in the higher frequency band. The changes predicted by the model are qualitatively consistent with experimental results presented by Destexhe *et al.* (1999) for their sleeping cat (Fig. 2.7).

The results of the time-series statistics computed for real EEG are presented in chapter 5 for human sleep patients, and also chapter 6 for laboratory rats.

In a three second, numerical simulation, Steyn-Ross *et al.* (2005) showed how V_e changes as the macrocolumn makes the SWS to REM transition. In Fig. 2.11 the macrocolumn starts on the bottom branch of Fig. 2.8 very close to the cusp and the white noise causes V_e to fluctuate about $V_e^{0,Bot} \approx -64.4$ mV. At about 2 s into the simulation V_e becomes greater than the unstable middle branch and rises up to $V_e^{0,Top} \approx -53.9$ mV on the top branch. Here the amplitude of the fluctuations has greatly diminished and only the high-frequency activity remains.

When the macrocolumn transitions from SWS to REM, the average excitatory

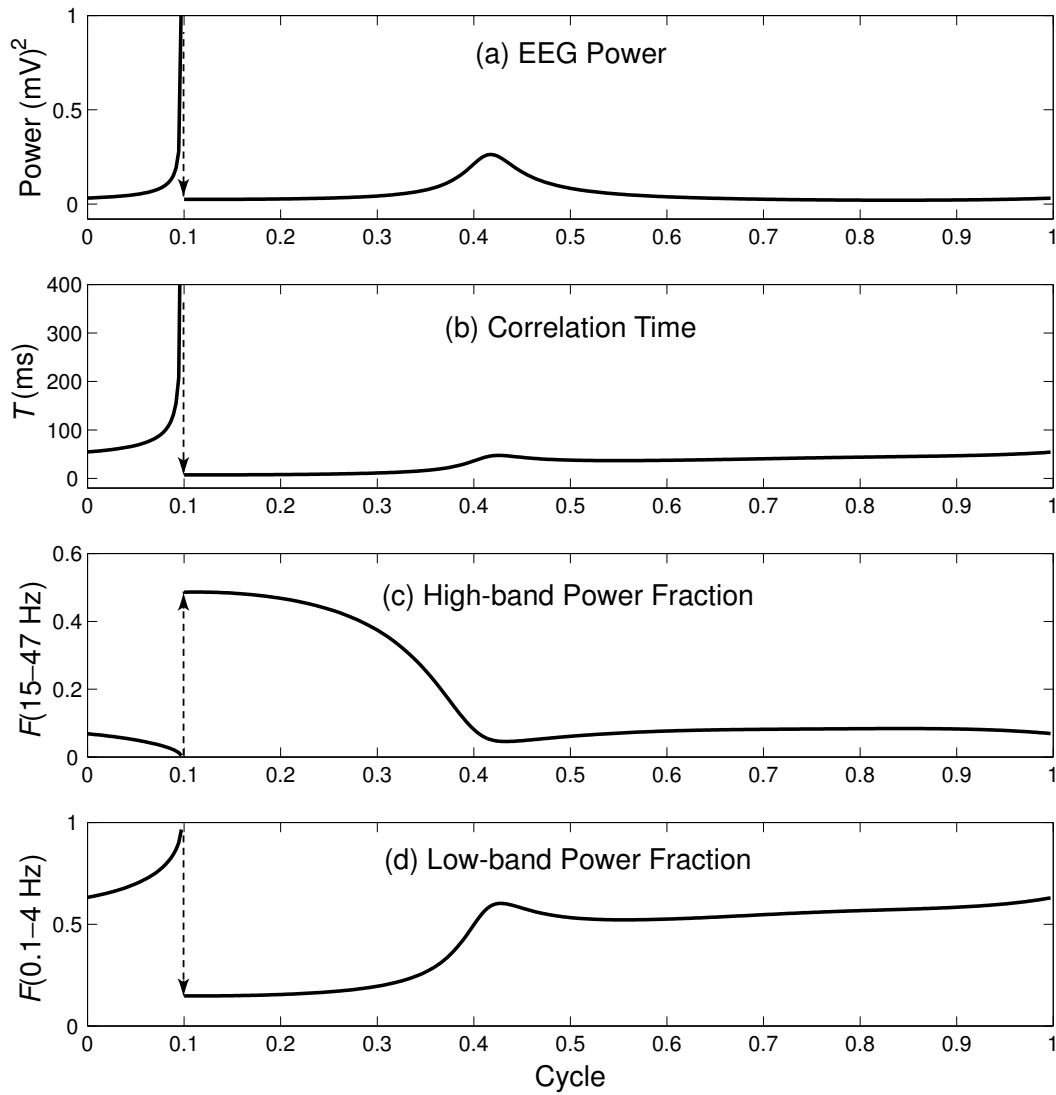


Figure 2.10: Time-series statistics that were predicted by taking the elliptical tour of the sleep cycle shown in Figs. 2.8 and 2.9. These statistics are calculated on real EEG in chapters 5 and 6. [image source: Steyn-Ross *et al.* (2005)]

soma voltage V_e increases by ~ 10.5 mV. This change in the DC voltage should be possible to detect if one can accurately record EEG down to zero frequency. The EEG of laboratory rats were recorded using a DC-coupled amplifier so that the predicted DC-shift could be observed. The results of the DC EEG are presented with the other rat results in chapter 6.




Figure 2.11: A numerical simulation of the nonlinear differential equations (2.19, 2.20) describing a macrocolumn as it transitions from SWS to REM. The starting position was an ellipse angle of $\theta = 34^\circ$ (with the $\theta_+ = 0^\circ$ and $\theta_o \approx 35.08^\circ$ in Fig. 2.9). (a) At $t \approx 2$ s the fluctuations become large enough to cross the unstable mid branch and the macrocolumn jumps up to the higher voltage level. (b) The fluctuations δV_e as what might be seen in an AC-coupled chart recorder. Before transition: $\delta V_e = V_e - V_e^{0, \text{Bot}}$. After transition: $\delta V_e = V_e^{0, \text{Top}} - V_e$. [image source: Steyn-Ross *et al.* (2005)]

Chapter 3

EEG Recording and Filtering

In this chapter the details of recording the raw EEG for a sleep study are described. The EEG electrode positions and the positions of electrodes for measuring other useful bio-potentials are given. The raw recordings contained 50 Hz interference from the mains power supply. The filter that was designed to remove the interference is also explained.

3.1 The Raw Sleep Recordings

3.1.1 Electrode Configuration

When recording the EEG from the scalp of a human subject, there is a standard set of positions for placing the electrodes. It is called the International 10-20 System of Electrode Placements. The 10-20 system labels certain areas of the human scalp with a letter that corresponds to the cortical lobe that is located in that area. Table 3.1 is the list of cortical lobes from which EEG signals are recorded.

3.1.2 Sleep Stages

For sleep studies, Rechtschaffen and Kale (1968) recommend that the recorded EEG should be the difference of the EEG signals from the central lobe and the mastoid on the opposite side (e.g., $EEG_1 = C_3 - A_2$ and $EEG_2 = C_4 - A_1$). The EEG record is usually displayed in time epochs of about 30 seconds. By observing the characteristics of the EEG signal (looking for certain amplitude and frequency waveforms) the stage of sleep can be scored. Some important EEG waveforms that are used as sleep staging markers are:

alpha waves 8–12 Hz oscillations usually seen when one is relaxing with their eyes

Table 3.1: The meanings of the labels used in the International 10-20 System of Electrode Placement. Odd numbers are on the left hemisphere while even numbers are on the right. A subscript “Z” indicates the line down the middle of the head, between the two hemispheres of the cortex.

Label	Cerebral Cortex Lobe
F _p	Front polar lobe
F	Frontal lobe
C	Central lobe
T	Temporal lobe
P	Parietal lobe
O	Occipital lobe
A	Ear (mastoid)




Figure 3.1: The electrode positions used for recording EEG as specified by the International 10-20 System of Electrode Placement. [image source: <http://butler.cc.tut.fi/~malmivuo/bem/bembook/13/13.htm>]

closed, but still awake.

spindles Coherent bursts of 12–15 Hz oscillations that rise and fall in less than one second. Commonly seen in stage 2. They are simultaneously detected over the cortex, and are believed to come from the thalamus.

K-complex A quick spike that over shoots on its recovery then settles back to normal EEG rhythms. The whole waveform lasts about 0.5 s. Most common to stage 2, and have been found to be evoked if the sleeper hears a startling noise.

delta waves large amplitude 1–4 Hz oscillations. Stage 3 and 4 are defined by the amount of delta waves.

The EEG is of primary importance, but the sleep stages can be difficult to discern with the EEG alone. Voltage signals are also recorded from other locations to assist with the sleep staging.

The very title of REM sleep gives away one of the staging signals. The eyes of the sleeper can move very fast during REM sleep, making the electro-oculogram (EOG) a good distinctive feature. REMs are much faster than normal movements, but they are not unnatural movements where, for example, one eye looks up while the other looks down. The eyeball has an electric dipole with the positive end at the cornea; the electrodes measure the change in electric field due to the moving dipole. Usually the EOG is recorded for both eyes. If one electrode is positioned upper-left of the left eye and the other is lower-right of the right eye, the two EOG channels will change in anti-phase with normal eye movements (both eyes looking in the same direction). Eye-blinks can distinguished from REMs because the two EOG signals change in-phase.

Another useful bio-signal is the muscle tone or electromyogram (EMG). During REM sleep the muscles become more relaxed so the EMG voltage becomes lower. It is also helpful for scoring wakefulness and movement. In sleep studies, the EMG is often recorded from under the chin.

The rules on categorising the stages of sleep as written by Rechtschaffen and Kale (1968) are summarised below.

wake Not a stage of sleep, but must be recorded when the sleeper is not sleeping. It can be identified by large EMG indicating movements, or more than 50% of the epoch contains alpha activity in the EEG.

stage 1 This is a transitionary phase lasting several minutes and containing many




Figure 3.2: Electrode placements for recording EEG, EMG and EOG during sleep studies. [image source: Rechtschaffen and Kale (1968)]

small arousals. The EEG contains low amplitude theta (4–8 Hz) and possibly delta, less than 50% alpha, and no spindles or K-complexes.

stage 2 Usually follows stage 1. The EEG contains mostly theta and some delta (less than 20% of the epoch). Spindles and K-complexes are seen. Scoring as stage 2 continues until there is good reason to change to another stage.

stage 3 The EEG has more than 20% delta, but less than 50%. The amplitudes are much larger than what has been seen in previous stages, greater than $75 \mu\text{V}$. Spindles and K-complexes are less common. There is no specific EMG level for stage 3 but it is typically lower than it was for stage 2.

stage 4 The EEG has more than 50% delta with high-amplitude. Stage 3 and stage 4 are often considered together as the transition is gradual and can the stage can alternate.

REM The EEG is low-amplitude and high-frequency, resembling wake or stage 1. Very sharp eye movements begin seen in the EOG and the EMG should be minimal, but may not be much lower than what was seen in stage 3 or 4. If a spindle or K-complex is seen during REM, then the epoch should be reclassified as stage 2.

In the figures below are samples of EEG from the various stages of sleep. Notice that from stages 1–4 there is an increase of low-frequency, high-amplitude activity. The sample of REM EEG is also platted with the left and right eye EOG. The very

sharp movements are also called saccades, and can have a peak angular velocity of greater than 500° s^{-1} (Steriade and McCarley, 2005).




Figure 3.3: Sample EEG of the different stages of sleep, 1–4 and REM. The EOG is included with REM to show the quick movements.

3.2 Human Sleep Recordings

The human EEG recordings that are analysed in chapter 5 came from patients who were all diagnosed with sleep apnea, a disorder in which the sufferer repeatedly stops breathing and wakes up during the night. It disturbs people's sleeping habits so that they always feel tired and, in extreme cases, they can fall asleep during everyday activities. The patients whose sleep records were supplied to me were being treated using Fisher & Paykel Healthcare's CPAP (continuous positive airway pressure) respirators. The original purpose of recording the patients' EEG was to determine how the CPAP respirators were helping their sleeping patterns. They had all agreed to the release of their sleep records for research purposes.

The patients' sleep records were taken at two locations, Fisher & Paykel Health-

care and Greenlane Hospital. The two locations had different recording equipment. Fisher & Paykel Healthcare had a reasonably new plant with up-to-date computer hardware. The sleep recordings made there used sampling rates of 256 Hz for the EEG. The EEG from the four scalp locations mentioned in the previous section, C_3 , C_4 , A_1 and A_2 , were all recorded independently (referenced to F_{pz}) and the $EEG_1 = C_3 - A_2$, $EEG_2 = C_4 - A_1$ were performed by their analysis software.

At Greenlane, the equipment was much older. In an effort to save electronic storage space, they made the EEG_1 and EEG_2 subtractions during the recording, presumably with a differential amplifier. The Greenlane data also used lower sampling rates. The EEG was sampled at only 125 Hz. They also had a low clipping level¹ which cut some of the higher peaks short.

3.3 The Binary Data Format

The commercial software that was used for acquisition and analysis of the sleep studies at Fisher & Paykel Healthcare and at Greenlane hospital, were not available at Waikato, so the data was exported in the European data format (EDF). The EDF format is a relatively straight forward binary format capable of storing many channels of data. The channels can even have different sample rates.

The binary data files for the human patients were opened and read into MATLAB with the use the BIOSIG toolbox. It is a set of MATLAB-compatible functions to aid in analysis of EEG and other bio-signals written by Alois Schloegl.²

3.4 Filtering Electrical Interference

3.4.1 Filtering Human EEG

The unfiltered data contained a lot of 50 Hz electrical interference from the mains power supplies. An elliptical low-pass filter was designed to remove the 50 Hz. Elliptical filters can provide large attenuations over a very small frequency range with only a small filter order. The cost of using such a filter is that there is ripple attenuation in the pass-band. The filter specifications used were:

¹clipping level: the limits of maximum and/or minimum values that can possibly be recorded

² BIOSIG toolbox: <http://sourceforge.net/projects/biosig>

pass band frequency cutoff:	f_p	=	47 Hz
stop band frequency cutoff:	f_s	=	49 Hz
pass band ripple attenuation:	α_p	=	1 dB
stop band attenuation:	α_s	=	40 dB

The minimum filter order was calculated to be 7 by the function `ellipord`, which is found on the MATLAB signal processing toolbox. The magnitude and phase response are plotted in Fig. 3.4. Both the attenuation and the phase response of the filter above ~ 50 Hz greatly fluctuate, which is typical for elliptical filters. This is not a problem, however, because most of the EEG patterns studied (delta, alpha, spindle etc.) have a frequency less than 50 Hz. The exception is gamma frequency band which goes up to 60 Hz (Steriade and McCarley, 2005). To better observe the ripple attenuation in the magnitude, a zoomed-in graph has also been plotted.

In Fig. 3.5 the power spectral density (PSD) of the EEG is shown. Before filtering there is a very strong spike at 50 Hz. In fact, the interference is so large that the EEG just looks like a pure 50 Hz sine wave. After filtering, the PSD still shows the 50 Hz spike, but it is low enough that it will not be noticed in the EEG. The filter specifications required that the stop band attenuation was 40 dB, which can be observed on the magnitude response in Fig. 3.4. The PSD of the filtered EEG has the 50 Hz spike reduced to -80 dB. This is because the filter was applied twice via the function `filtfilt`, which filters the data forward then backward to remove any phase shifts.




Figure 3.4: The magnitude and phase response for the low-pass filter designed to remove the 50 Hz mains interference. As seen in the zoomed in plot, there are minimums in the ripple at ~ 19 , 40 and 46 Hz where the attenuation reaches -1 dB (scaled down by ~ 0.79)




Figure 3.5: The PSD of an EEG channel before and after low-pass filtering to remove 50 Hz electrical interference. The spike at 50 Hz has been removed

3.4.2 Filtering Rat EEG

As the rat EEG was recorded, the mains interference was filtered electronically by a 50 Hz notch. However, the notch was not strong enough to completely remove the mains interference. In Fig. 3.6 the PSD of a sample of rat EEG is shown. The 50 Hz peak, is so large that it emerges out of the notch. The harmonics of the 50 Hz signal, 100, 150 and 200 Hz, are also present in the PSD. An elliptical, low-pass filter, similar to the human EEG filter, was used to further lower these peaks. The only difference was that the stop band attenuation was only set to 20 dB, as the interference peaks were already at a value of about -50 dB




Figure 3.6: The rat EEG was filtered during the recording using a 50 Hz notch filter. The PSD of Rat EEG shows that the notch was not deep enough to remove all the 50 Hz. The PSD reveals that the harmonics of 50 Hz are also present in the EEG.

Chapter 4

Time Series Signal Processing

The number of data points that come from a sleep study can be very large. If the scalp voltages are sampled at 256 Hz, then for a 6 hour overnight sleep recording, there are about 5.5 million data points stored per channel. In order to analyse and display such large volumes of data, it is necessary to compute statistical averages that can be compared with cortical modelling predictions. The averaging calculations described here are designed to split each data channel into time-blocks or epochs of length, T , and then to evaluate selected statistics for each epoch. Four such statistics were used: total power, fractional band power, correlation time and spectral entropy.

4.1 Total Power

The total power of the EEG, voltage signal $V(t)$, averaged over the epoch period T , is defined to be:

$$P_{\text{tot}} = \frac{1}{T} \int_0^T V(t)^2 dt. \quad (4.1)$$

In discrete form, this equation becomes:

$$P_{\text{tot}} = \frac{1}{T} \sum_{i=1}^N V_i^2 \Delta t. \quad (4.2)$$

Remembering that the epoch length is $T = N\Delta t$, the equation can be rewritten as

$$P_{\text{tot}} = \frac{1}{N} \sum_{i=1}^N V_i^2 \quad (4.3)$$

Equation (4.3) can be implemented efficiently in MATLAB, as the formula can be applied to an entire data channel after a reshape command.

I find that the total power is a valuable first analysis tool. It shows how the squared amplitude of an EEG (or any other) signal changes with time, so it is useful for detecting artefacts, as these often show up as large power spikes that are generated when the sleeper moves causing the EEG to become corrupted with high-amplitude EMG.

4.2 Fractional Band Power

In order to get results comparable to Destexhe *et al.* (1999) (Fig. 2.7, Pg. 22), I needed to calculate the amount of power concentrated in different frequency bands and then normalise these against the total power. This was achieved by first taking the modulus-squared of the Fourier transform, \mathcal{F} , for each time epoch to give the power spectral density (PSD), $S(f)$:

$$S(f) = |\mathcal{F}[V(t)]|^2 \quad (4.4)$$

Then the fractional band power F in a specified spectral band $f_1 \leq f < f_2$ is given by

$$F^{\text{band}} = \frac{1}{P_{\text{tot}}} \int_{f_1}^{f_2} S(f) df \quad (4.5)$$

The superscript “band” indicates which frequency band f_1 – f_2 is being analysed. e.g., the alpha band would give $f_1 = 8$ Hz and $f_2 = 12$ Hz. In discrete form, Eq. (4.5) becomes

$$F^{\text{band}} = \frac{1}{P_{\text{tot}}} \sum_i S(f_i), \quad f_1 \leq f_i < f_2 \quad (4.6)$$

The fractional band power is usually calculated twice or more, with different values for f_1 and f_2 , giving a low frequency band such as 0.1–4 Hz and a high frequency band, 15–47 Hz. In the analysis, the high band limit of 47 Hz was chosen because a low-pass filter had been applied to the raw data to remove the 50-Hz mains power supply interference. The low band limit of 0.1 Hz was chosen to remove very low frequency artefacts arising from eye-blinks and movements. If $f_1^{\text{lo}} = f_{\text{min}}$ and $f_2^{\text{hi}} = f_{\text{max}}$, and the other frequency limits f_1^{hi} f_2^{lo} , are equal, then the two traces will be exact complements, changing in anti-phase with their sums adding to unity.

As we shall see in chapter 5, the fractional band power is very good at making




Figure 4.1: A demonstration of the results of the fractional band power calculation. $T = 2$ s, $f_s = 400$ Hz, window overlap = 80%, $f^{lo} = 0\text{--}3$ Hz and $f^{hi} = 3\text{--}8$ Hz

distinctions between SWS and REM sleep. During SWS, most of the energy in a sleeper's EEG is concentrated below 4 Hz. After the transition into REM sleep, the PSD of the EEG changes to contain a larger proportion of the total energy in the high frequencies bands. By calculating both a high and a low frequency power fraction, the two will change in anti-phase as the sleeper cycles between SWS and REM sleep.

4.3 Correlation Time

The autocorrelation of a time series can be used to estimate a signal's temporal coherence. The autocorrelation $C_{\text{true}}(\tau)$ of a zero-mean signal $V(t)$ is defined as

$$C_{\text{true}}(\tau) = \int_{-\infty}^{\infty} V(t) V(t + \tau) dt \quad (4.7)$$

The variable τ is referred to as the *lag time* ($C(t)$ measures the degree of similarity between $V(t)$ and $V(t+\tau)$, a copy of $V(t)$ which has been delayed (lagged) by time τ). $C(\tau)$ is always maximum when τ is zero because the signal is perfectly correlated with a zero-lagged copy of itself. The autocorrelation of a noisy or random signal drops very quickly and can be approximated by an exponential decay. The *correlation time* τ_C is then defined to be the value of τ at which the autocorrelation has fallen to $1/e$ of its zero-lag value:

$$\frac{C(\tau_C)}{C(0)} = \frac{1}{e} \approx 0.368 \quad (4.8)$$

In practice, Eq. (4.7) is not very helpful, as it is impossible to have an infinite data set. Another problem with this form of the autocorrelation is that in time-series analysis, we are looking for changes in the signal over time, so calculating the autocorrelation for the entire signal would be unhelpful, as we would miss all of changes we are looking for. It is better to take many autocorrelation samples, each of period T , where T is chosen to be small enough to track the changes in correlation statistics, but large enough to smooth out the random fluctuations in the statistic. For these reasons a *biased*, finite autocorrelation is used to estimate the true autocorrelation:

$$C(\tau) = \frac{1}{T} \int_0^T V(t) V(t+\tau) dt \quad (4.9)$$

In a discrete form, Eq. (4.9) becomes:

$$C_j = \frac{1}{N} \sum_{i=-N+1}^{N-1} V_i V_{i+j}, \quad j = -(N-1), \dots, N-1 \quad (4.10)$$

Here j is not the lag time, it is the lag index, and N is the number of samples in the epoch, $N = f_s T$. To calculate the lag time, j must be divided by the sampling rate of the signal:

$$\tau = \frac{j}{f_s} = j \Delta t \quad (4.11)$$

Using MATLAB, Eq. (4.10) is evaluated with the `xcorr` function, found in the signal processing toolbox:

```
[C,lags] = xcorr(data,'biased');
```

the 'biased' option tells `xcorr` to divide by N as in Eq. (4.10). The second output argument `lags` is optional, but it is useful, as it corresponds to all the possible values of j , $-(N-1) \dots (N-1)$.

To demonstrate the calculation of the correlation time, I created some random

data with mean 0, standard deviation 1, and time step 1 ms. The data was then smoothed by using a Whittaker smoother (Eilers, 1994) with different smoothing strengths, to give three new signals, labelled 1, 2 and 3. The autocorrelation function for each signal was calculated. As shown in Fig. 4.2 the correlation time is longer for the smoother, less random signal-3 than it is for the less smoothed, signal-1. This means that signal-3 is more coherent than signal-1.

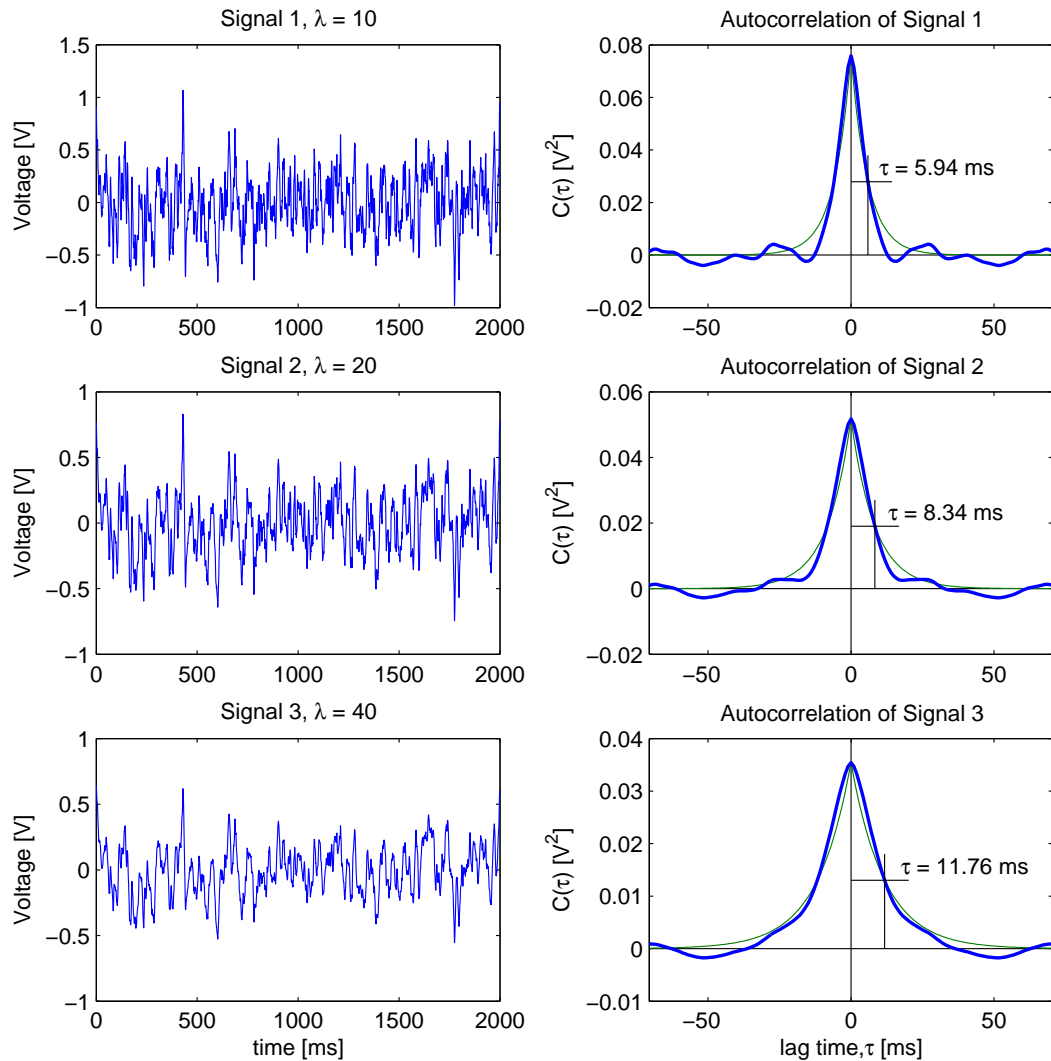


Figure 4.2: Left: A random, Gaussian signal smoothed using a Whittaker smoother (Eilers, 1994); λ specifies the degree of smoothing. Right: The autocorrelations of the random signal and the best-fit exponential decay curves. The smoother signals have longer correlation times. The time axis extends all the way from -1999 ms to 1999 ms, but we are only interested in the decay centred at $t = 0$

4.4 Spectral Entropy

Entropy is a measure of the disorder in a system. The *spectral* entropy is then a measure of the disorder in the frequency spectrum of a signal. White noise, which has a flat spectrum, is highly disordered, and therefore has a high entropy. A pure sine wave, on the other hand, has only a single spike in its spectrum, meaning all the energy is concentrated in a very ordered state, and hence its spectrum has a low spectral entropy.

A starting point for the spectral entropy is the discrete Shannon information entropy H_1 (Shannon and Weaver, 1949),

$$H_1 = - \sum_{i=1}^N p_i \ln p_i. \quad (4.12)$$

where p is the probability density function approximated from the PSD $S(\omega)$ with a histogram which has a bin width of $\Delta\omega$, and N is the total number of spectral bins. Hence $\omega_i = i\Delta\omega$, $i = 1, 2, \dots, N$. Because each p_i is a *probability*, $\sum_i p_i = 1$.

Steyn-Ross (2002) showed that the discrete Shannon form is unsatisfactory since it gives a *biased* estimate of the spectral entropy of a continuous signal that has been discretely sampled. It turns out that the bin width $\Delta\omega$ plays an important role in correctly calculating the entropy. Steyn-Ross shows that the *histogram* spectral entropy H_2 gives an *unbiased* estimate of the underlying continuous spectral entropy:

$$H_2 = -\Delta\omega \sum_{i=1}^N s_i \ln s_i \quad (4.13)$$

where

$$s_i = \frac{p_i}{\Delta\omega} = \frac{S_i}{\Delta\omega \sum S_i} \quad (4.14)$$

and

$$\sum_{i=1}^N s_i \Delta\omega = 1. \quad (4.15)$$

In order to normalise Eq. (4.13), Steyn-Ross then shows that the maximum entropy possible is $H_2^{\max} = \ln[N\Delta\omega]$. The normalised histogram spectral entropy is now

$$H_2^{\text{norm}} = \frac{-\Delta\omega}{\ln[N\Delta\omega]} \sum_{i=1}^N s_i \ln s_i. \quad (4.16)$$

To demonstrate the spectral entropy, I have used the same original functions as in Fig. 4.1 and then also added zero-mean noise whose standard deviation is

exponentially growing in amplitude. Figure 4.3 shows the spectral entropy grow and reach a local maxima as the 5 Hz pulse reaches the same magnitude as the 2-Hz sinewave. The entropy then drops again as more energy is concentrated in the pulse. Later, at ~ 5 s, the random noise starts to increase and the spectral entropy grows towards unity.




Figure 4.3: A test signal formed by summing a 2-Hz sinewave, a 5-Hz and exponentially growing noise (top two panels). The spectral entropy of the composite signal is shown in the bottom panel.

4.5 Epoch Overlap

In time-series analysis, the calculations are usually performed on small time epochs of length T . For example when looking at 6 hours of data, perhaps one point for each 30 seconds is appropriate and then a smooth curve can be fitted to the data. However, when looking at smaller time periods of only a few minutes, then epochs lengths of 30 s, or even 5 s, may not give enough data points to observe changes accurately. To overcome this paucity of observation points, I decided to *overlap* the time epochs. The use of overlapping epochs is quite common and in the literature,

the simplest method would be to make a calculation on data in the time range of $[t, t+T]$ and then increase t by $(1-X) \times T$, where X is the overlap fraction, before making the next calculation.

```

X = overlap fraction
T = time epoch
i = 1
while (t+T) < length(data)
    answer(i) = function(data(t:t+T))
    t = t + (X-1) × T
    i = i + 1
end while

```

This method is fine for (compiled) languages such as C or FORTRAN, but in MATLAB it does not make use of the very fast array and matrix operations available. In fact, in MATLAB this method runs very slowly.

4.5.1 A New Overlapping Algorithm

I thought of a better way to overlap the time epochs in MATLAB. My idea was to calculate the result for the whole data channel, as if there was no overlap, then re-calculate with the starting position moved forward by $X \times T$, and repeat until all the required overlaps have been processed. Figure 4.4 shows how a data vector of length $17 \times \frac{T}{4}$ seconds would be calculated in four iterations if the epoch overlap were set to 75%.

After the overlapping algorithm has finished, the resulting array from figure 4 would look like

$$\begin{pmatrix} 1a & 1b & 1c & 1d \\ 2a & 2b & 2c & 2d \\ 3a & 3b & 3c & \text{NaN} \\ 4a & 4b & 4c & \text{NaN} \end{pmatrix}. \quad (4.17)$$

The array is then reshaped columnwise back into a vector and the NaN's are removed from the end as they are just place holders.

$$(1a \ 2a \ 3a \ 4a \ 1b \ 2b \ \cdots \ 1d \ 2d) \quad (4.18)$$

This is now a new time-series with four data points every epoch T .

The relationship between the fraction of overlap X and the number of iterations




Figure 4.4: An illustration showing how the overlapping algorithm, set to 75%, would perform calculations on a data vector. The full calculation over the vector takes four iterations, with the starting position increased by $\frac{1}{4}T$ each time. Note the two NaN's where 3d and 4d should be, which are there because there is not enough data in the vector to define a full epoch (NaN means “not a number”).

comes from looking at the fraction that the epoch *advances* δ each iteration:

$$\delta = 1 - X. \quad (4.19)$$

For example, in Fig. 4.4 the overlap fraction is $X = 0.75$, so the advance per iteration is $\delta = 0.25$. The number of iterations M is then

$$M = \frac{1}{\delta} = \frac{1}{1 - X} \quad (4.20)$$

Numerically, it helps if X is a rational fraction, whose numerator and denominator are both factors of N the number of samples in the epoch. This is because at each iteration the code moves the starting position by $\delta \times N$ samples, and errors can occur if indices are not integers. In the beginning of the algorithm I have a few lines of test code to convert δ into a rational fraction and change its value if necessary.

Usually X is set to a fraction of the form $\frac{x-1}{x}$, such as $\frac{1}{2}$ or $\frac{3}{4}$. This means δ has the form $\frac{1}{x}$. If this is not the case, i.e., the numerator is greater than 1, then the example in Fig. 4.4 is not correct. Figure 4.5 shows an example of the overlap, X set to $\frac{1}{3} \implies \delta = \frac{2}{3}$.

As the eye follows down the first column a the epoch is moved to the right by $\frac{2}{3}$ each time. At the start of the next column b, one notices that 1b occurs *before* 3a. The epoch at 1b cannot exist with the overlap set to $\frac{1}{3}$. The next epoch $\frac{2}{3}$ after epoch




Figure 4.5: With the overlap set to 33.3%, the second and fourth columns (b and d) are never used.

3a is epoch 1c. It is clear from the diagram in Fig. 4.5 that the whole of column b is redundant and the same holds for the column d. The resulting array will for this example will look like

$$\begin{pmatrix} 1a & 1c & 1e \\ 2a & 2c & 2e \\ 3a & 3c & \text{NaN} \end{pmatrix} \quad (4.21)$$

As columns b and d are not required, there is no point in calculating them. The way to find which columns will be redundant comes from the advance fraction δ . The fraction δ can be expressed as, $\frac{\text{numerator}}{\text{denominator}} = \frac{n}{d}$. Then only every n^{th} column after the first needs to be calculated. The best way to do this is to remove the unwanted data at the start of each iteration. In MATLAB this is easy as the data is first reshaped from a vector to an array. Then the unwanted data can be removed with the command

```
newdata = data(:,1:n:end);
```

To demonstrate the effectiveness of overlapping the epochs, I have plotted the total power for 2 minutes of rat EEG in Fig. 4.6. It has been calculated with `sleppower` and the amount of overlap was set to 0%, 50%, 80% and 95%. At 0% and 50%, the plot looks a little sparse, and one is left wondering what happens in between the points. At 80%, there is five points in every epoch, and the graph looks good. Not much more is gained by going up to 95%, but the computation time now begins to become annoying, because $\frac{1}{1-0.95} = 20$ iterations. I have found that for most small data sets (about 2 min long) 75% or 80% overlap is the best compromise between detail and computation time.

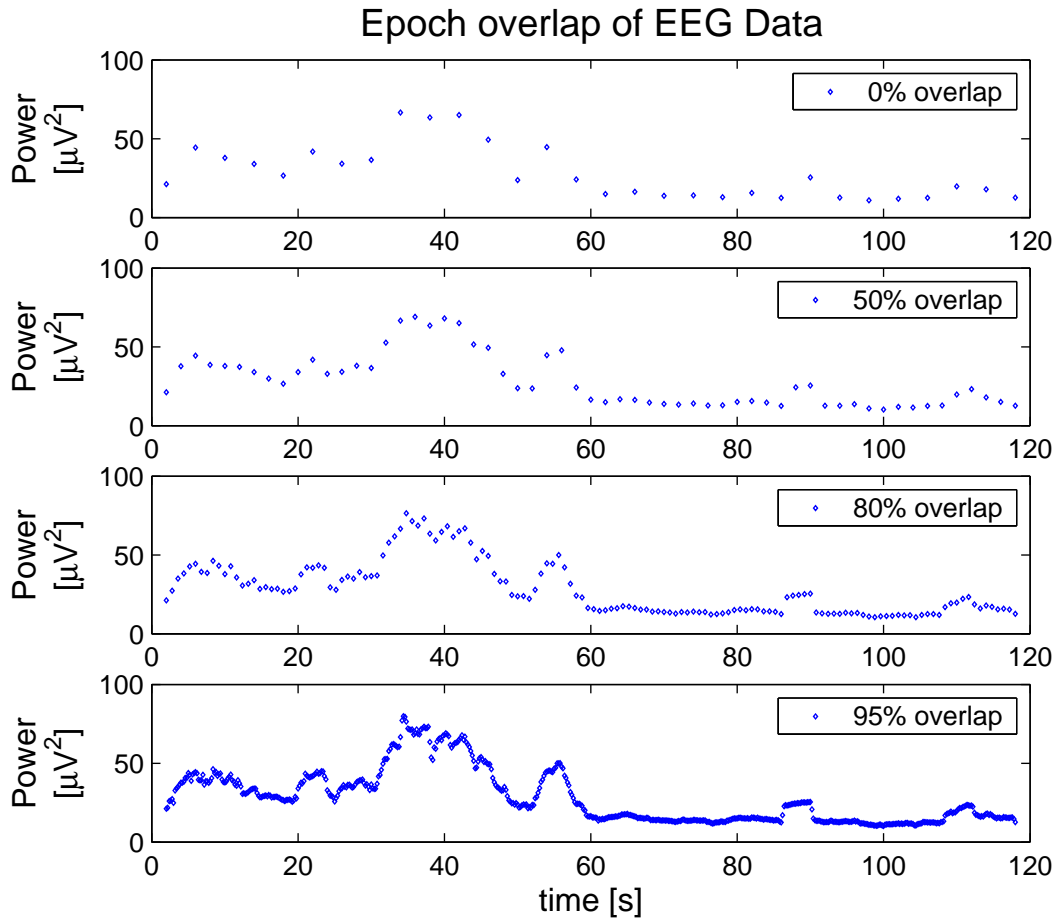


Figure 4.6: The power of some rat EEG data calculated by `sleeppower`. The epoch was set to 4 s and the amount of overlap was varied: 0%, 50%, 80%, 95%

4.5.2 Overlap Time Trials

To test the effectiveness of the epoch overlapping algorithm, the total power of a artificial data set, of length 10^7 samples, was computed. My overlapping algorithm was compared with the simple one outlined on page 46. The power was calculated and the computation time was recorded. The results are shown in table 4.1.

The simple algorithm was slightly faster only for the case where there was no overlap, but for high overlap fractions, it is outclassed. Figure 4.7 shows a plot of the data. For both methods, the computation time increases with the amount of overlap, but the growth is much slower for my new method.

Table 4.1: The time for the power to be calculated for a 10^7 element data vector with varying degrees of epoch overlap.

Overlap X	Simple Algorithm time (s)	New Algorithm time (s)
0	0.581	1.542
0.5	3.368	2.751
0.75	15.32	4.927
0.8	23.27	5.998
0.875	62.59	9.304
0.9	97.32	11.51
0.95	363.7	22.33




Figure 4.7: The time taken to compute the power of a large data vector (10^7 samples) using two different epoch overlapping methods

Chapter 5

Human Data Results

The Steyn-Ross model makes predictions on how the EEG recorded from a sleeping cortex will change over time. The single prediction that is the focus of this thesis is that when the cortex makes the transition from EEG-synchronised SWS and goes into the active, paradoxical REM sleep, this is in fact a first-order phase transition, analogous to ice melting into water. When any system undergoes a first-order phase transition, there are measurable changes in power, temporal/spatial correlation and entropy. With the use of the time-series statistics described in chapter 4, we can attempt to detect the SWS to REM transitions and compare them with the theory common to all first-order phase transitions.

For our cortical phase transition, the model predicts a surge in the low-frequency power of the EEG during SWS as the approach towards REM is made. Simultaneously we should see an increase in the correlation time and a reduction in the spectral entropy. At the SWS to REM transition, the model predicts a sudden drop in the total power and the energy should spread out into higher frequencies. This is accompanied by the correlation time becoming shorter and the spectral entropy increasing (Steyn-Ross *et al.*, 2005).

In this chapter I present the results that came from the human sleep volunteers. The subjects were all suffering from obstructive sleep apnoea, but were being successfully treated using Fisher & Paykel Healthcare's CPAP (Continuous Positive Airway Pressure) respirators. The patients' EEG were recorded from four locations of the scalp. Using the 10-20 system of standard electrode positions, the electrodes were placed at C₃, C₄, A₁ and A₂, and were referenced to F_{pZ}, at the front of the forehead.

5.1 Sleep Staging

Before the time-series of the sleep records could be analysed, I needed to look through the data set and double-check the stage of sleep that had been assigned to each 30-second epoch of the time-series. The staging had already been scored by the sleep technician. However, I still needed to go through the data to get a feel for the number of artefacts. I wrote the software tool `sleepwork` (described in Appendix A.2.2) to assist with the viewing of the time-series, assigning of stages and flagging of artefacts. The staging was scored through observation of the EEG, EMG and EOG. From these three bio-signals, one can discern the correct stage of sleep.

Because sleep staging is scored manually, the application of the standard rules by Rechtschaffen and Kale (1968) are subject to human error. Different people who score the same time-series will usually come up with discrepancies. There are times when it is difficult to assign one particular sleep stage. The EEG may not show the telltale signs of any particular stage (spindles, slow-waves, etc). The EOG can show light activity that causes the epoch to resemble REM. It seems that sometimes epochs are scored as stage 2 because of a lack of anything else to call them. An example of this is shown in Fig. 5.1.

Examples of possible misclassifications in the sleep records have been shown in Figs. 5.1 and 5.2. Both figures show epochs that have some high-amplitude and low-frequency waves and neither contain sleep spindles. The epoch in Fig. 5.1 has been classified as stage 2, but it contains more activity in the EOG (labelled LOC: left ocular, ROC: right ocular) than the epoch in Fig. 5.2 which has been scored as REM.

Another problem with manual sleep staging comes from the scorer sometimes relaxing the rules depending on the need for accuracy and speed. For instance, sometimes every epoch is not given its own score; the scorer just notes the time and new stage every time the sleeper changes state. Recording of sleep stages in this way means that brief changes are ignored. To make matters worse, when the sleeper is alternating between two stages, or is very close to the boundary, a rough average of the two is given. Stage 3 and 4 are often combined to become 3.5, while stage 1 and 2 become 1.5.

It seems that the sleep staging can be somewhat arbitrary and for the purpose of this analysis it should be used only as a guide.




Figure 5.1: This epoch was classified as stage 2. The EEG is a little high-amplitude and contains some low frequencies, but it does not have any of the stage-2-defining spindles. LOC and ROC are the left and right EOG channels. The EOG has enough movement for the technician to consider scoring the epoch as REM.




Figure 5.2: This epoch has similar EEG to Fig. 5.1. It has less activity in the EOG, but it was classified as REM by the sleep technician (the downward spikes are ECG artefacts).

5.2 Pre-processing – Artefacts Detection

When an artefact arises in the EEG, it is worthwhile checking to see what could have caused it. It is helpful then, to have access to signals other than the EEG, EOG and EMG. The sleep records I received contained a leg movement channel, an airflow channel and even a microphone. These were all very useful in detecting artefacts. The leg movement channel was good for detecting full body movements which cause the EEG to completely overrun with high-amplitude noise. Although the sleepers were all being successfully treated for sleep apnoea, they would still occasionally stop breathing and wake up. The airflow channel was helpful in explaining artefacts arising from these awakenings. The microphone channel was helpful if an artefact came from a cough or snoring.

While the patients were awake, the amplitude of their EEG was rather low, but there were often many artefacts that came from eye-blinks and other movements. In terms of flagging unwanted or bad data, the waking state and artefacts were handled in the same way. I assigned them a sleep score of zero = awake. If the patient was actually awake, the whole 30-second epoch would be scored awake. It was unusual for an artefact to last 30 seconds, so for the flagging of small artefacts, the 30-second epochs were broken into six five-second epochs. Then only a subset containing the artefact needed to be scored as awake and the rest of the epoch can remain as its correct status. Figure 5.3 is a screen shot taken from *sleepwork*. It shows a large movement artefact in the middle of a stage 3 epoch. By splitting the 30 seconds into six epochs, only 10 seconds need to be flagged as an artefact.

The total EEG power of patient one has been calculated and graphed in Fig. 5.4. Here one can see the need for artefact detection, as the sleep cycles are not well observed among the high-amplitude artefact spikes. In contrast, the same data is plotted again in Fig. 5.5 with the artefact detection turned on. There are still large amounts of fluctuation in the time-series, so a smooth curve has been fitted to the data.

One observation that can be made directly from Fig. 5.5 is the consistency over the four channels. The electrodes were positioned approximately 5 cm apart on the patient's scalp, but on this time scale the power traces are extremely similar, showing a high degree of synchrony across the cortex.

During REM there are many high-amplitude eye-movements that can corrupt the EEG. They occur so often that to flag each one as an artefact would leave out a great deal of the time-series. I found that a reasonable way of removing the REMs was to use a 1 Hz high-pass filter. In Fig. 5.6, 30 seconds of patient two's sleep record are plotted. The left and right EOG (LOC and ROC) are wired in such a




Figure 5.3: The patient is in stage 3 sleep. There is a movement artefact at time = 1:30:55. There is artefact signal on the microphone and the leg channel, indicating that the patient has momentarily stirred then gone back to sleep. Notice how much the EEG, LOC/ROC and EMG channels are corrupted. This artefact would produce a large spike in the results and therefore must be flagged so that it can be eliminated from the analysis. (The EEG channels 1 and 2 are the difference between the EEG recorded from two scalp locations as recommended by Rechtschaffen and Kale (1968): $EEG1 = C_4 - A_1$, $EEG2 = C_3 - A_2$.)

way that REMs produce voltage signals that change in anti-phase, while blinking of the eyelids makes the signal change in phase. In calculations of the total power of the EOG, the ROC was subtracted from the LOC so the in-phase eye-blinks and high-amplitude slow-waves from the EEG would cancel out. Meanwhile, the out-of-phase REMs would combine to create a stronger signal.

5.3 Full Sleep Record Results

In this section I present the results of the time-series analysis. In the following subsections (5.3.1–5.3.4), I have plotted the EEG time-series statistics described in chapter 4. The total power, fractional band power, correlation time and spectral entropy are all graphed as functions of time. The power of the EOG and EMG are also included as these can help in distinguishing SWS from REM sleep. The sleep staging that was scored by the sleep technician is also included.

When calculating these statistics, I used a consistent set of parameters. The the




Figure 5.4: The total power calculated from four EEG channels without the artefacts being flagged. There was one point calculated every 5 seconds. The interesting parts of the time-series are the clear sleep-cycles at 1–3 hours. The large spikes are from the unflagged artefacts.




Figure 5.5: The same data as shown in Fig. 5.4 with the artefacts all flagged and removed. The two sleep-cycles at the beginning become much more clear. A smooth curve has been fitted to the data using the Whittaker smoothing technique described in Appendix A.3.2.




Figure 5.6: Rapid eye-movements during REM sleep are the cause of some artefacts in the EEG. A 1 Hz high-pass filter removes most of the artefact without losing much of the real EEG signal.

epoch T length for each calculation was 5 s. The overlap fraction was set to zero, as the raw time-series data are so long I did not require intermediate points. The sample frequency f_s was 256 Hz for the first three patients, whose recordings were taken at Fisher & Paykel Healthcare. The other three patients (four, five and six) had their sleep records taken at Greenlane Hospital and the sample rate was only 125 Hz.

The frequency bands used in the fractional band power calculation had the default value: $F^{\text{lo}} = 0.1\text{--}4$ Hz and $F^{\text{hi}} = 15\text{--}47$ Hz. These were chosen because they were the same as Destexhe *et al.* (1999) used for their sleeping cat. (Their high-frequency limit was actually 75 Hz but I was low-pass filtering to remove 50 Hz mains power supply interference.) The frequency limits used for evaluating the spectral entropy were 0.1–47 Hz. The minimum frequency bin width is determined by the epoch size, $\Delta f = \frac{1}{T}$. The 5 s epoch gave a bin width of 0.2 Hz (see Appendix A.1.4).

Because the resulting graphs were so noisy, a Whittaker smoother was used to draw a smooth curve to the data. The difference order was $d = 2$ and the smoothing strength was $\lambda = 100$ (see Appendix A.3.2).

Graphs showing only the fractional band power calculated for five different frequency bands are also plotted. This is because, in humans, many different EEG patterns occur in the range 4–15 Hz. Spindles, one of the defining characteristics of stage 2, have frequencies in the range of 12–15 Hz. Alpha activity, 8–12 Hz, is associated with the subject being awake and relaxed, with eye closed (Steriade and McCarley, 2005, Pg. 256). The alpha band can serve as an indicator to see if sleep has begun, as it disappears when the subject has fallen asleep.

The common frequency bands used in EEG analysis are listed in table 5.1. The exact values of the frequency limits may vary from one author to another. The frequency limits used for this analysis came from Marshal *et al.* (2003) with the exception of the gamma band stopping at 47 Hz from my low-pass filter (Marshal has a 30 Hz low-pass filter so leaves out gamma altogether), and the slow-wave band beginning at 0.1 Hz (Marshal uses 0.2 Hz).

Table 5.1: Frequency bands commonly used in EEG analysis.

Frequency Band	Frequency range (Hz)
slow-wave	0.1–1
delta	1–4
theta	4–8
alpha	8–12
spindle	12–15
beta	15–25
gamma	25–47

5.3.1 Patient One

The time-series analysis of patient one is shown in Fig. 5.7. The top graph is the total power time-series. When the patient is in deep SWS (stage 3 or stage 4) the EEG amplitude is high, resulting in a large power. REM sleep, on the other hand, has a much smaller amplitude than SWS and therefore a lower power. Patient one went through two clear sleep cycles during the first few hours of the recording. He then woke up at around the time his next REM phase should have occurred. The last cycle consists of mostly light, stage 2 sleep, then a long REM phase.

We are particularly interested in the transitions from SWS to REM sleep. These are marked by the vertical dashed lines. The first of the three transitions into REM is the clearest. There is a sudden drop in the total power, and from the fractional band

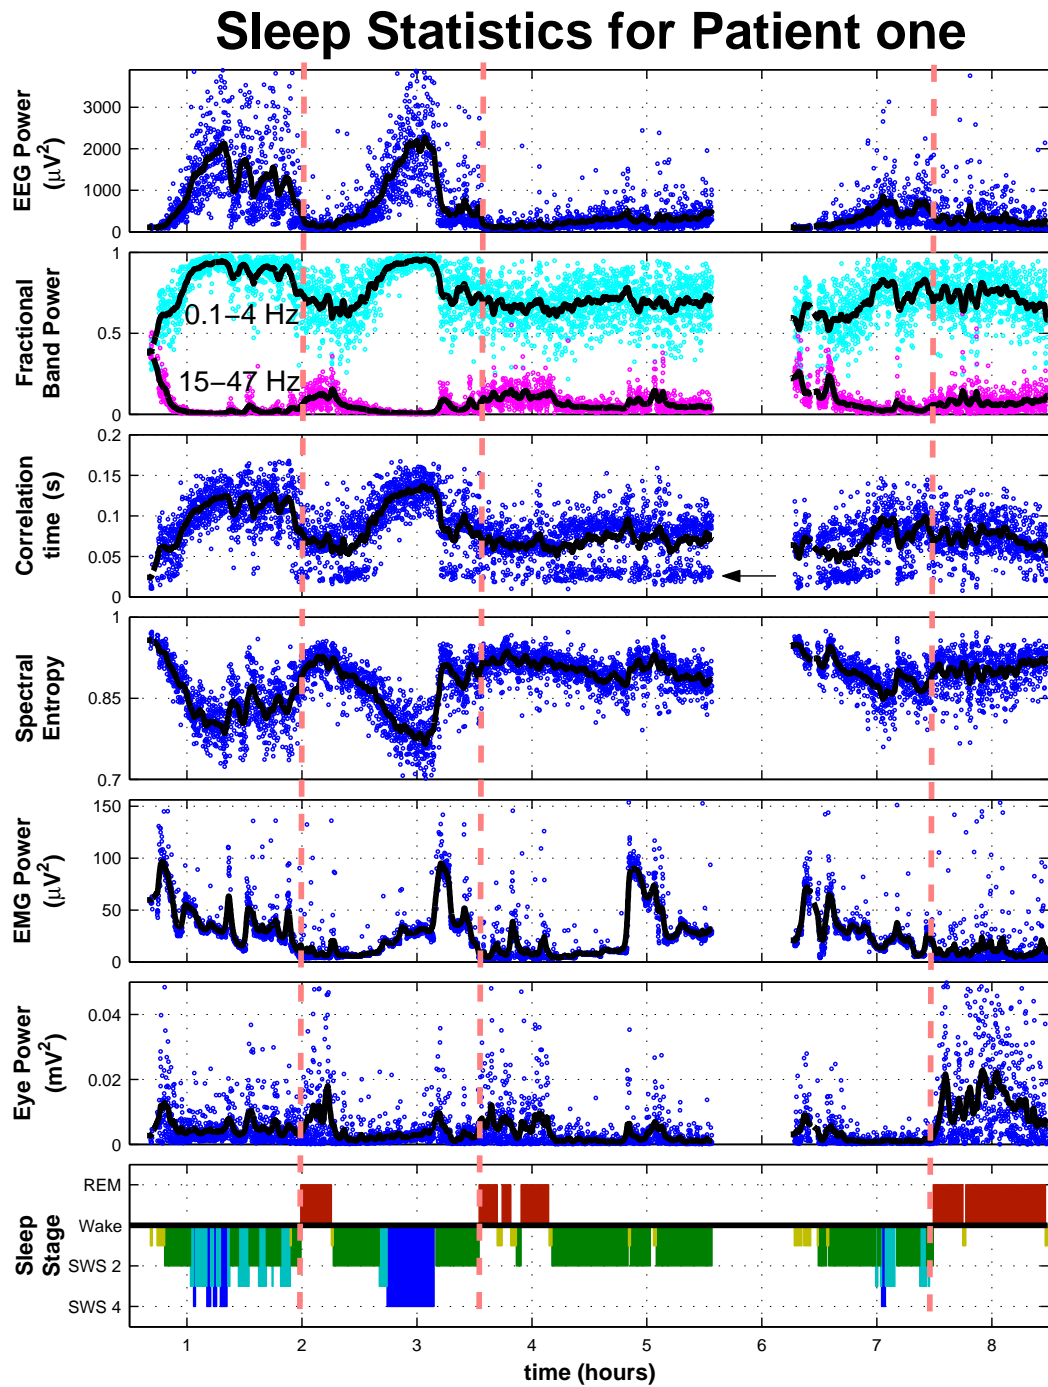


Figure 5.7: The time-series analysis of the sleep recordings of patient one

power we see there is a change in the spectral distribution of the energy. Before the transition, almost 100% of the energy is concentrated below 4 Hz. After the transition, however, it is reduced to about 75%. This is also seen by the 15–47 Hz band increasing from almost zero before the transition, to about 10% of the total power after the transition. (The middle band of 4–15 Hz also increases, but it was excluded from the graph to avoid cluttering.) In Fig. 5.8 the fraction band power calculated for several frequency bands has been graphed.

The high correlation time during SWS is due to the high-amplitude and low-frequency of the slow-waves. Once the cortex has transitioned into REM, the correlation time is greatly reduced. This is because the slow-waves go away with the drop in low-frequency power, and all that is left are the higher frequencies that have a shorter correlation time. The work of Steyn-Ross (2002) suggests that spectral entropy will change in anti-phase with the correlation time. This is exactly what we observe here. During the highly coherent SWS phase, the spectral entropy is low, meaning the energy in the signal is well ordered in the low-frequency band. After the transition, when the energy is more evenly spread across all frequencies, the spectral entropy is higher, because the spectrum is in a more disordered state.

The second sleep cycle contains an even more dramatic change in the time-series statistics at $t \approx 3:10$ hours. However, this is not a transition from SWS to REM. It is a transition from stage 4 to stage 2. After inspecting the raw EEG time-series, I believe the scorer has correctly staged this period as stage 2 as there are many spindles present. The transition into REM occurs slightly later. The predicted changes in the time-series statistics are still present, but not nearly as distinguishable as they were in the first sleep cycle.

The third REM period occurs near the end of the sleep record at $\sim 7:30$. Here the patient only briefly enters deep, stage 3 or stage 4 SWS, then goes back into stage 2. At the transition into REM, the total power reduces slightly, but it is not very noticeable, as the power was not very large before the transition.

If attention is turned back to the graph of the correlation time in Fig. 5.7, one will notice a large concentration of data points, marked by the arrow, at a value of ~ 25 ms which is lower than all the other data points. These occur primarily during stage 2 and arise from the autocorrelation of spindles. In Fig. 5.9 (a) the autocorrelation of some typical stage 2 EEG is plotted. The exponential decay approximation matches the drop in correlation reasonably well down to the point $1/e$. On the other hand, Fig. 5.9 (b) is the autocorrelation of stage 2 EEG that contains a spindle. The exponential decay follows the first downward oscillation of the autocorrelation. This gives a correlation time that is much quicker than it should




Figure 5.8: The fractional band power of patient one calculated for five frequency bands. The spindle band is helpful in distinguishing low-amplitude stage 2 from REM, as REM does not contain sleep spindles. The alpha band can assist with detecting the relaxed but awake period before the onset of asleep.

be. The exponential decay curve should pass through the middle of the oscillations in the autocorrelation decay.




Figure 5.9: The correlation time of stage 2 sleep EEG is around 0.6–0.8 s as seen in (a). When there is a spindle in the EEG epoch the autocorrelation oscillates at the spindle frequency, ~ 13 Hz, and a much shorter correlation time is measured. The oscillations in the autocorrelation should be removed before calculating the correlation time.

Spindles have a frequency of ~ 13 Hz which is a period of 0.077 s. The spindle is about $1/3$ or $1/4$ through its period when it reaches $1/e$, which gives correlation time as 0.019–0.026 s. These numbers are consistent with what is seen on the

graph. Ideally, the spindle resonance in the autocorrelation should be removed before computation of the correlation time for a more accurate result. This is one of the unfinished tasks that I would like to complete in the future.

5.3.2 Patient Two

The sleep record for patient two shows similar trends to that of patient one. During stage 3 and stage 4 SWS the power is high and concentrated in the low frequencies. At 70 minutes into the recording, there is what looks like a transition, but the patient does not enter REM; he wakes up. The waking and REM states are very similar, and this event could still be a phase transition. In earlier work, Steyn-Ross *et al.* (1999, 2001a,b) give evidence that the transitions from wakefulness to induced unconsciousness via an anaesthetic, and back again, are also both first-order phase transitions.

At $t = 2:10$ hours, the patient enters REM for the first time. This is an interesting transition, as the patient was in stage 3 SWS and seems to transition down to stage 1 for about five minutes, then makes the full transition into REM. All of the calculated statistics show a kind of step pattern. The power drops as stage 3 goes to stage 1, then it drops again as stage 1 goes to REM. The fractional power shows two successive reductions in low-frequency power. There are also drops in correlation time and jumps in spectral entropy. In Fig. 5.11, where different fractional power bands have been plotted, there is a jump in the spindle band but then it falls away after REM begins.

The second SWS to REM transition occurs at about four hours. Although the patient was only in stage 1 and stage 2 prior to the transition, there is still a good deal of change in the sleep statistics. Part-way through the REM phase the total power of the EEG seems to increase and die away again when the patient leaves REM and starts the next sleep cycle. The other statistics show trends that are consistent with power increases. This increase in power is actually caused by REM artefacts. This can be verified by a comparison with the eye-power time-series, which has a very similar shape.

The third SWS to REM transition also shows the same changes in sleep statistics that are consistent with a phase transition, although there is not a great deal of contrast.

It should be noted that the raw EEG time-series contains a lot of ECG signal which can be seen as sharp spikes at very regular intervals. The ECG included in an EEG signal does slightly change the results in the sleep statistics, but the sleep cycles are still visible and not changed in time. This is investigated later in

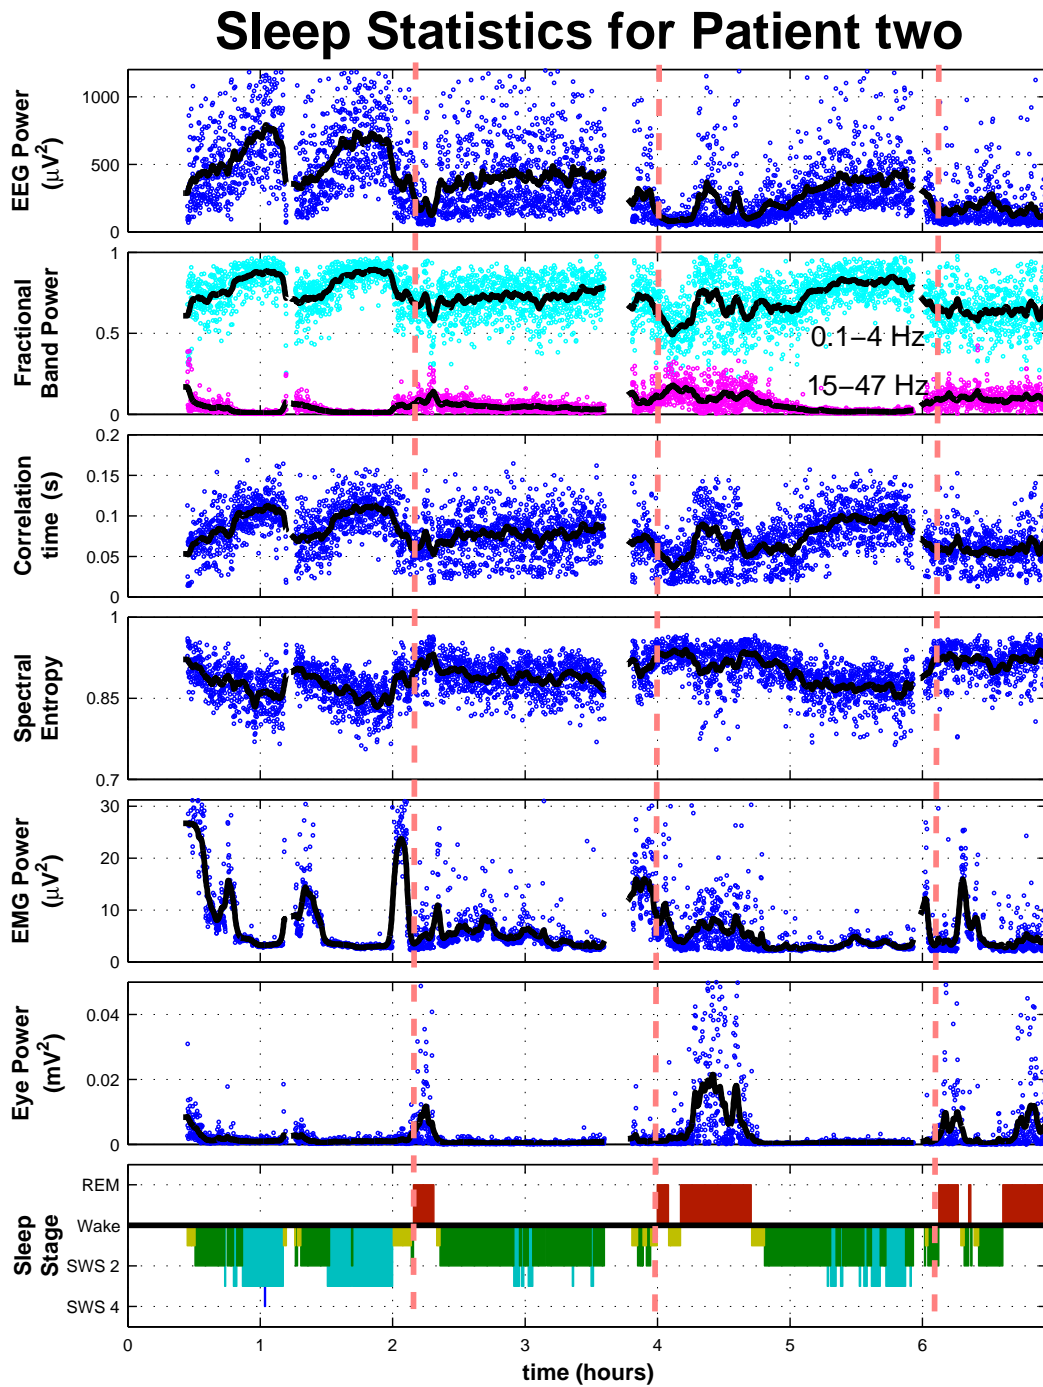


Figure 5.10: The time-series analysis of the sleep recordings of patient two




Figure 5.11: The fractional band power of patient two calculated for five frequency bands.

section 5.4.

5.3.3 Patient Three

Patient three's sleep record is shown in Fig. 5.12. He has four REM phases. However, the last is very short and the preceding SWS is low-amplitude stage 2, making the transition difficult to determine from the statistics. The other three transitions are much clearer.

The second sleep cycle contains a very sharp SWS to REM transition. The power steadily grows and then suddenly drops away. The same abrupt change is seen in the other statistics. During SWS the EEG amplitude and the correlation time are increasing, while the spectral entropy is decreasing. Then, while in deep sleep, the cortex abruptly transitions into REM.

The fractional band power has been plotted again in Fig. 5.13 showing more frequency bands. This figure is different from the similar graphs of patients one and two (Fig. 5.8 and Fig. 5.11). They show strong differences in the 8–12 Hz, 12–15 Hz and 15–47 Hz bands during stage 2 and REM. For patient three, however, all these bands show very similar trends. They all are high during REM, very low during stages 3 and 4, and at an intermediate level during stage 2.

In Fig. 5.12 the smooth curve is perhaps too smooth. To show how abrupt the transition is, I have plotted three minutes of the raw EEG and EOG time-series in Fig. 5.14. Here we see on the left of the figure, that the EEG amplitude is pulsating high and low with a period of ~ 10 s. The slow-waves switching on and off are possibly evidence of a phenomenon called an *infraslow oscillation* described by Vanhatalo *et al.* (2004). During the SWS phase, the EOG remains low and uneventful. At $t = 2:48:50$ hours, the last pulse of slow-waves ends and large REMs begin. The EEG remains low in amplitude, while the EOG becomes active with the sharp, unpredictable movements that are characteristic of REM sleep.

The transition shown in Fig. 5.14 is the best SWS to REM transition I have found out of the six human sleep records that were provided to me. The change in state occurs very quickly, as if a switch has been flipped.

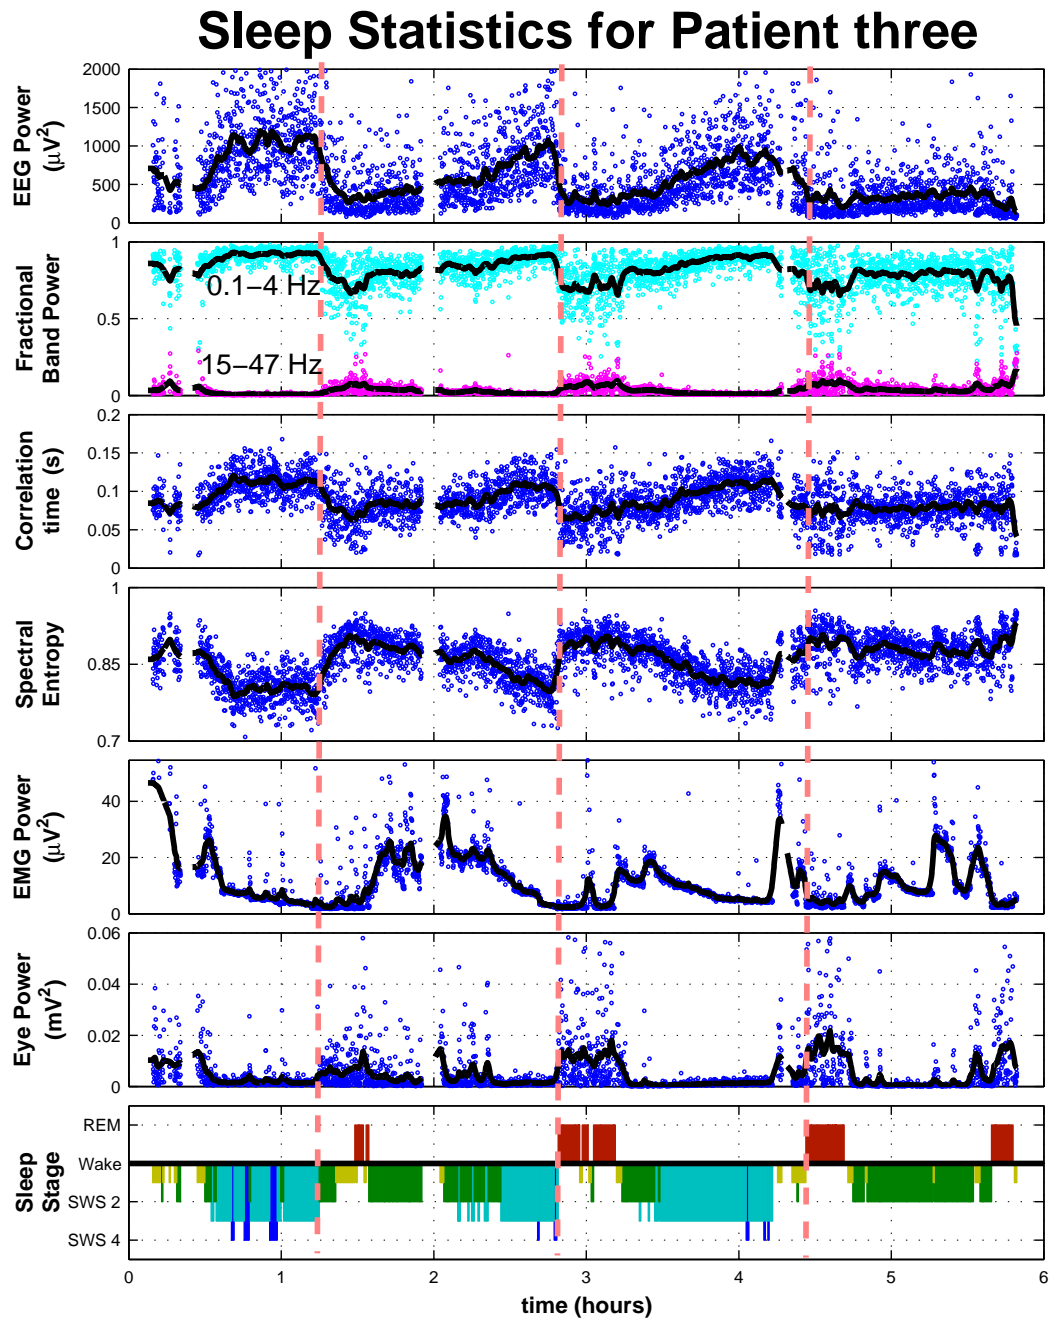


Figure 5.12: The time-series analysis of the sleep recordings of patient three




Figure 5.13: The fractional band power of patient three calculated for five frequency bands.




Figure 5.14: The raw time-series of patient three. At the start of the graph, the patient is in deep stage 3 SWS. Then, very suddenly, the EEG amplitude reduces and large REMs begin. The thick lines in the EEG time-series are total power (amplitude-squared) scaled to fit on the graph. The oscillations between high and low-amplitude during SWS have a period of ~10 s and are reminiscent of infraslow oscillations. (Vanhatalo *et al.*, 2004) (The repeating downward spikes in EEG2 are ECG heartbeat artefacts)

5.3.4 Patients Four, Five and Six

The three patients four, five and six all had their sleep recordings taken at Greenlane Hospital. At Greenlane, the data acquisition equipment is older than that used at Fisher & Paykel Healthcare, so electronic storage space is saved by lowering the sampling rates, and clipping levels.

Patient four has only one clean SWS to REM transition in his sleep record which is shown in Fig. 5.15. At 5:50 hours, he leaves stage 4 SWS and enters REM with only a few minutes of stage 1 in between. At this transition all the signal processing statistics have trends that are consistent with our phase transition theory. In Fig. 5.16, we see that for this transition there is a great increase in the fractional band power for the theta, alpha, beta and gamma bands, but the spindle band increases only slightly at the start of REM then drops away.

Patient four's sleep record is far from ideal, as he woke often throughout the night, and one of the EEG channels is bad due to a loose electrode.

Figure 5.17 is the sleep record of patient five. It has three good sleep cycles, the first of which contains the cleanest SWS to REM transition at 2:00 hours. For this transition, the theoretical predictions are observed. In Fig. 5.18 the fractional band power of all the bands higher than 4 Hz increases for all of this REM phase. The other two transitions do not show such clear changes in state. As with patient two, the EEG of patient five has a lot of ECG artefacts.

The sleep staging file that came with this sleep record contained many non-integer labels (1.5, 2.5, 3.5) indicating the patient is often alternating between two stages. The staging record for this patient is more arbitrary than records for the other patients.

Patient six's sleep record in Fig. 5.19 contains one long REM phase at four hours. The expected drop in low-frequency power and correlation time combined with the jump in spectral entropy are all present. There is another very short REM period at 2:35 hours. This is interesting as it shows the 0.1–4 Hz low-frequency power goes well below 50%, and also the correlation time drops very low. In Fig. 5.20, we see that the fractional band power for this transition rises the most in the alpha, beta and gamma bands.

Sleep Statistics for Patient four

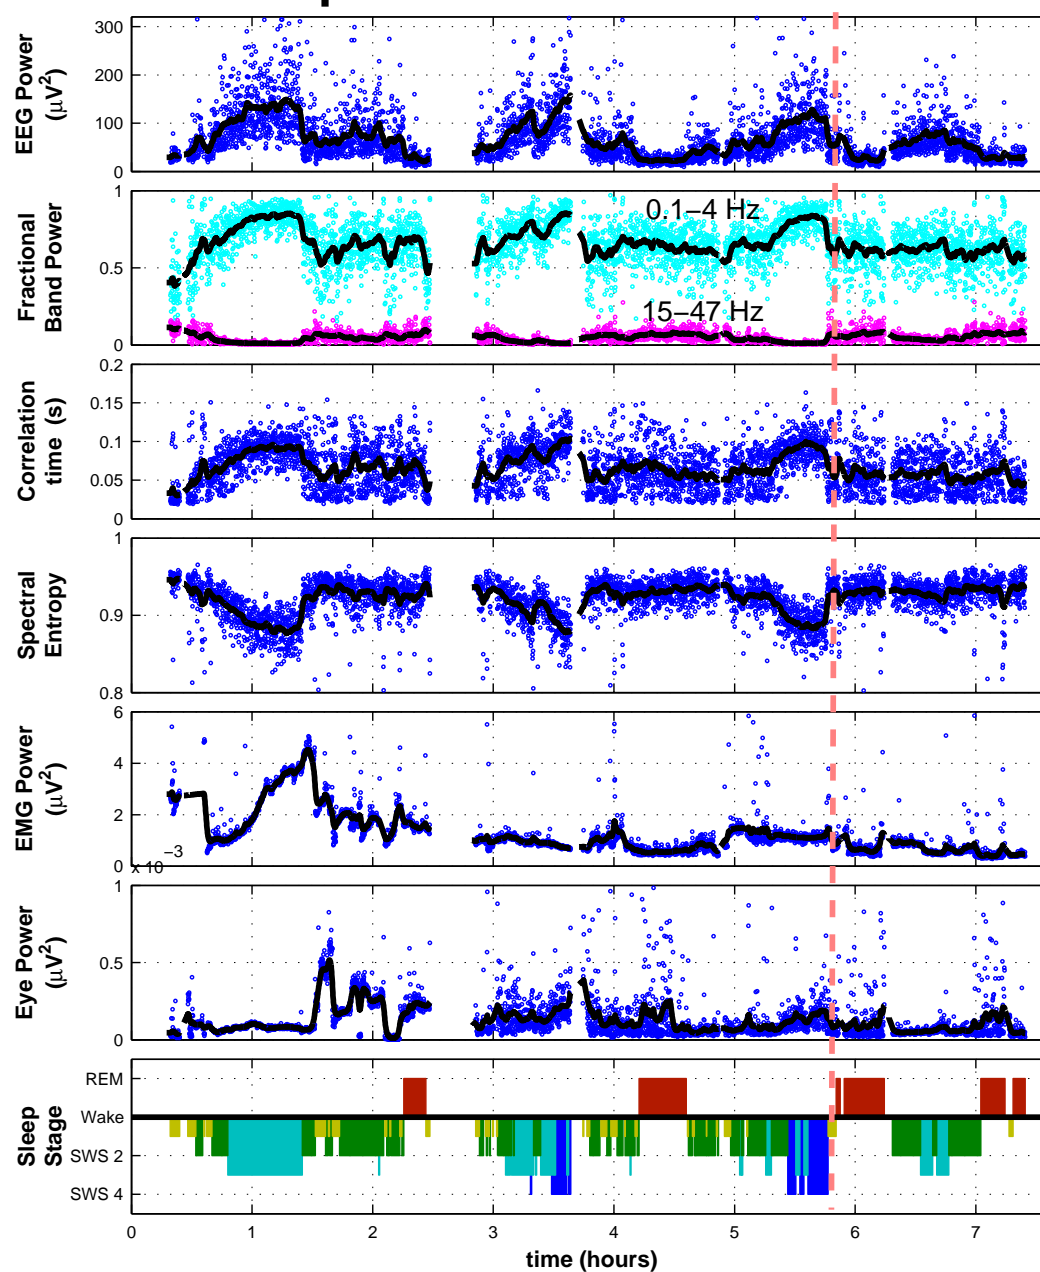


Figure 5.15: The time-series analysis of the sleep recordings of patient four




Figure 5.16: The fractional band power of patient four calculated for five frequency bands.

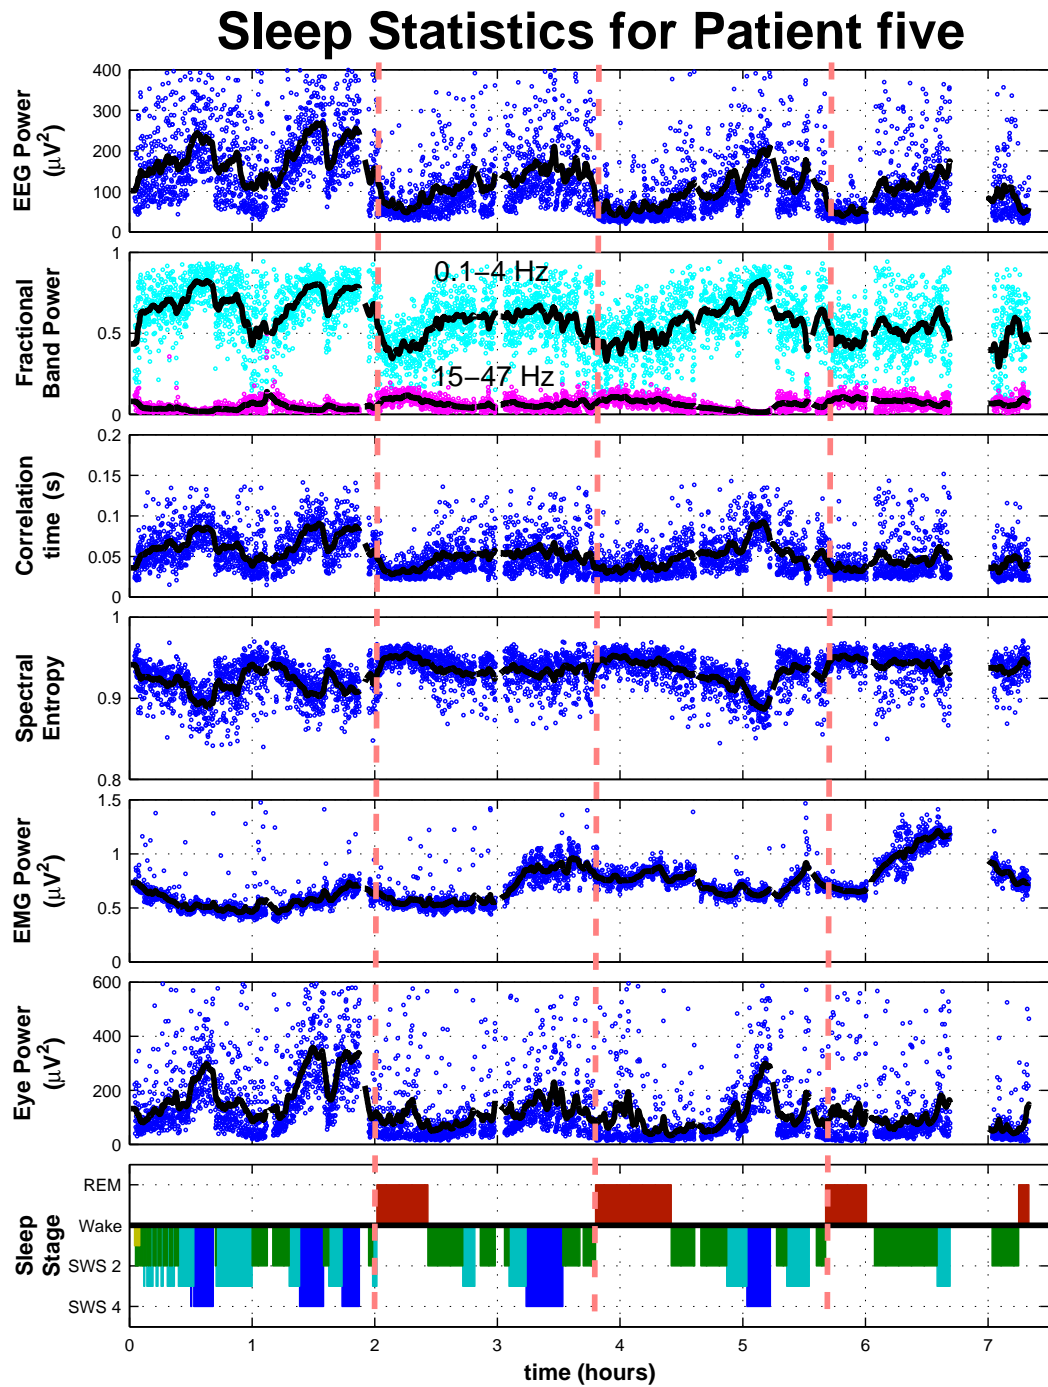


Figure 5.17: The time-series analysis of the sleep recordings of patient five




Figure 5.18: The fractional band power of patient five calculated for five frequency bands.

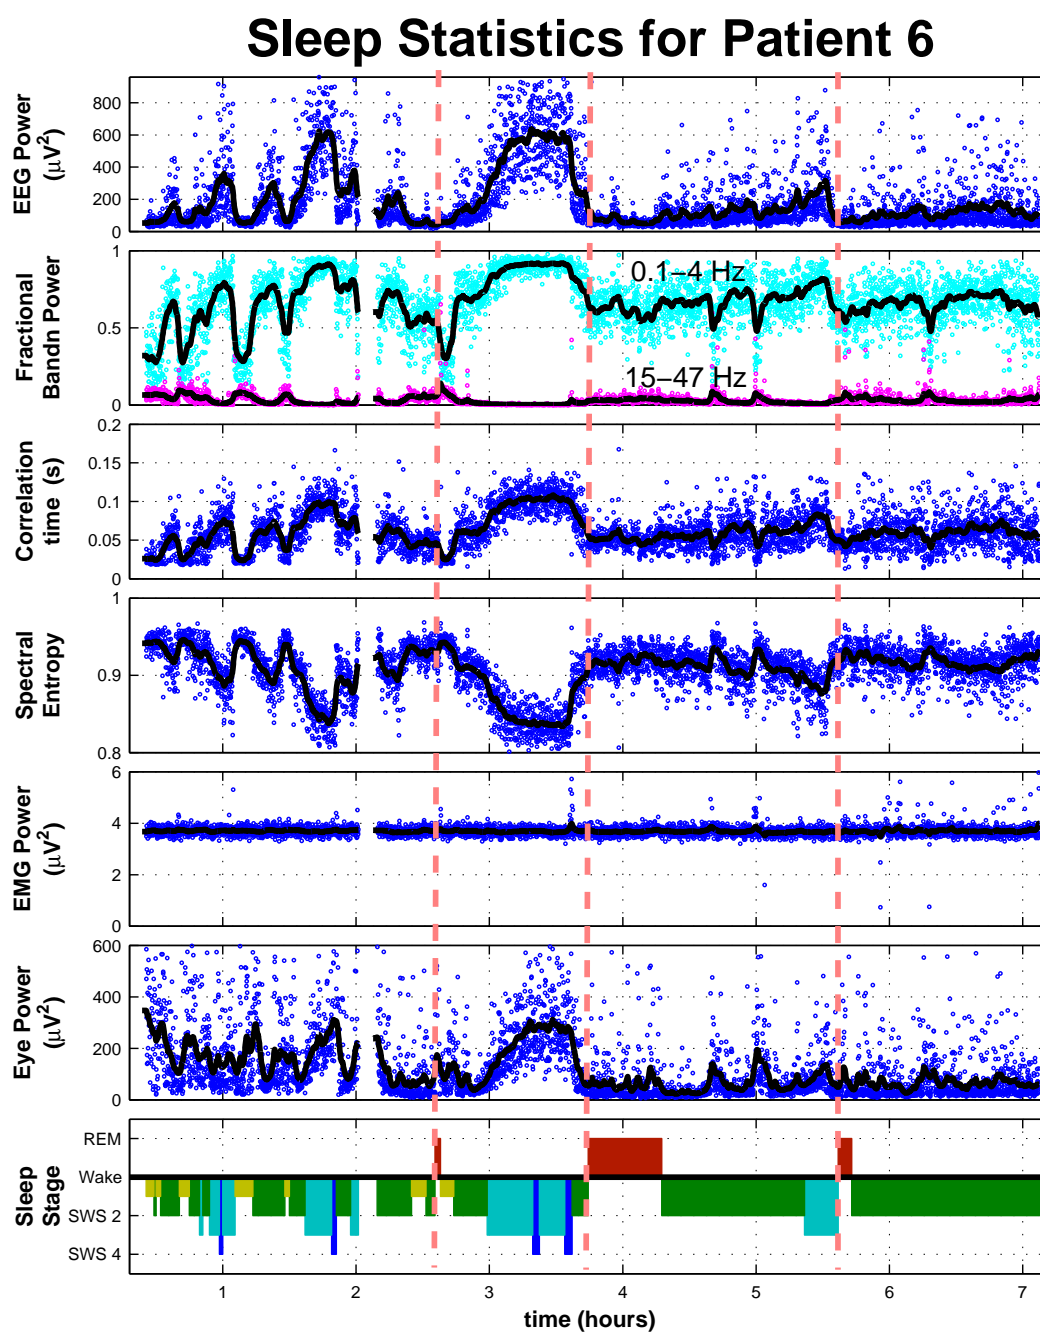


Figure 5.19: The time-series analysis of the sleep recordings of patient six




Figure 5.20: The fractional band power of patient six calculated for five frequency bands.

5.4 How Does ECG Affect the Results?

Several of the patients had ECG artefacts in their EEG time-series. The ECG is difficult to remove without losing information from the EEG, so I did not attempt to remove it. Eye-blink and movement artefacts usually appear as sharp spikes in the calculated statistics, making them easy to identify. ECG artefacts are different. They are a constant periodic spiking signal that is usually most noticeable when the sleeper has low-amplitude EEG. One quick prediction is that the total power will increase, but it is difficult to say what would happen to the other statistics.

To investigate how much the sleep statistics are affected by ECG, I took the sleep record of patient one, whose EEG was free of ECG corruption, and deliberately added his ECG channel to the EEG. The ECG channel records the signal with much greater amplitude than is seen in the EEG, so ECG was scaled by $\frac{1}{25}$. This is shown in Fig. 5.21. The sleep statistics were then calculated for the new time-series.




Figure 5.21: The EEG was corrupted with ECG to see how the sleep statistics are changed when ECG is present.

Figure 5.22 shows the time-series statistics for both the clean EEG and the EEG that was corrupted with ECG. The ECG changes the time-series statistics in different ways. The total power is simply displaced upwards to a higher level. This is because the ECG is so regular and at a constant amplitude. In the three fractional band power graphs (labelled F B P) the same trends are followed, but there is more high-frequency energy in the signal. The 15–47 Hz band has lost much of its fine detail. It seems to miss out many of the sharp peaks. The correlation time has been

shortened and slightly smoothed. The spectral entropy has been affected the most by the ECG. It has been increased, but unevenly. Most of the features in the graph have gone except the high-amplitude SWS centred at three hours.




Figure 5.22: The sleep statistics are calculated again for a portion of patient one's EEG. A new time-series was made by adding ECG to the EEG to compare how the artefacts in the signal change the statistics.

fix heading

From these results, we see that the overall trends of the sleep cycle and the SWS to REM transition are still observed. The total power and the low-frequency frac-

tional power are still usable, but, some of the other time-series statistics, especially spectral entropy, are changed to the extent that quantitative measurements would be questionable. Only qualitative statements could be made about results produced from EEG with these types of artefacts.

5.5 Summary of the Human Results

The results of this chapter show a discontinuous change in the time-series statistics calculated from a sleeping subject's EEG. When the cortex is in the SWS state, the power of the EEG is high and concentrated at low frequencies (below 4 Hz). The correlation time of the EEG is relatively long and the spectral entropy is low. When the transition into the REM state occurs there is a sudden drop in the power and correlation time. The spectral distribution of the EEG power also becomes flatter, which is apparent by the increase in spectral entropy, the drop in relative low-frequency power and the rise in high-frequency power. These abrupt changes in the time-series statistics give strong evidence verifying that the SWS to REM transition is indeed a first-order phase transition, as predicted by the Steyn-Ross *et al.* (2005) model.

From the four different time-series statistics that were computed from the patients' EEG, the total power and spectral entropy were found to be the fastest. Both were useful for displaying the sleep cycles and detecting the SWS to REM transition. The fractional band power was often even better at distinguishing the SWS and REM states. The correlation time also showed the sleep cycles well, but its computation time is much longer than for the other statistics. The correlation time (by accident) detected spindles, a defining characteristic of stage 2 sleep. Epochs containing spindles had a shorter-than-usual correlation time and stood out from the rest of the data.

It was evident to me that the transition from SWS to REM was best defined on people who slept well throughout the recording. All of the raw data came from people who were diagnosed with sleep apnoea and were using the CPAP respirators manufactured by Fisher & Paykel Healthcare. They were being successfully treated, but most of them still woke several times during the recording. When people wake often in the night, as is the case with sleep apnoea, their sleep cycle is disturbed and they may not enter the REM phase of sleep at all. In order to analyse more well-defined SWS to REM transitions to further verify the model, we should record the EEG from people who do not suffer from sleep dysfunctions.

Chapter 6

Rat Data Results

After some encouraging results with the human sleep recordings, Alistair Steyn-Ross and Jamie Sleigh decided to take recordings of laboratory rats. Then I could apply a similar analysis and determine whether the SWS to REM transition of a rat also showed evidence of a phase transition. High-quality amplifiers, able to record down to very low frequencies, were purchased. This was necessary to verify the theoretical prediction of a positive DC-shift in the potential of the excitatory neurons V_e at the point of transition from SWS to REM (see section 2.4.3 on Pg. 24).

Logan Voss, a colleague of Jamie Sleigh, did the hands-on work with the lab rats. To record bio-signals from the rats, Logan surgically implanted the electrodes directly into the rats' cortices. Each set of electrodes was fixed in place by a small headset screwed into the skull. A photograph of one of the rats wearing a headset is shown in Fig. 6.1.



Figure 6.1: A photograph of one of the laboratory rats wearing its surgically attached electrode headset. [image source: Logan Voss (personal communication)]

A bio-signal recorded from cortical electrodes is not called EEG. It is called the electrocorticogram, or ECoG. When recording voltages from the scalp (EEG), the skull and skin act as low-pass filters. So when recording the ECoG, much more high-frequency energy is expected.

6.1 Full Rat Time-Series

When I received the rat time-series data, I tried an analysis similar to the human data analysis, to plot all the sleep statistics as functions of time. There was not a channel reserved for recording the EOG, but it was not required. With close observation, REM can be identified by twitches in the rat's whiskers. A video camera was synchronised with the recording equipment to assist with scoring the stages of sleep. Logan Voss scored the rats' sleep into three stages: waking, SWS and REM.

When humans sleep, we have a sleep-cycle of approximately 90 minutes (Rechtschaffen and Kale, 1968). In general, smaller mammals have a shorter sleep-cycle. Our rats were found to have a cycle of only about 5–10 minutes, and after each cycle they would often wake up and become active. When the rats wake up and move around, their ECoG signals become overrun with EMG artefacts.

Many of the recordings contained too much noise or too many artefacts to perform a full time-series analysis. The three best (i.e., least noisy and most artefact-free) recordings came from two rats. One recording came from rat 6 and two recordings came from rat 12. Their data were clear of artefacts and the rats were generally healthy. I have used a naming convention for the form: $rat\#_{recording\#}$. The three recordings used are: 6_3 , 12_2 and 12_3 .

The statistics graphed in the following figures were computed using the time-series statistics described in chapter 4. They all had the same input parameters. The epoch size T was 4 seconds and the epoch overlap fraction was set to 50%. The sample rate of the raw data was 400 Hz. The time-series statistics have been smoothed using a Whittaker smoother (section A.3.2) with difference order $d = 1$ and smoothing strength $\lambda = 500$.

Figure 6.2 is the time-series analysis for the full recording of rat 6_3 . The first observation is the speed of the rat's sleep cycles. The quick cycling makes it difficult to detect anything of interest in the power except the large-amplitude noise during the long wakeful periods (shown in lighter shade). The fractional band power efficiently distinguishes REM from SWS and wakefulness. In the human results of chapter 5, the energy during SWS was concentrated below 4 Hz. With the rats, however, I have found that the 4–15 Hz energy is higher during SWS than it is during REM

sleep. To get the fractional band power to show the distinctive anti-phase pattern, I used only two frequency bands, 0.1–15 Hz and 15–47 Hz. The correlation time and spectral entropy both show the predicted changes, but the range is very small. To see the contrasts properly, the y-axis limits of both plots have to be greatly reduced. When in SWS the total power is high and concentrated above 15 Hz, the correlation time is high and the spectral entropy is low. When REM sleep begins the power drops away and is distributed into higher frequencies, the correlation time becomes shorter and the spectral entropy increases.

The DC ECoG contains a lot of amplifier drift. When combined with the movement artefacts, it makes identifying any positive jumps in potential difficult. If there are any of the predicted DC-shifts, they are too small to be seen at the resolution of this graph.

The next two figures are separate recordings, 12₂ and 12₃, of the same rat. Fig. 6.3 and Fig. 6.4 contain time-series statistics that have similar trends to those seen in Fig. 6.2. The results in Fig. 6.3 are cleaner and the SWS to REM transitions can be identified easily. Again, the DC ECoG does not contain any clear, positive shifts.




Figure 6.2: The time-series analysis of rat recording 6₃. The top four subplots show the results of the signal processing statistics described in chapter 4: power, fractional band power, correlation time and spectral entropy. The next three subplots are the time series of the three raw data channels: DC ECoG, AC ECoG and EMG. The last subplot is the sleep stage. The portions of the sleep statistics that are shaded lighter are movement artefacts from when the animal is awake. The vertical dashed lines mark some of the SWS to REM transitions.




Figure 6.3: The time-series analysis of rat recording 12₂. The spectral entropy and fractional band power are the most effective means of detecting the REM state.




Figure 6.4: The time-series analysis of rat recording 12₃. This time-series is noisier than the other two. There are some changes in the DC channel, but they seem to occur when the rat wakes and moves around.

6.2 Time-Aligned Transitions

The three previous figures of the full rat time-series all show the basic changes in the ECoG statistics that are consistent with the human results of chapter 5. However, the rats sleep-cycle so quickly and so much data is plotted, that it is hard to look at one specific transition and determine exactly when it occurred. We were unable to find any DC-shifts that corresponded with a transition from SWS to REM. Jamie Sleigh and Alistair Steyn-Ross thought that it would be helpful to take many data subsets containing SWS to REM transitions from the recordings and align them in time so the point of transition for all subsets occurred simultaneously. Then averages could be computed across all of the transitions, and DC-shifts might be easier to detect.

6.2.1 Locating Transitions

Logan Voss, who did the hands-on work with the rats, went through the raw data and found 88 clean SWS to REM transitions (and many more not so clean). For each one, he saved two minutes of both AC and DC ECoG, with the transition approximately at the centre.

My first task was to complete the time-aligning to mark the point of transition as accurately as possible. As it can be difficult to pinpoint the transition using only the ECoG time-series, the time-series statistics were calculated as they show a good distinction between the two states. A program was written to calculate the sleep statistics and graph them along with the raw AC and DC ECoG. The two minutes of data were originally plotted from $t = [-60, 60]$, with the SWS to REM transition approximately in the centre. The time-series were slid up or down the time axis to set the point of transition to $t = 0$. I wrote a function named `ratrem_zero` to load, calculate and plot data. The program had several user input features to help with finding and setting the transition point. It supported the following features:

- Zoom capability in the x or y direction.
- Ability to change the frequency band displayed.
- Ability to remove slope and y -intercept from raw DC ECoG.
- Point of transition set to $t = 0$.
- Ability to change to another transition.

After the transition had been correctly identified, the program would save the index of the closest point to $t = 0$, so that the data could be used in another program to calculate averages.

I found that certain statistics were more helpful than others in determining the transition point. The power, spectral entropy, and spindle and gamma fractional power bands showed the best contrast in the SWS and REM states. The correlation time was not calculated because it computes much slower than the other statistics. Figure 6.5 is a screenshot of the display of `ratrem_zero`. The data was calculated with an epoch of 2 s and overlap of 90%. At $t = 0$, there is a clear SWS to REM transition.




Figure 6.5: A screenshot of `ratrem_zero` displaying one of the transitions. The time-series statistics were computed to assist with finding the position of the SWS to REM transition. This was done for all 88 transitions. The total power fluctuates high and low while in SWS and then drops away when REM begins. The spindle band is high in SWS and then is much lower in REM. The gamma band is opposite. It is nearly non-existent in SWS, then becomes high during REM. The other frequency bands did not show a consistent change at the transition. The spectral entropy is lower in SWS and jumps to a higher level at the transition. Notice that the time-series starts before -60 s and ends before $+60$ s. It has been slightly shifted in time so that the transition occurs at $t = 0$ s.

6.2.2 Differences in ECoG Amplitude

When displaying the statistics, I noticed that the amplitudes of the ECoG time-series varied from one recording to another. Figure 6.6 (a) is a plot of the rms voltage amplitude of the AC ECoG versus the transition number (1–88). The amplitudes of the transitions are consistent within a single recording, but there is a great difference

between recordings. The main explanation for the different rms voltage levels was that different amplifier gain settings were used. The recordings 6_1 , 6_3 , 7_1 and 10_1 whose rms voltages are all at a similar level, had an amplifier gain of 100. The other recordings all had a gain of 500, but there is much more variance in the recordings. Figure 6.6 (b) is a plot of the rms ECoG voltage, with the amplifier gain removed.




Figure 6.6: The rms voltage levels of the ECoG of the 88 transitions. Transitions belonging to the same recording have been plotted with the same marker (e.g., \circ , \times , $+$, ∇ , etc). (a) Early recordings had the amplifier gain set to 100, while the later recordings had a gain of 500. This was a major cause of differences in the ECoG amplitudes. (b) The gain is removed, but there are still differences between the recordings. The outlier in recording 6_1 is from a large artefact at the very end of the two-minute data set.

Even with the gain removed from the signal, there is still quite a lot of difference between the rms voltages. Even the three recordings of rat 12 all have different rms voltages. Perhaps these differences could be from the electrodes moving from their original position. Before each recording, the electrode wires would have to be connected to the headset. Because the electrodes actually go through the rat's skull and directly into the cortex, a slight change in depth can alter the strength of the recorded signal. Any movement by the rat could change the amplitude of the ECoG.

When combined to calculate averages, the transitions were grouped by their recording number so that their ECoG amplitudes would be all at the same level.

6.2.3 Results of the Time-Aligned Transitions

To display the results of the time-aligned transitions, the raw DC and AC ECoG and the time-series statistics have been plotted. The fractional band power was calculated for the seven frequency bands previously listed in table 5.1, but the slow-wave (0.1–1 Hz) was always very close to zero so it was not plotted. The results from rat recording 6₁ are shown in Fig. 6.7. In each subplot, the main point of focus is the average, which is superimposed on the top and rendered as a thick, dark line. The data in the background are plotted to show the amount of fluctuation between transitions.

In Fig. 6.7, the expected drop in power and rise in spectral entropy at the SWS to REM transition are present. The correlation time also becomes lower at the transition, but it is more vulnerable to artefacts, so there are many large spikes. The spikes in the correlation time make the SWS and REM states not as distinguishable as they are in the power and spectral entropy. The fractional band power is interesting. The only frequency bands to show a clear change at the transition are the 12–15 Hz spindle band and the 25–47 Hz gamma band. Most of the other, lower-frequency bands slightly decrease, but it does not happen suddenly as with the spindle and gamma bands.

The graphs of the other rat recordings are all very similar to Fig. 6.7, and provide little extra information. They are included in Appendix B. All the figures are consistent in that the power, spectral entropy and the spindle and gamma bands of the fractional power most clearly show a sudden change at the transition from SWS to REM.

The results of the time-aligned transitions show that the SWS to REM transition of a rat is very quick. The abrupt change in the time-series statistics is strong evidence that the cortex undergoes a first-order phase transition as it switches from SWS to REM, as predicted by the Steyn-Ross model.

6.3 Looking for the DC Shifts

As mentioned at the beginning of this chapter, the Steyn-Ross model predicts a positive DC-shift in the EEG/ECoG as the cortex undergoes the phase transition from SWS into REM. One of the goals of recording from the rats was to observe and measure this DC-shift. With the time-aligned transitions from the previous section, the corresponding DC ECoG could also be averaged to see if there was any change in the potential.

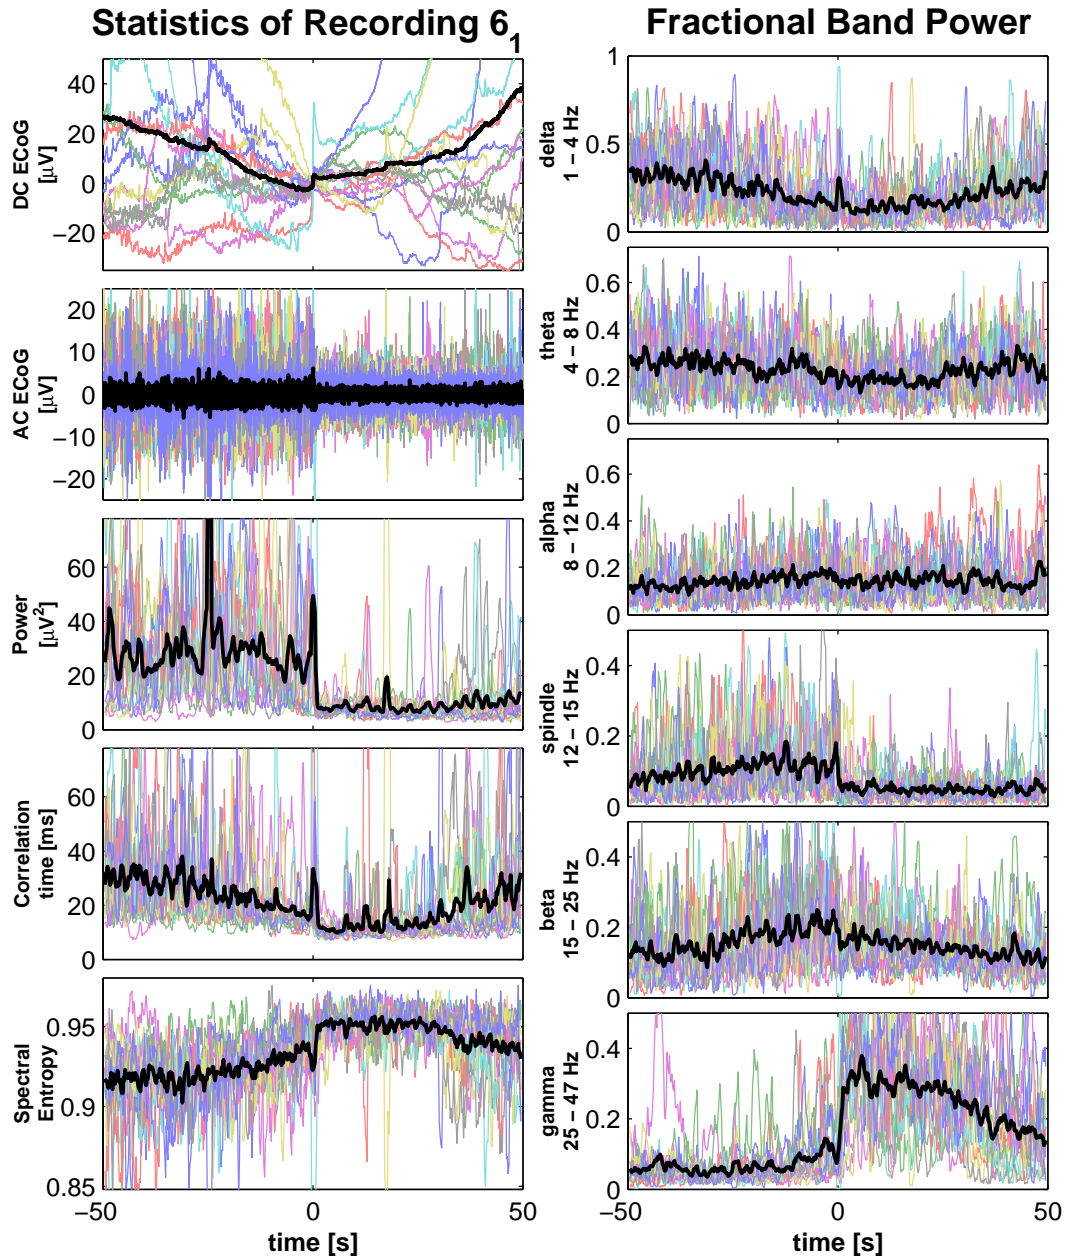


Figure 6.7: The 15 time-aligned SWS to REM transitions for rat recording 6₁. The dark line is the mean of the 15 data sets. The SWS to REM transition is centred at $t = 0$ where the AC ECoG and time-series statistics change abruptly. The power, spectral entropy and the spindle and gamma frequency bands give the best evidence of a change in state. In the DC channel, a small positive DC-shift is present in the mean. The DC ECoG analysis follows in sec 6.3.

6.3.1 Removing the DC Drift

Time-aligning the transitions alone was not enough to prepare the DC ECoG for averaging. The drift in the DC channel made comparisons between the transitions difficult as they were all at different voltage levels. To make the DC transitions comparable, each DC ECoG time-series needed to be modified. First, a subset of the time-series that contained drift was selected. The slope of that subset was then computed via linear regression. The slope and y -intercept of the fitted line were then subtracted from the DC ECoG time-series leaving a new time-series that was close to zero, but still contained any DC-shifts that were present. An example of this procedure is shown in Fig. 6.8 with the DC ECoG from the 73rd transition. The program `ratrem_zero` had an option to allow the user to select the data subset; then it would calculate the straight line and subtract it from the time-series. The coefficients of the straight line were later saved to a data file.




Figure 6.8: A demonstration of how the drift was removed from the DC ECoG. (a) A linear portion of the time-series is selected. (b) A linear, least-squares line is computed from the data points inside the selected region. The line is extrapolated to be the same length as the time-series. (c) The line is subtracted from the time-series. After the drift is removed, there is a slight increase in DC potential at transition, but it occurs slowly over ~ 10 s, not as fast as predicted by our model.

Figure 6.9 shows how modifying the DC data helps. The DC ECoG time-series of the SWS to REM transitions of rat recording 6.1 are plotted in Fig. 6.9 a. Each transition has a different average DC voltage. By bringing them back towards zero, they can be compared and averaged. In Fig. 6.9 b the DC time-series have been modified. They now all fit into a smaller range and an average can be easily calcu-

lated.




Figure 6.9: (a) The DC ECoG of 15 transitions of recording 6.1 are all at different voltage levels because of the drift in the DC channel. (b) The slopes and y-intercepts are removed to aid in comparison. Several of the transitions have positive DC-shifts at $t = 0$, but only one can be easily seen here. Subplot (b) was plotted earlier in Fig. 6.7 where the mean was also calculated.

By altering the DC ECoG in this way, the final and initial potentials of any DC-shift were ignored. At this point in the research, however, we were more interested in finding whether there was any DC-shift associated with the SWS–REM transition.

6.3.2 How Many Transitions Show a DC Shift?

Most of the transitions in Fig. 6.9 do not have a DC-shift. To find those that do have a jump in the potential, the time-series from recording 6₁ have been plotted along with their first derivative in Fig. 6.10. If a DC-shift is present, the derivative should show an upward spike. Because the DC-shift is predicted to occur quickly, only 10 s either side of zero are shown. The data has been smoothed to remove the AC components of the ECoG. Smoothing is applied because if the numerical derivative is computed on the raw DC ECoG, the fluctuations are amplified, which hides the spike that is to be detected. Unfortunately, the smoothing also smooths out the sharp DC shifts, making the spike in the derivative lower than it should be. The parameters used for the Whittaker smoother were $d = 3$ and $\lambda = 5 \times 10^7$.

Five of the 15 time-series shown here have clear DC-shifts. They are all of different amplitudes, the greatest of which is number 11. The change in potential at the SWS to REM transition in number 11 is very large and it heavily influences the mean, which was shown previously in Fig. 6.7. There are many other DC-shifts that have not been marked because they do not occur at $t = 0$. There are also other events that are noticeable in the DC time-series that do occur at $t = 0$, but are not well distinguished DC-shifts. One of these other events is a “down-up-down” kind of waveform in number 8. Another is seen in number 7 where a quick, small jump at the transition follows large, slower change in potential. It is hard to tell whether these events are related to the SWS to REM transition.

The problem with detecting the DC-shifts by calculating the first derivative is that each jump in potential is found manually. By eye, it is often easier to see the DC-shift directly in the ECoG, rather than in the derivative. When the detection is performed manually, there will always be biases. If a detection algorithm could be implemented, ideally it would not be subject to human biases. The noise in the derivative, coming especially from the large SWS fluctuations, makes automatic software detection difficult. With some early help, I have developed an algorithm for detecting DC-shifts. This is the subject of chapter 7.

6.3.3 Should Transition 11 be Included?

As mentioned earlier, the DC-shift in transition number 11 is much larger than the others. It is so large that it is almost too good to be true. The means across all of the DC ECoG transitions, and their derivatives for numbers 1–15 have been calculated both with and without transition number 11. In Fig. 6.11, which includes number 11, there is a large, quick jump in the potential, which gives a tall spike in the




Figure 6.10: The first 15 (smoothed) DC ECoG transitions and their first derivatives. Circles mark the spikes in the derivatives that clearly indicate that there is a DC-shift at $t = 0$.

derivative. Number 11 is excluded from the average in Fig. 6.12 and the positive DC-shift is still present, but it is just not as large, and goes over a longer period of time. In the derivative, there are two smaller peaks at $t = 0$, but they do not stand out well against the rest of the fluctuations.

The 11th transition has such a large effect on the means that one questions its validity. Using the MATLAB command `findstr`¹ the position of the 11th transition was located in its original recording. As seen in Fig. 6.13, this particular DC recording has some large disturbances which change the DC potential by about 5 volts. It is highly unlikely that the rat's cortex could change by 5 V. There must have been some electrical interference to cause such a huge change.

The region where transition 11 occurs looks to be on the decay curve of the amplifier as it recovers from the huge shift of ~ 5 V. In the two minutes containing the DC-shift, there is not much fluctuation in the time-series apart from the DC-shift itself, and another smaller one at 190:30 minutes. If the DC-shift is really an electrical artefact, and not caused by the cortex changing state, then it is an unfortunate coincidence that it happens simultaneously with the SWS to REM transition.

Because the change in DC potential is so large, and the signal seems to be influenced by electrical interference, transition 11 should be excluded from the analysis. It is encouraging that when 11 is excluded from the mean, we still see the rise in the DC potential. At the moment, I can say that sometimes there is a positive DC-shift as the cortex makes the transition from SWS to REM. Still to be answered is the question of why some transitions do not contain any sign of a DC-shift.

To get more conclusive evidence, we need to make more recordings and take better precautions to stop electrical and movement artefacts. A Faraday cage would help to remove any electrical interference, but stopping movement artefacts would be a more difficult matter. In a study by Marshal *et al.* (1998), she used a clamp attached to the bed to restrain the heads of her human volunteers to stop them from moving during their sleep. This eliminated all the movement artefacts which, can drastically change the DC potential. In contrast, our rats were free to move around in a container during the recording process.

¹The function `findstr` is designed to find the indices of small strings hidden inside larger strings, but it also works for numerical arrays provided they are arranged as row-vectors.




Figure 6.11: The means of the DC ECoG and their derivatives for transitions 1–15. When transition number 11 is included in the mean, the DC-shift is very clear in both plots.




Figure 6.12: The means of the DC ECoG and their derivatives but excluding transition 11. The change in potential is still present in the DC ECoG, but it is not very distinguishable in the derivative, as the shifts do not all happen simultaneously, which causes the shift in the mean to occur more slowly than in Fig. 6.11.

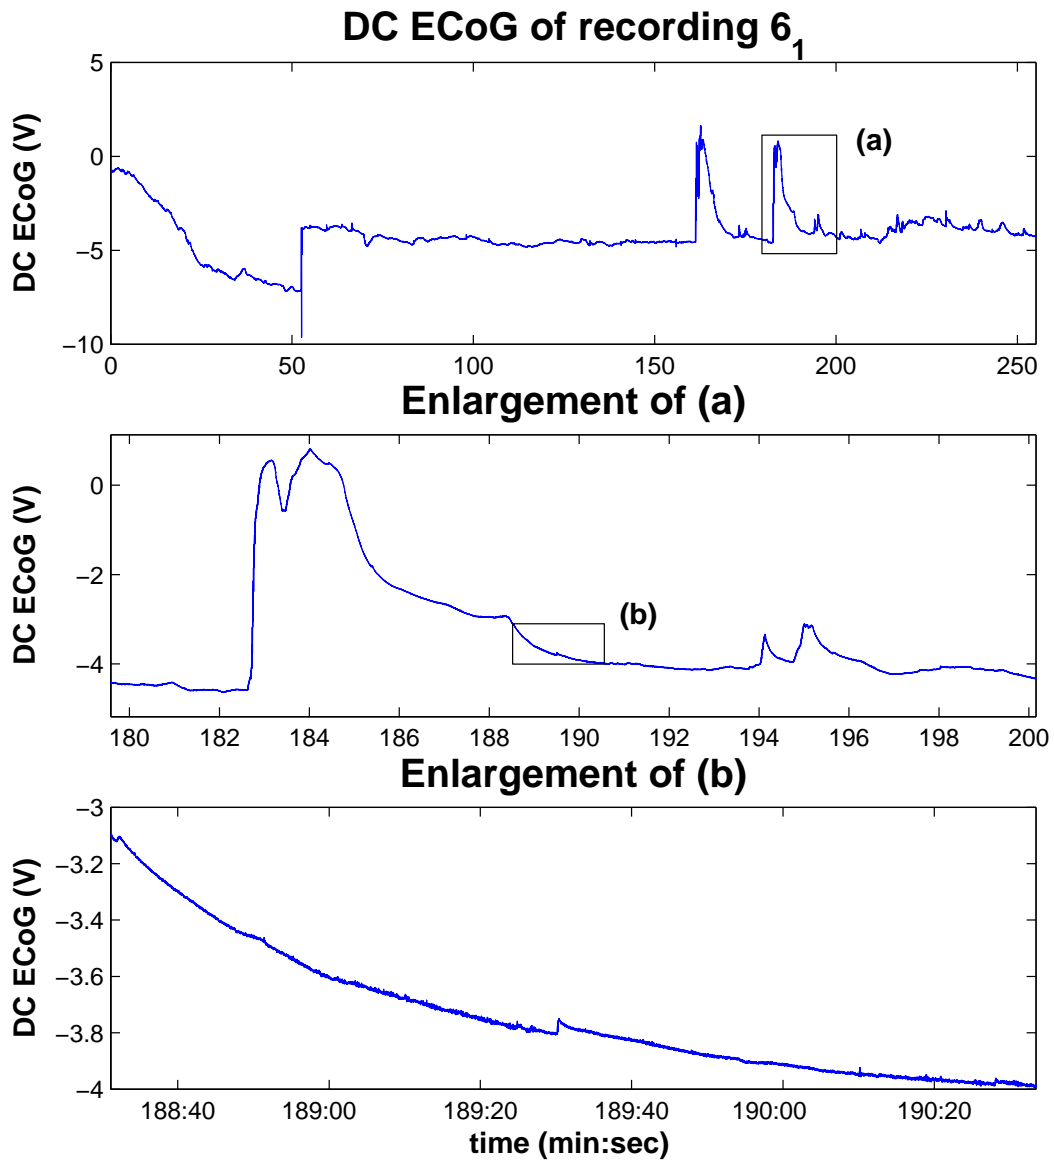


Figure 6.13: The location of the 11th transition comes from a region in recording 6₁ where there seem to be large electrical fluctuations.

6.4 A K-complex or a DC-shift?

While examining the rat transitions, particularly number 11, I made an interesting observation. Whenever there is a sharp change in the DC potential, from either the cortex or an electrical artefact, the corresponding AC ECoG always has a waveform that looks very similar to a K-complex. I believe that these waveforms are the result of high-pass filtering of the true ECoG signal that contains all frequencies.

All amplifiers used for data acquisition have a built-in high-pass filter (unless specifically measuring DC). When filtering pre-recorded, digitised data on a computer, one often uses a function such as `filtfilt` found on the MATLAB signal processing toolbox. What is useful about `filtfilt` is that the signal is filtered twice, once forward and once backward. This eliminates any phase shift that would come about from filtering in only one direction. During the data acquisition on the other hand, the amplifier has to apply its filter in real time, and therefore cannot filter in the backward direction.

In Fig. 6.14 there is a comparison between a real K-complex taken from human EEG, the apparent K-complex in the ECoG of transition number 11, and a K-complex type waveform that was simulated by filtering noise containing an artificial step. The shape and time-scale of all three are very similar. The filter parameters needed some optimisation to get the simulated waveform to match the shape of the other two. An elliptical filter was used with a cutoff frequency of 0.1 Hz and high-pass frequency of 1 Hz. The amount of ripple allowed in the pass band was 0.5 dB. The sharpness of this filtering required a filter order of 2. The magnitude and phase response of the filter used are plotted on Fig. 6.15.

From the results of Fig. 6.14, it is quite clear that filtering a sharp step in the DC potential, with a filter whose magnitude and phase response are similar to those in Fig. 6.15, will result in a waveform that is very similar to the K-complex. The filtered DC-shifts and real K-complexes are so alike that I believe that misclassifications could arise. If the Steyn-Ross model is correct, and there are DC-shifts at the transition from SWS to REM, then a fake K-complex, arising from filtering out the DC, could be mistaken for a real K-complex. Because K-complexes belong to stage 2 sleep, the sleep stage would be assigned to stage 2 instead of REM.

Perhaps all K-complexes are actually some kind of DC-shift, and the shape that is so familiar is really just the impulse response of the amplifier filter.




Figure 6.14: Different waveforms that look similar to a K-complex. (a) A real K-complex from a human sleeper. (b) The DC-shift from the transition 11 of the rat time-aligned transitions. (c) The AC ECoG that was recorded after the amplifier filtered the DC component from the true signal. (d) Artificial data made to look similar to a DC-shift. (e) The result of high-pass filtering of the artificial data in (d) in only the forward direction. The phase response of the filter distorts the sharp DC shift and leaves the waveform that looks very similar to the K-complex.




Figure 6.15: The magnitude and phase response of the high-pass filter that was used to filter out the DC component of the artificial data in Fig. 6.14 d.

6.5 Summary of Rat Results

The changes in the power, fractional band power, correlation time and spectral entropy are consistent with both the human results of chapter 5 and the predictions of the phase transition model. At the transition from SWS to REM the power in the ECoG greatly drops and it is distributed more into the lower frequencies. This is accompanied by a drop in correlation time while the spectral entropy increases.

One difference in the results between the rats and the humans is the frequency distribution. During SWS, there is much more energy in the 12–15 Hz spindle band and much less in the 0.1–4 Hz slow-wave and delta bands for a rat. With a human, almost 100% of the energy is concentrated below 4 Hz.

The quick cycling of the rats makes it difficult to see if there is a slow build-up in power during SWS as there is with a human, but the qualitative changes in the time-series statistics are still consistent with our phase transition model.

By aligning many time-series subsets that all contained SWS to REM transitions, so that the transitions occurred simultaneously, averages were calculated. The

averages very clearly showed that there was a sharp drop in power and a rise in spectral entropy. The correlation time was not as effective at detecting the transition, but the drop was still visible. Two of the six different frequency bands showed an easily observable change at the transition. There was a rise in the fractional power in the 25–47 Hz gamma band and a drop in the 12–15 Hz spindle band. The 15–25 Hz beta band also reduced slightly at the transition. Some of the other fractional power bands, particularly the 4–8 Hz theta band, gave no hint of a change in state. This was surprising as I had previously assumed that all the energy concentrated in the low-frequency bands would be redistributed evenly throughout the higher-frequency bands.

With the DC ECoG, it was easier to look at each transition separately. Five of the 15 SWS to REM transitions showed a clear positive jump in potential which was also seen as a spike in the first derivative. However, the jump was so large for transition number 11 that it really should be excluded from the analysis, as it seems to be an electrical interference artefact with very good timing. In the calculation of the mean of the other 14 transitions, an increase in the potential was still present, but it was too slow to be well distinguished above the noise in the mean of the derivative. To get better, more reliable results, I believe more precautions need to be taken to avoid artefacts when recording down to zero frequency.

One of the observations made while searching for DC-shifts was that the AC recording contained a waveform that looked very similar to a K-complex. These are the result of the high-pass filtering of the true bio-signal in real time, which is done electronically by the amplifier. The sharp DC-shifts are phase-shifted by the filter and produce what looks much like a K-complex. It could be that there are DC-shifts at the transition into REM that appear to be K-complexes. This similarity could lead to a misclassification. According to the sleep staging rules by Rechtschaffen and Kale (1968), this epoch would have to be scored as stage 2.

Overall, the results of this chapter provide good evidence that the Steyn-Ross model is correct in predicting that the transition from SWS into REM can be classified as a first-order phase transition. So far, there is some encouraging evidence of a DC-shift at the transition, but it is not observed consistently enough to give us any certainty.

Chapter 7

DC-Shift Detection Algorithm

In the previous chapter some of the DC ECoG was analysed in the hopes of finding positive jumps in the DC potential at the time of the SWS to REM transition. From 15 transitions, 5 contained DC-shifts that were easily identified by inspection of both the DC ECoG itself, and its numerical derivative. As mentioned earlier, manual identification of the DC-shifts is not ideal. Presented in this chapter is an algorithm that I have developed for detecting the DC-shifts.

7.1 Developing the New Algorithm

At first glance the problem of detecting a DC-shift is similar to step detection that would commonly occur in image processing, such as the ‘weak string’ by Blake and Zisserman (1987). However, this and many other image processing step-detection techniques require that a threshold be set, making them unsuitable for my data as we do not know what the magnitude of the DC step will be, and therefore cannot set a meaningful threshold.

Another possibility, would be the first derivative. If the signal is generally smooth then any sudden step should show up as a large spike in the first derivative. But, because the EEG is full of noise-like fluctuations, the first derivative seems to amplify the fluctuations, making the voltage step difficult to discern. In chapter 6 the first derivative was computed after the DC ECoG had been smoothed. Smoothing the data did remove the fluctuations, but it also smoothed out any DC-shifts, thus reducing height of the spike in the derivative.

7.1.1 The Two Least-Squares Lines

With the help of Jamie Sleight and Alistair Steyn-Ross, we thought of another method. Two linear, least-squares lines were fitted to the data, the first at $[t - \Delta t, t]$ and the second at $[t, t + \Delta t]$. There would be a difference in the vertical position, or DC voltage ΔV , of the two lines at time t . Then t could be varied along the time axis with the lines being re-fitted for every position, giving a function of $\Delta V(t)$. If there is a DC-shift present in the signal, then it would be detected by the maximum value of ΔV . This is illustrated in Fig. 7.1 where the two lines have length $\Delta t = 2$ s and t is varied to different positions along the time axis. Where there is no DC-shift in the ECoG, there is only a small difference between the ends of the straight lines. At $t = 0$, where there is a DC-shift, the two lines have a large separation and ΔV is measured.

In Fig. 7.1, only a few seconds either side of the SWS to REM transition have been plotted. This is because the model predicts that the transition will be very sudden. The data has already been time-aligned, so we know that the SWS to REM transition is at $t = 0$ (or very close to it). By varying t to every time sample between -2 and 2 s and evaluating the two least-squares lines within the region of interest, the function $\Delta V(t)$ is formed. The function corresponding to the example of transition 5 shown in Fig. 7.1, is graphed in Fig. 7.2. The DC ECoG is plotted in the interval $[-4, 4]$ while ΔV only has the range $[-2, 2]$. This is because the first point of the change in potential, $\Delta V(-2)$, is calculated from 4 s of ECoG data, -4 s to 0 s.

7.1.2 Varying Δt to Make a Surface

A clear peak in ΔV shows the position of the DC-shift, and the maximum value gives its size. The value of $\Delta t = 2$ s is somewhat arbitrary. Jamie Sleight originally suggested to try to fit two lines of *about* 2 s each to the DC data to see if there were any noticeable changes in the potential. While developing the algorithm further, I wondered if it would make much difference to vary the length of the least-squares lines. Shorter values of Δt would make the fitted lines become more sensitive to small changes in the DC signal, and would then introduce more fluctuations into $\Delta V(t)$. Conversely, if Δt is increased the signal would be less susceptible to change and ΔV would become smoother. Figure 7.3 shows how ΔV changes when different line lengths are used.

It is hard to say exactly what value of Δt is optimal, but between 0.5 and 2 s seems to work well. Much bigger than a few seconds would smooth out ΔV so




Figure 7.1: Two least-squares lines, each of length $\Delta t = 2$ s, slide along the time axis of the (smoothed) DC ECoG of transition 5. The vertical difference between the two lines gives the change in the DC potential. By looking at the time-series, there is obviously a voltage step at $t = 0$. The two lines do a good job at detecting the step at $t = 0$ in plot (d).




Figure 7.2: The function $\Delta V(t)$ is maximum at the same position as the DC-shift in the ECoG. This is expected as the example shown contained a clear DC-shift.




Figure 7.3: By changing Δt , the length of the least-squares lines, the sensitivity of the DC-shift detection changes. When Δt is large, only the clear DC-shifts are detected. As Δt is lowered, it becomes more sensitive to small fluctuations

much that only very large DC-shifts (such as transition 11) would be detected. On the other hand, lowering Δt down to only a few samples would make ΔV fluctuate so much it would be almost as noisy as the numerical derivative. Because the ideal value of Δt is unknown, it can be made a variable, rather than a constant parameter. The definition of ΔV has now changed, it is a function of two variables and has become a surface: $\Delta V = \Delta V(t, \Delta t)$.

In Fig. 7.4 the surface of $\Delta V(t, \Delta t)$ for transition 5 is plotted. Δt is varied in reasonably large steps from 0.5 s to 2 s. The main feature of the surface is the narrow hill centred at $t = 0$, where the DC-shift has been detected. It is encouraging that the maximum value of the hill is almost constant along the Δt axis. It shows that when there is a clear DC-shift, it is detected no matter the length of the lines. Away from $t = 0$, the surface changes a lot more. When Δt is short, there is more fluctuation and a smaller peak emerges at $t \sim -1.6$ s. At the high end of Δt , either side of the peak, V drops below zero, making wide troughs.

The mean value of the surface ΔV is then calculated across Δt to compress it back into a one variable function of t , $\Delta \bar{V}(t)$. Then the DC-shift and its position can be measured with a simple `max`¹ command. The mean of the surface, $\Delta \bar{V}(t)$, is plotted in Fig. 7.5. The vertical, red line denotes the position and size of the DC-shift.

7.2 Applying the Algorithm to the DC Transitions

This method works very well when it is applied to data with a clear jump in the potential, such as transition 5, which was the test data that was used in generating the last few diagrams. It accurately determines the position of the maximum change in the DC potential. The algorithm seems elaborate considering that taking the derivative also found the position of the DC-shift in only a fraction of the time. But this algorithm does have advantages over the derivative. It also finds the height of the DC-shift and is more resistant to noise and artefacts. To test how well the algorithm works, it needs to be applied to data with less obvious DC-shifts. In the diagrams that follow, Figs. 7.6 and 7.7, the transitions 1–15 (except 5 as it was used previously) from rat recording 6₁ are plotted.

As seen in the Figs. 7.6 and 7.7, wherever there is a true DC-shift, it is accurately detected. In a few cases (transition 2, 7 and 12), the algorithm marked DC-shifts occurring close the SWS to REM transition at $t = 0$ that could not be clearly marked

¹The MATLAB command `max` can return the maximum value and its index:
`[mx, i] = max(data);`




Figure 7.4: A plot of the surface $\Delta V(t, \Delta t)$. t is the time axis, and Δt is the length of the least-squares lines used to measure the DC-shift. The maximum point of the peak that runs along $t = 0$ is the DC-shift. For this transition, varying the length of the lines does not change the position of the DC shift.




Figure 7.5: The surface $\Delta V((t, \Delta t))$ has been averaged over Δt , to form a one variable function, $\Delta \bar{V}(t)$. The red vertical line is the DC-shift whose position and length were calculated by finding the maximum of $\Delta \bar{V}(t)$. The line has been copied to the second subplot of the DC ECoG to show how well the algorithm detects the DC-shift.

when using the first derivative. The idea of using the two least-squares lines to find the shift definitely works when there is a shift. But the problem is that when there really is no DC-shift, the max command forces one to be found. Some kind of test needs to be included so that data without a shift gets marked as having no shift.

More logic needs to be programmed into the algorithm to exclude shifts that are not actually there. Such as in transition 11, for example, the positive shift is very clear, but the algorithm also detects a lower amplitude negative shift after the positive one. This seems to happen wherever there is a clear, positive (or negative) spike in ΔV . Preceding and following the spike there are always smaller amplitude local minima (or maxima), which are not at all associated with their own DC-shift. Even more logic needs to be included into the algorithm so that other real shifts that, show up as local maxima in $\Delta \bar{V}$, are correctly marked as DC-shifts. After these have been implemented, it would be simple to include detection for negative shifts too. The simple solution would be to set a threshold that a shift must overcome before it is detected. However, there is a large range in the magnitudes of the good DC-shifts seen in Figs. 7.6 and 7.7.

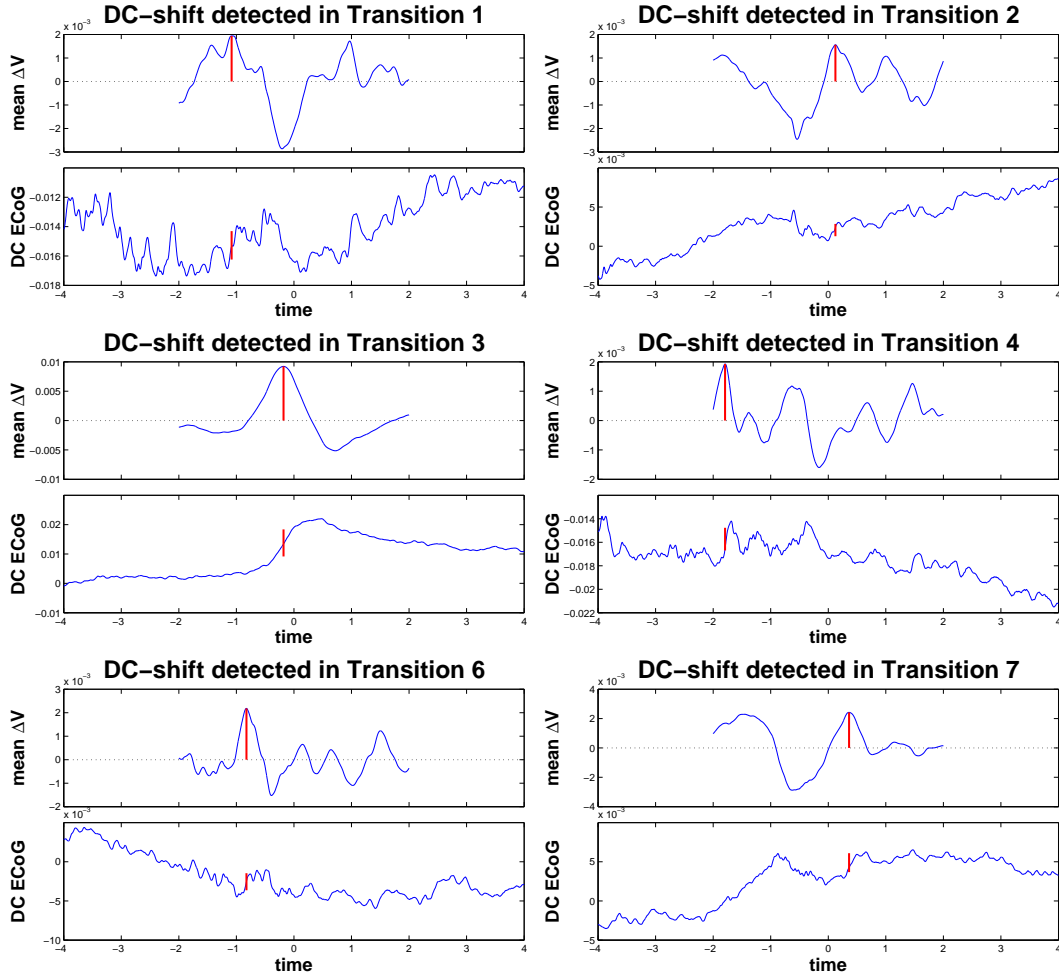


Figure 7.6: Because the algorithm finds the maximum value of $\Delta \bar{V}$, only one DC-shift is detected in each data set. Transitions 2, 3, 6 and 7 have positive DC-shifts that were detected within 1 s of $t = 0$. Number 3 has a clear DC-shift, but its slow rise makes the predicted value lower than it should be. The other transitions, 1 and 4, both have other peaks showing smaller positive shifts. They also contain large negative shifts (the deep troughs in $\Delta \bar{V}$) that are closer to $t = 0$.

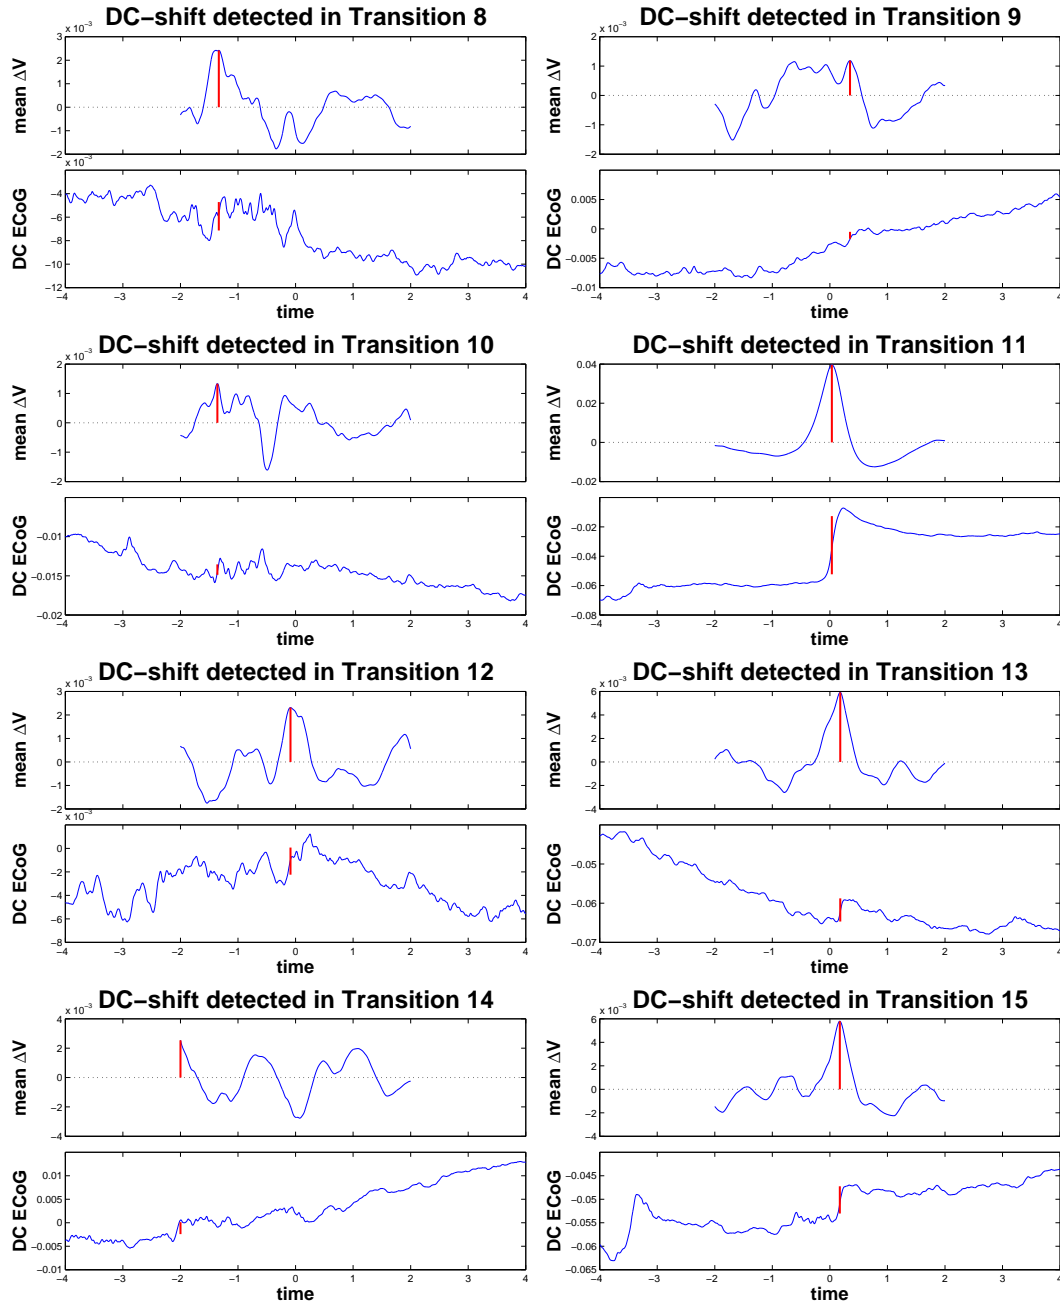


Figure 7.7: Three very clean DC-shifts are detected in transitions 11, 13 and 15. Although earlier it was decided that number 11 is probably an electrical artefact and should be excluded, it is still a good step which is accurately detected by the algorithm. Transition 12 is also very close to $t = 0$. The overall slope of the DC ECoG seems to change at the transition at $t = 0$, and seems to boost the amount of the DC-shift. In transition 9 a shift has been identified that is very small. Here the DC-shift occurs over a long period of about 1.5 s. The at the point in the DC ECoG that is marked as the shift, there is a small ‘down-up’ which gives it the maximum value. Transitions 8, 10 and 14 show large negative shifts close to $t = 0$. The positive shifts that are marked, are quite far away from the transition point.

7.3 How Well do the Lines Fit the Data?

One way of further increasing the sensitivity of the algorithm that could be used to eliminate some of the questionable or arbitrary maxima in ΔV , is to include a measure of “goodness of fit” of the least-squares lines into the algorithm. When there is a clear shift in potential, the two lines fit the data points best when either both are on one side of the shift or one line is before the shift, and the other is after. If a straight line is fitted to the corner of a step, the fit will be poor, as the residuals will be large. In Fig. 7.8 transition 5 is plotted again with the two least-squares lines calculated at three positions, $t = -0.5, 0$ and 0.5 s. The norm^2 of the residuals have been calculated for each of the straight lines. A low norm value indicates a better fit. For each pair of lines, the average of the two norms has been divided by the square-root of the number of elements in the line. For example if the number of elements is n , then

$$R_{\text{norm}} = \frac{R_{\text{norm},1} + R_{\text{norm},2}}{2\sqrt{n}} \quad (7.1)$$

The MATLAB function `polyfit` is used to calculate the least-squares lines. It can return the norm of the residuals as an optional output, so to compute R_{norm} takes only a few extra lines of code, and it does not increase the computation time.

In the same manner as $\Delta V(t)$ was computed earlier, a new function $R_{\text{norm}}(t)$ can also be computed. The new function will have a minimum at the position of the shift where the two least-squares lines fit to the data well. The new function $R_{\text{norm}}(t)$ is plotted in Fig. 7.9 along with $\Delta V(t)$. The two are then combined by dividing ΔV by $R_{\text{norm}}(t)$ to make a new function with a more distinct peak. Note this new function can only be used to find the time value of the DC-shift. To find the magnitude of the DC-shift, the time index would have to be referred back to ΔV .

In Fig. 7.9 the functions ΔV and $R_{\text{norm}}(t)$ are exactly the same size. This stays true when Δt is varied and the functions become surfaces. Without altering the process for locating the DC-shift, except for dividing ΔV with $R_{\text{norm}}(t)$, the DC-shifts of transitions 1–15 were located again. Figures 7.10 and 7.11 show the results.

When there is a clear shift present in the data, using the norms of the residuals of the least-squares lines does make the peak of the maximum more sharp. When the data is noisy, and there is no obvious shift, using the norms does not particularly help. This is because the ‘goodness of fit’ measure is usually better for the short lines than it is for the long ones, as they are more free to move to better fit the data.

So far only results from the first 15 transitions have been shown. To condense

²The norm of a data set is calculated the same way as finding the length of a vector: $\sqrt{(x_1^2 + x_2^2 + \dots + x_n^2)}$. If the elements of \mathbf{x} have a Gaussian distribution, then $\text{norm}(\mathbf{x}) \propto \sqrt{n}$

Residuals are Lower at the DC-shift




Figure 7.8: The “goodness of fit” can be quantified by using the norm of the residuals. The quantity R_{norm} is the sum of the norms of the residuals of the two least-squares lines. Plot (b), where the lines best detect the DC-shift, has a lower value of R_{norm} than the other two plots, (a) and (c).

the results of all 87 transitions (number 11 has been excluded), the times of the detected DC-shifts have been graphed as a histogram. The algorithm was still only using the max command, so one position was returned per transition whether there was a real DC-shift present or not. They were detected both with and without the use of the norms. For each case, two histograms have been plotted. One with a bin width of 0.129 s, giving 15 bins, and the second has a bin width of 0.267 s amounting to 31 bins. In all the histograms, shown in Fig. 7.12, there is a higher than average number of shifts detected close to $t = 0$. When R_{norm} is used in the detection, there is a slight increase in the number of DC-shifts close to zero. It seems that there are some DC-shifts at the SWS to REM transition, but with most of the transitions analysed here, it has not been observed. With more precautions to reduce movement and electrical artefacts during the recording process, cleaner and more reliable DC data could be gathered.

Something that is interesting is that number of detected shifts on the left side




Figure 7.9: The functions ΔV and R_{norm} . Dividing ΔV by R_{norm} produces a new function that has a very sharp peak that is easy to detect.

(during SWS) is decreased when R_{norm} is included. This is probably due to the straight lines having a poor fit to the large fluctuations that are seen in SWS. Even though the time-series was smoothed first to remove the fluctuations, sometimes they were still not smooth enough.

7.4 DC-shift Detection – A Work in Progress

I have come to the conclusion that this algorithm shows promise and could be developed further. The two least-squares lines, which slide along the time-axis, nicely detect all the obvious DC-shifts by tracing out local maxima (or minima for negative shifts) on the ΔV curve. But it really needs some logic to choose which maxima on ΔV should be identified as coming from a DC-shift, rather than just the global maximum. This could be difficult to implement without setting thresholds, which were originally avoided because magnitudes of the DC-shifts we hope to detect are unknown.

Even more important than detecting secondary shifts, is eliminating poor quality maxima all together so they are not mistaken as DC-shifts. Using the norms of the

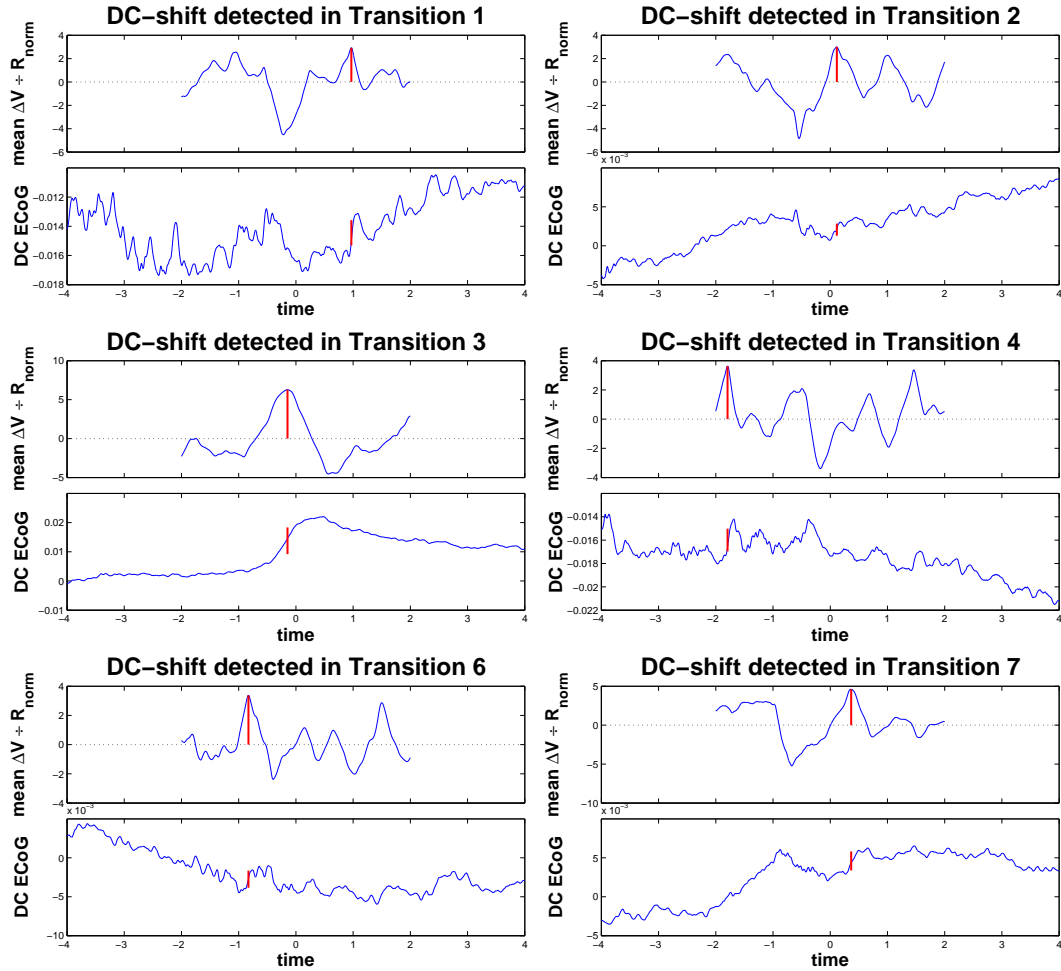


Figure 7.10: The norms of the residuals, which is a measure of the “goodness of fit” of the two least-squares lines, have been included into the DC-shift detection algorithm to increase its sensitivity. When the lines fit the data well, the peaks that indicate DC-shifts should be made more distinct. The same transitions as seen in Fig. 7.6 are shown here to see if the positions of any of the DC-shifts have changed. From these six, only the positive shift in transition 1 has changed. The peak on the left side has been reduced to a lower amplitude than the peak on the right. The shapes of the graphs of the other transitions have changed too, but not enough to mark a different position for the DC shift.

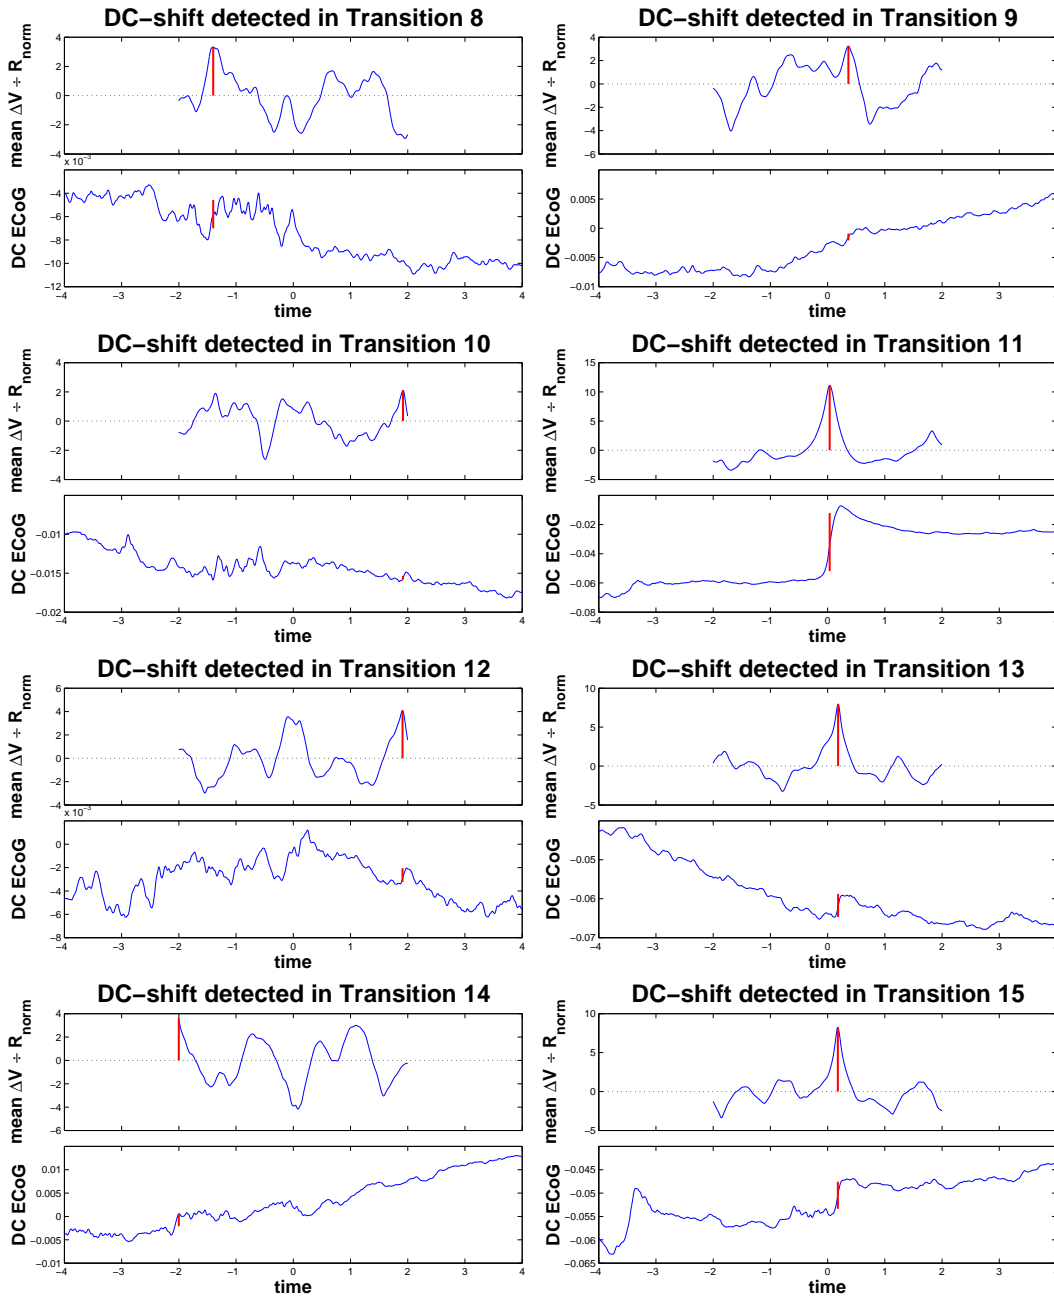


Figure 7.11: The transitions 8–15 that are shown here were also shown earlier in Fig. 7.7. After modifying the algorithm to include the norms of the residuals, 10 and 12 have got DC-shifts that are marked at different positions. A least-squares line passing through the fluctuations, between $t = -1.4$ s and $t = -0.5$ s, of transition 10, would have large residuals and therefore, a large value of R_{norm} at $t = -1.4$. The high R_{norm} causes the maxima at $t = -1.4$ to be reduced in height. In contrast the value of R_{norm} corresponding to the small “blip” in the time-series at $t = 2$ s would be much smaller. In transition 12, it is a similar situation. The DC-shift was detected close to $t = 0$, but all the fluctuations in that region would cause the spike to be lowered. The maxima corresponding to three clearest shifts, in numbers 11, 13 and 15, have all been enhanced by the inclusion of R_{norm} into the detection algorithm

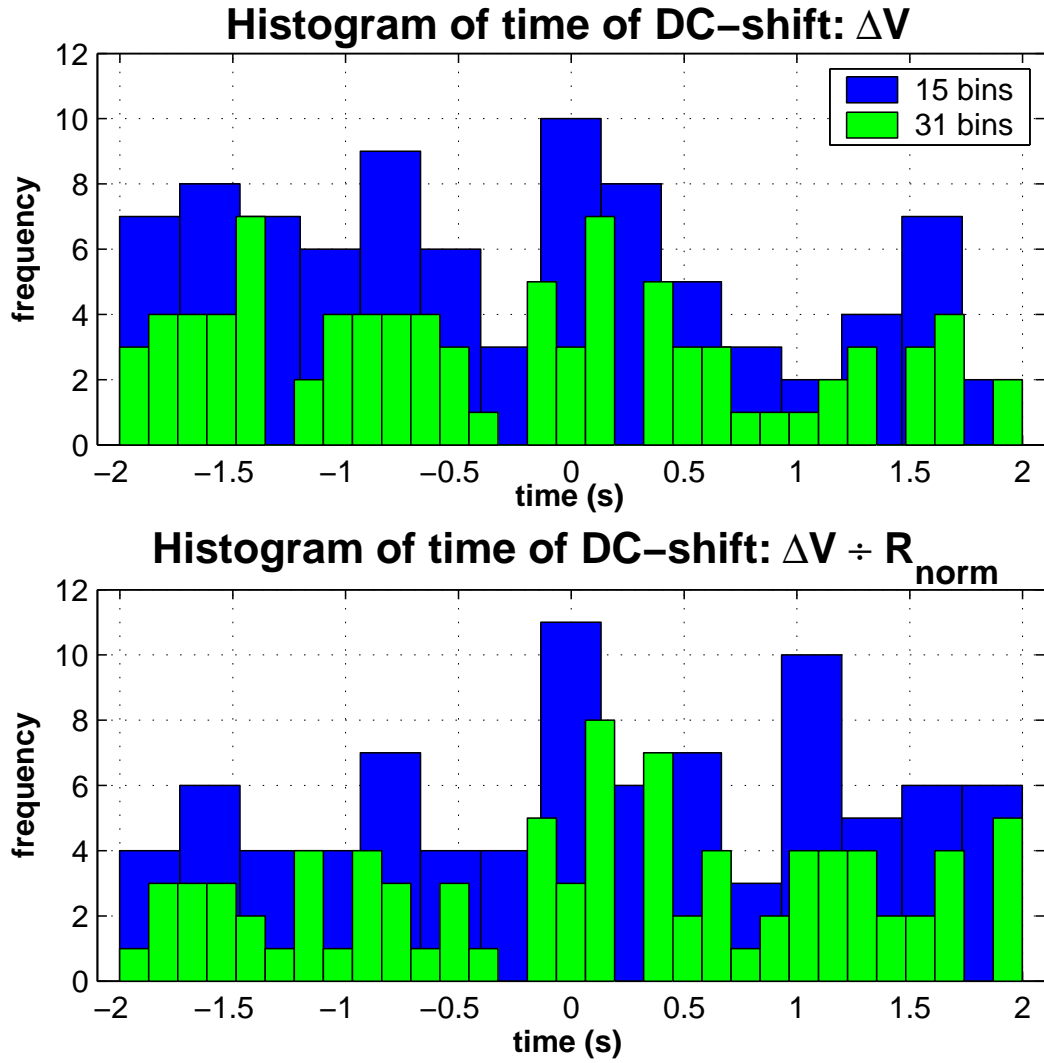


Figure 7.12: The positions of the detected DC-shifts when R_{norm} is and is not included in the detection algorithm. There is a small increase in the number of DC-shifts that occur close to $t = 0$, but it does not stand out well against the fluctuations of the rest of the graph and cannot be considered as evidence that a shift at $t = 0$ is more likely than anywhere else. When R_{norm} is used in the detection algorithm, less DC-shifts are detected on the right hand side, $t < 0$, which is during SWS.

residuals as a measure of the quality of the fit was intended to help reduce the likelihood of detecting an unwanted shift, but it did not work as well as I had hoped. The shorter lines are always able to fit the data better than the longer lines, and the distribution of the residuals is far from Gaussian, so dividing by \sqrt{n} did not help matters. Perhaps varying Δt is not required and an optimal value could be found. Only using one value of Δt would also drastically reduce the computation time.

Another way to test for good DC-shifts would be to measure the sharpness of the corresponding peak in ΔV . This could be done with a FWHM (full-width half-maximum). After the maximum of a peak is found and its height is measured, width of the peak at half its maximum value is also measured. Then if the width is greater than some threshold, then the peak is discarded and a different peak would be located. As mentioned earlier, setting thresholds has been avoided, but in this case it is a threshold in the time domain, not in voltage. After seeing the results of the algorithm so far, meaningful time thresholds for a FWHM could be set. The shape of the peaks that correspond to good DC-shifts are well known now and the minimum FWHM could be set to ~ 0.6 s.

I was pleased that the algorithm correctly identified the obvious shifts, but it seems that in most cases the data used either did not have any change in the DC potential, or it had so many changes that pinpointing only one was difficult.

Chapter 8

Conclusions and Future Work

The goal of this thesis was to find evidence supporting the Steyn-Ross *et al.* (2005) model of the sleeping cortex. The major prediction of the model is that the transition from SWS to REM sleep is a first-order phase transition.

The model predicts that as the cortex undergoes the phase transition, statistics describing the EEG fluctuations will dramatically change. During SWS, on approach towards the phase transition, there will be an increase in power and correlation time of the EEG and the energy will become more concentrated at low frequencies. Simultaneously the spectral entropy will decrease. At the point of transition the model predicts that power and correlation time will surge toward infinity and the energy will be concentrated at zero frequency. After the transition into REM has been made, the power and correlation time of the EEG will both drop to a much lower value and the frequency distribution will be much more even, leading to a higher spectral entropy.

In the human results presented in chapter 5, the time-series statistics — power, correlation time, fractional frequency band power and spectral entropy — were calculated on real EEG. The sharp drop in power, correlation time and fractional low-frequency power and the rise in spectral entropy after the transition were observed consistently. The large surge in power and correlation time were not seen every time. This was usually because a short period of stage 1 or 2 existed between the high-amplitude stage 3 and stage 4. One SWS to REM transition in particular showed an exceptionally good agreement with the predictions of the model. This occurred in the second sleep cycle of patient 3 (Fig. 5.12 on Pg. 67). This transition was so clear that the raw EEG and EOG time-series were also shown (Fig. 5.14 on Pg. 69). The transition was very abrupt, and the time of the transition can be pinpointed to within a few seconds.

In chapter 6, the same time-series statistics were also calculated on the ECoG

(similar to EEG but electrodes touch the surface of the cortex) of lab rats. Full time-series were difficult to analyse because the rats complete their sleep cycles so quickly. The best results came from time-aligning many time-series subsets so that the SWS to REM transitions occurred simultaneously. Averages were calculated and sharp drops in the power and correlation time and the jump in spectral entropy were all clearly visible. The fractional band power behaved slightly differently for the rats from the way it behaved for the humans. The low-frequency (0.1–4 Hz) did not show a consistent change. The frequency bands of most interest were the spindle (12–15 Hz) and gamma (25–47 Hz) bands. At the transition, the spindle band suddenly reduced and the gamma band increased.

The DC-coupled ECoG of the rats were also recorded in an attempt to detect the jump of V_e from the bottom branch of the manifold (Fig. 2.8 on Pg. 24) to the top branch that is predicted by the Steyn-Ross *et al.* (2005) model. These results were less convincing. In 14 SWS to REM transitions, a clear DC-shift was observed in only four of them. There were many irregular fluctuations in the DC ECoG that could possibly have been from electrical interference. To get more reliable results, more precautions need to be taken when recording down to zero frequency. DC signals are very sensitive and even a slight movement of the electrode wire can drastically change the DC potential.

I made the observation that one-way, high-pass filtering of a DC-shift returns a phase-shifted impulse response that looks very similar to a K-complex. Most EEG time-series are recorded using such a filter built into the amplifier. Perhaps DC-shifts do occur as the cortex transitions from SWS into REM, but in the AC-coupled EEG, they would appear to be K-complexes. According to the rules by Rechtschaffen and Kale (1968), seeing a K-complex is reason to score an epoch of sleep as stage 2.

Overall, the results presented in this thesis give good evidence supporting the phase transition theory of Steyn-Ross *et al.* (2005), with the most convincing evidence in the time-series statistics of patient 3 shown in Fig. 5.12 (Pg. 67).

8.1 Unfinished Work

During my own research, I needed to write many codes and functions to assist with the time series analysis. There are pieces of code that I am unhappy with in their present form.

One of the larger unfinished projects is the DC-shift detection algorithm that was the subject of chapter 7. I believe that the idea of this method shows promise,

but at the moment the max command is used to find the DC-shift. It does detect the obvious DC-shifts, but when there are several possible shifts, or no obvious shifts, the position and size of the DC-shift returned are sometimes questionable. There needs to be better logic conditions for specifying a “good” DC-shift.

In the human results, the correlation time showed a set of data points that were much lower than the rest (Fig. 5.7 on Pg. 59). These were the result of computing the autocorrelation of EEG containing spindles. The correlation time, as it stands, is defined as the lag time where the autocorrelation, Eq. (4.9), drops to $1/e$ of its maximum value. When the autocorrelation function oscillates as was shown in Fig. 5.9 (Pg. 62) the lag time recorded is only about $1/4$ of the period of the resonant oscillation. Marcus Wilson later informed me of another definition of the correlation time:

$$\tau_C = \frac{1}{C(0)} \int_0^{\infty} C(t) dt \quad (8.1)$$

Here the area of the autocorrelation function is calculated from zero lag right out to infinity, but the correlation time is usually only calculated on small epochs of only ~ 4 s. In the results, the correlation time came to be less than 200 ms. If an exponential decay with a time constant of 200 ms fits the autocorrelation well, then the decay will have fallen to less than 99% of its starting value after 1 s. For the EEG analysed in this thesis, setting the upper limit of Eq. (8.1) to 4 s is a fair approximation to infinity. If this new definition were implemented, it would probably evaluate faster than the current implementation, because this uses only a sum-and-divide operation. In the current form, the lag time is interpolated from the point where $C(\tau)/C(0)$ reaches $1/e$.

8.2 Future Work

The recordings of the DC-coupled ECoG made from the rats contained a lot of artefacts. These include movements and electrical interference. The evidence supporting a DC-shift as the cortex transitions from SWS to REM was not conclusive. To find conclusive evidence, we need to obtain DC recordings that are more reliable and free of artefacts.

It has been observed that K-complexes can be evoked in the EEG by the sleeper hearing a startling sound. If the DC-coupled EEG were also recorded then evoked K-complexes could be compared with their corresponding DC-EEG waveform. It is conceivable, though unlikely that all K-complexes are actually DC-shifts.

In some of the human sleep records, the ECG was clearly visible in the EEG as sharp spikes at a rate of about one per second. In sleep studies where only

the amount of time spent in SWS and REM is considered, the ECG artefact is not a problem. But the ECG does change the spectral distribution, so the results of any time-series statistics will also be changed. I would like to try to remove the ECG from the EEG. If the ECG is also recorded on its own channel, then using a correlation function or differential filter could be useful.

Appendix A

Matlab Codes Written

The most useful codes that were written are explained in this Appendix. The first section gives the details of the functions that were written to compute the time-series statistics from real EEG. The second section describes some very useful utility functions that were written to assist with displaying data. The third section discusses two different methods of data smoothing.

A.1 Time-Series Statistics Functions

A.1.1 Function: `sleppower`

```
P = sleppower(data, fs, epoch);  
P = sleppower(data, fs, epoch, overlap);  
[P, t] = sleppower(...);
```

I implemented the total power calculation in MATLAB with a function named `sleppower` (power is already a built-in MATLAB command). It requires 3 basic inputs. The first is the original time-series data, which can contain multiple channels of the time-series, but it should be arranged as a column-vector. The second input is the frequency that the data was sampled at, f_s . The third input argument is the epoch size, which is the time T (see Eq. (4.1)) that the calculation is to be performed over. There is a fourth input argument, which is optional. The `overlap` argument is a fraction between 0 and 1 that specifies how much each epoch should overlap with the previous one. The overlapping feature was described in Chapter 4.5. If `overlap` is not entered, it defaults to be zero.

The output of `sleppower` is the result of the power calculation P . It will be returned as an array with the new time-series going downwards in columns, with as many channels as were originally input. The units of P are the square of the units of

the original data. For example, if the EEG is measured in μV , then the power will be returned as $(\mu\text{V})^2$. There is also an optional output argument, `t`. This is a time vector, in units of seconds, which is for plotting against `P`. The time vector does not start at zero. Its first element has the value of one-half the epoch time. This is so that every data point will be plotted on the time axis at the centre of the of its epoch.

A.1.2 Function: `fract_pwr`

```
F = fract_pwr(data, fs, epoch);
F = fract_pwr(data, fs, epoch, overlap);
F = fract_pwr(data, fs, epoch, overlap, bands);
[F, t] = fract_pwr(...);
```

The fractional band power is calculated with the function `fract_pwr`. It has the same basic inputs and the optional time vector output as `sleppower`. There is a fifth input argument, `bands`, which specifies what frequency bands are to be calculated. Its default value is:

```
bands = [0.1, 4, 15, 47];
```

which means that `F` will be calculated for 3 different frequency bands,

```
low : 0.1 ≤ f < 4 Hz,
mid : 4   ≤ f < 15 Hz,
high: 15  ≤ f < 47 Hz.
```

The variable `bands` must have at least two elements to specify the upper and lower frequency limits. I have also provided an option to quickly set bands to the commonly used Greek alphabet frequency bands (see table 5.1 on Pg. 58). This is done by setting `bands` to the string 'greek'.

When `F` is returned, it is arranged as a three dimensional array. The time-series runs down the columns, the different channels run across the rows, and the frequency bands go back into the third dimension. This is illustrated in Fig. A.1 for an array containing 4 channels and 3 frequency bands. If only one channel was entered into the function, however, rather than returning an `n` by 1 by 3 array, the function will eliminate the singular dimension and return a two dimensional, `n` by 3, array.

A.1.3 Function: `correlation_time`

```
tau = correlation_time(data, fs, epoch);
tau = correlation_time(data, fs, epoch, overlap);
[tau, t] = correlation_time(...);
```

The function `correlation_time` has the same basic inputs and optional time vector output as `sleeppower`. Evaluating the correlation time is much more time consuming than the other time-series statistics described here. This is because the autocorrelation calculation must be performed for every epoch in the time-series. To make matters worse, the correlation time must be interpolated from the autocorrelation. Because the calculation runs so slowly, I would not use the correlation time for early analysis.

A.1.4 Function: entropy

```
H = entropy(data, fs, epoch);
H = entropy(data, fs, epoch, overlap);
H = entropy(data, fs, epoch, overlap, df);
H = entropy(data, fs, epoch, overlap, df, f_lim);
[H, t] = entropy(...);
```

The function `entropy` has the same basic data, sample rate epoch length and overlap input arguments as `sleeppower`. It also has two extra inputs.

The first input is the frequency bin width `df`, whose default value is $\frac{1}{2\pi} \approx 0.159$ Hz. This is to make $\Delta\omega = 1 \text{ rad s}^{-1}$. This value for `df` is not always possible, and needs to be checked.

When performing a discrete Fourier transform (DFT), the number of data points returned is the same as the number of points that was entered, N . Because the DFT of a purely real function (such as an EEG time-series) gives a function whose real part is symmetric, the negative frequencies are all redundant. The spacing of the




Figure A.1: The arrangement of the array that is returned by `fract_pwr`

frequencies after the DFT is

$$\Delta f = \frac{1}{N\Delta t} = \frac{f_s}{N}, \quad (\text{A.1})$$

and the number of samples in the epoch T is

$$N = \frac{T}{\Delta t} = T f_s. \quad (\text{A.2})$$

Now rearranging we find

$$T = \frac{N}{f_s} = \frac{1}{\Delta f} \quad (\text{A.3})$$

and therefore,

$$\Delta f = \frac{1}{T}. \quad (\text{A.4})$$

So the minimum bin width is $\frac{1}{T}$. To get Δf to stay at $\frac{1}{2\pi}$, the epoch period must be greater than $2\pi \approx 6.3$ seconds.

The function contains a check to test whether the the specified (or default) value of `df` is less than the minimum possible $\Delta f = \frac{1}{T}$. If test is true, then `df` is changed to be equal to Δf .

The second optional input, `f_lim`, is a two-element vector containing the frequency limits. It has form $[f_{\min}, f_{\max}]$. If it is not specified, `f_lim` defaults to $[0.1, 47]$ Hz.

A.2 Utility Software Tools

During the data analysis, I quickly realised that there were some simple signal processing and display tasks that would need to be performed very often. I wrote several MATLAB functions to open data files and display results, allowing for user input from either the mouse or keyboard in order to zoom, scroll, and even toggle which data trace is being displayed. These utility functions are `sleepwork`, `showline`, `zoomall` and `extreme`. They are described below.

A.2.1 Matlab Graphics

Before I can proceed with the explaining how the functions work, the graphics handling in MATLAB needs to be explained.

The MATLAB graphics environment maintains a hierarchical system of ‘parent’ and ‘children’ objects within each figure. The root level is at the top, and all figure elements are children of the root. All subplots contained within a given figure are

children of that figure. The subplots are then parents of many more graphics objects such as lines, surfaces and text. Each graphics object has a double-precision, floating-point identifier associated with it called a 'handle'. The handle points to a large list of 'properties' and 'property values'. For example a text label inserted into a graph has the property 'position', and its property value is an (x, y, z) co-ordinate.

A.2.2 Data Display Software

The patient data I first received from Fisher & Paykel Healthcare had already been sleep-staged by a sleep technician. I was given summary printouts of the sleep recording which contained a chart of sleep stage versus time. One of my first tasks was to view the patients' sleep EEG and understand how the sleep staging rules of Rechtschaffen and Kale (1968) are applied. This meant plotting data in ~ 30 second epochs and double checking the score that each epoch had been assigned. I did not want a simple loop that in each iteration only plotted the next 30 seconds of data. I wanted to be able to scroll forwards and backwards, increase or decrease the epoch size and easily choose what channels were to be displayed. There were also other operations that I wanted to include into the code, such as displaying the power spectral density or autocorrelation of the current epoch.

The final code had the following structure:

1. Initial user input
 - raw data file name
 - which channels are to be displayed, total = n
 - where in the file to start displaying from
2. Initialise figure
 - create figure
 - create n sub-plots
3. Begin while loop, stop when Escape is pressed
 - read and filter raw data
 - plot each channel in relevant subplot

- wait for user input
 - or ← scroll forward or backward
 - ↑ or ↓ zoom in or zoom out
 - 0 ... 5 assign sleep stage. 0 = wake, 1 = stage 1, 2 = stage 2, 3 = stage 3, 4 = stage 4, 5 = REM (these are disabled once staging is finalised)
 - f then i display PSD of channel i (where i is one of the numerical keys)
 - c then i display autocorrelation of channel i
- reset the start and stop index for reading in next epoch of data

4. Save the current state of the sleep-staging file

This code proved to be very useful for viewing data and helping me get a feel for EEG sleep staging and artefact detection. As this was the first work I did on the project, I (perhaps foolishly) named the program `sleepwork`. Figure A.2 is a screenshot of the display produced by `sleepwork`. It displays each different channel in its own subplot. When fewer channels are specified, it will adjust the height of each subplot so that the display fills the screen. Figure A.3 shows examples the PSD and autocorrelation that are displayed after pressing the `f` or `c` keys.




Figure A.2: A screen shot of the display of `sleepwork`. The bold face labels at the left of each subplot are the names of the channels. LOC/ROC = left/right electrooculogram (EOG). LEG = leg movements, MIC = microphone, FLOW = airflow.




Figure A.3: Screenshots of the new figures that are displayed when using the psd or autocorrelation options of sleepwork.

A.2.3 Function: showline

```
showline
showline(L)
showline(L,string)
```

The function `showline` operates on a set of subplots containing many plotted traces, clearing all but one from the display. Pressing the left or right arrow key causes the visible trace to disappear and the next trace in the sequence to appear. This is very useful when examining cluttered graphs.

It has two main ways of starting. The first is to automatically detect the handles l of all n traces in each of the m subplots and arranges the handles in an array:

$$\mathbf{L} = \left(\begin{bmatrix} l_{1,1} \\ l_{2,1} \\ \vdots \\ l_{n,1} \end{bmatrix} \begin{bmatrix} l_{1,2} \\ l_{2,2} \\ \vdots \\ l_{n,2} \end{bmatrix} \cdots \begin{bmatrix} l_{1,m} \\ l_{2,m} \\ \vdots \\ l_{n,m} \end{bmatrix} \right) \quad (\text{A.5})$$

The subscript i and j , of $l_{i,j}$, mean the i^{th} trace of the j^{th} subplot. This method fails if there are a different number of traces in each axes, because the columns of L would have different lengths.

The second method of starting `showline` is to input \mathbf{L} as an argument. In that case, it does not matter how the trace handles are arranged in \mathbf{L} , just as long as \mathbf{L} is rectangular. The function then goes through the process of making all traces whose handles are in the j^{th} row of \mathbf{L} visible, while all the other traces are rendered invisible. This means that if the number of traces contained in each subplot are not the same, \mathbf{L} can be made rectangular by entering one handle into \mathbf{L} more than once. For example, Fig. A.4 shows one subplot with three traces, while the second subplot

has only one trace. For `showline` to alternate only traces l_a , l_b and l_c in the first subplot, the array \mathbf{L} could be constructed as:

$$\mathbf{L} = \left(\begin{bmatrix} l_a \\ l_b \\ l_c \end{bmatrix} \begin{bmatrix} l_d \\ l_d \\ l_d \end{bmatrix} \right) \quad \text{or} \quad \mathbf{L} = \begin{pmatrix} l_a \\ l_b \\ l_c \end{pmatrix} \quad (\text{A.6})$$




Figure A.4: An example described above of three traces in the first subplot and only one in the second subplot. To get l_a , l_b and l_c to alternate using `showline`, \mathbf{L} would have to be constructed as in Eq. (A.6).

In both cases for \mathbf{L} , `showline` will perform in the same way. The plots in the first subplot will alternate between the three traces, l_a , l_b and l_c , while the second subplot will continuously show the one trace, l_d . The function will run slightly slower for the first \mathbf{L} , as it will make the trace in the second subplot invisible then visible every time. Using the second version of \mathbf{L} , `showline` will leave the l_d trace alone and it will stay visible all the time.

The function `showline` contains a *while* loop that sets the `visible` property to `off` for the handles contained in \mathbf{L} , then turns it back on for l_i , where $i = 1$ as an initial condition. This displays all traces whose handles are in the i^{th} row in the array \mathbf{L} , and hides all the others. The function then goes into a “paused while waiting for user input” state with the use of the built-in MATLAB function `ginput`. Only the \leftarrow , \rightarrow , space and Esc keys are acceptable input in this state. If a \leftarrow or \rightarrow is input, i is decreased or increased by 1. The function also checks if $i < 1$, then $i = n$, and also if $i > n$, then $i = 1$, which gives the variable i periodic boundary conditions. The function then goes back to the start of the *while* loop with a new value for i . The loop is ended when the Escape key or space bar is pressed and the function ends. The difference between Esc and space is that space will leave the figure with all the traces except l_i invisible, while Esc will make all traces become visible again.

There is a second optional input argument that was included to provide a label for each different trace. This label was often just the channel number of the trace being displayed. The second argument needs to be a character array, with the same number of rows as **L**. When the i^{th} trace is displayed, the i^{th} row of the character array would be displayed in the top right corner of the plot.

An unexpected, but useful result from using the `showline` function was that the scale of the axes changes automatically to fit the currently visible trace. This occurs because of MATLAB's default behavior to set the axis properties 'xlimmode' and 'ylimmode' to 'auto'. It can be overwritten by the command: `axis(axis)1`, which holds the limits at their current values.

Figure A.5 shows three curves of the form

$$y = \frac{\sin(\omega t)}{t} \quad (\text{A.7})$$

where $\omega = 1, 2$ and 4 radians per second. All three curves were plotted on the same axes and then the `showline` function was used to hide two of the curves from the display. I also made use of the string argument to show the value of ω on each axes. These plots also demonstrate the convenient auto-scaling action, as the y-axis limits change for the smaller amplitude functions.

A.2.4 Function: `zoomall`

```
zoomall
zoomall('x')
zoomall('y')
zoomall('xy')
```

When working with time-series analysis, often many graphs are displayed and all are plotted as a function of time. It is common to want to zoom in to take a closer look at a selected subset of the data. A difficulty arises in that the zoom function built into MATLAB, needs to be performed separately and manually on each subplot within the graphics window. This is done by selecting a region of interest with the mouse; after the button is released, the axes zooms in to display only what was selected in the box. This selection procedure must then be repeated for each subplot within the figure if all subplots are to share a common axis scaling. Because it is extremely difficult to draw the same sized box for each subplot, the zoomed subplots will usually have slightly different limits.

¹`axis` has two modes of operation, giving an input argument: `axis(limits)`, sets the current axis limits to the values specified by `limits`, while `axis` by itself returns the current axis limits

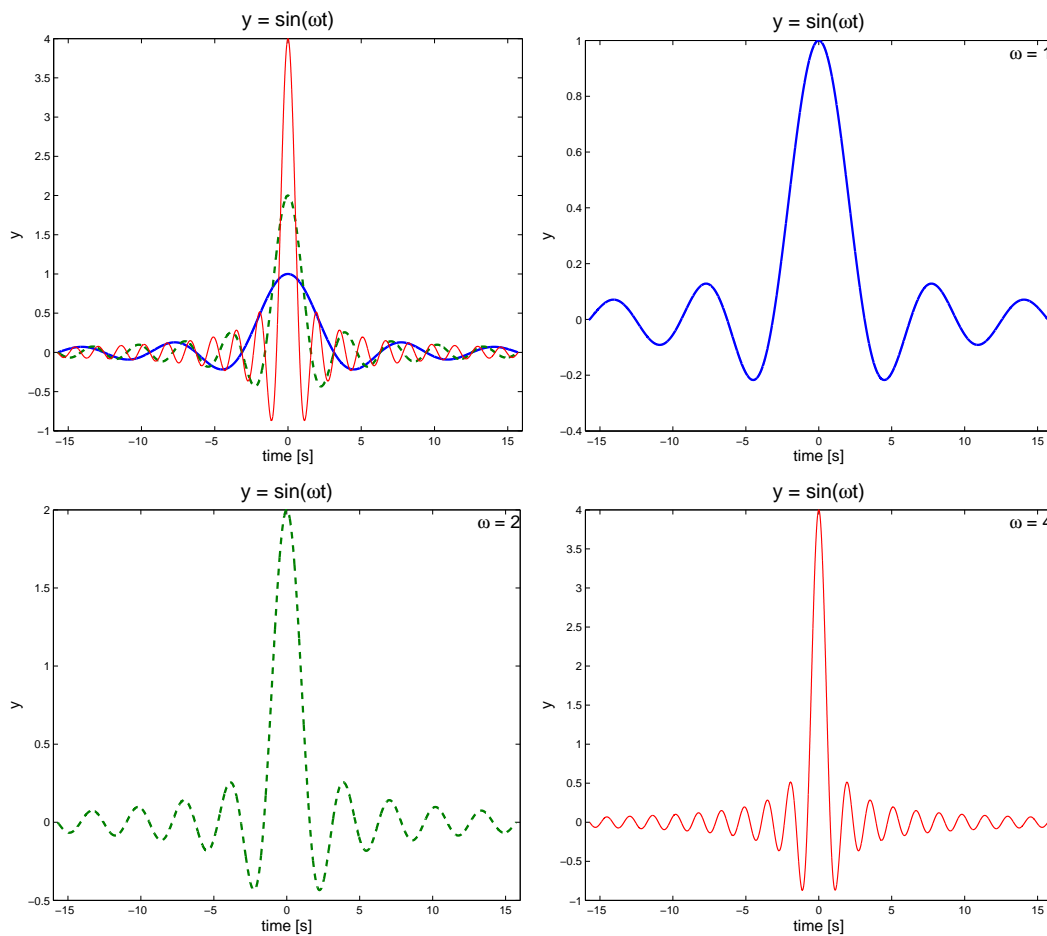


Figure A.5: An example of showline in action. Upper-left: Three traces plotted with different values of ω . Upper-right: showline removes two traces so only the $\omega = 1$ trace is plotted. Lower-left: The $\omega = 2$ trace is made visible, the other two are suppressed. Lower-right: Only the $\omega = 4$ trace is visible.

I wrote the `zoomall` function to allow the user to zoom within one subplot and have the function automatically zoom all the other axes to the same limits. This was achieved by first turning on the zoom attribute, then calling the function `waitforbuttonpress`, which pauses the running function until a mouse button or key is pressed. Fortunately, if zoom is turned on, the `waitforbuttonpress` function does not un-pause until the mouse button is released, and the axes has zoomed in. The function finds the limits of the axis which was last ‘clicked’ in (the `axis` command will return the limits of the current subplot). The `zoomall` function then sets all the axis limits to the correct values.

I found that it was more useful to implement the zooming operation by handling the x -axis and y -axis separately. This meant that an option of zooming in x or y or both could be included in as an input argument. In fact, as most of my data is plotted with time on the x -axis, I set the default behavior to zoom only in the x direction.

A.2.5 Function: extreme

```
x = extreme(data)
x = extreme(data,n)
```

At times it is desirable to plot the full \sim six hours of raw data, but at a sample rate of 256 Hz, this corresponds to $\sim 5.5 \times 10^6$ sample points per channel. Attempting to display this quantity of data can be slow and unwieldy if using an older computer. To help with this, I decided to *subsample* the data set when plotting raw time-series. Often subsampling is performed by taking every i^{th} sample, to reduce the data length from N samples by $\frac{1}{i}$ to give a subsampled data set of length $n = \frac{N}{i}$. This, however, may leave out a sharp peak, especially if i is large (comparable to the sample rate).

I decided to subsample the data in a different way. In the function `extreme`, instead of taking one sample per i points, I took *two* samples per $2i$ points, where the two sample values were selected to be the local maximum and minimum values within that data subset. This was achieved by reshaping the data vector of N samples columnwise into an array of size R rows by C columns, where

$$R = 2i = \frac{2N}{n} \quad (\text{A.8})$$

and

$$C = \frac{N}{2i} = \frac{n}{2} \quad (\text{A.9})$$

Notice that n must be even and a factor of N to keep the values all integers. The new R by C data array D looks like:

$$D = \begin{pmatrix} \begin{bmatrix} d_{1,1} \\ d_{2,1} \\ \vdots \\ d_{R,1} \end{bmatrix} & \begin{bmatrix} d_{1,2} \\ d_{2,2} \\ \vdots \\ d_{R,2} \end{bmatrix} & \cdots & \begin{bmatrix} d_{1,C} \\ d_{2,C} \\ \vdots \\ d_{R,C} \end{bmatrix} \end{pmatrix}. \quad (\text{A.10})$$

Using the `max` and `min` commands, the maximum and minimum values for each column in the array D are calculated and returned as row-vectors.

$$\begin{aligned} D^{\max} &= \max[D] \\ D^{\max} &= \left(\max \begin{bmatrix} d_{1,1} \\ d_{2,1} \\ \vdots \end{bmatrix} \max \begin{bmatrix} d_{1,2} \\ d_{2,2} \\ \vdots \end{bmatrix} \cdots \right) \\ D^{\max} &= \begin{pmatrix} d_1^{\max} & d_2^{\max} & \cdots \end{pmatrix} \end{aligned} \quad (\text{A.11})$$

and

$$\begin{aligned} D^{\min} &= \min[D] \\ D^{\min} &= \begin{pmatrix} d_1^{\min} & d_2^{\min} & \dots \end{pmatrix}. \end{aligned} \quad (\text{A.12})$$

Then the D^{\max} and D^{\min} row vectors are combined to give a new matrix:

$$\begin{bmatrix} D^{\max} \\ D^{\min} \end{bmatrix} = \begin{bmatrix} d_1^{\max} & d_2^{\max} & \dots \\ d_1^{\min} & d_2^{\min} & \dots \end{bmatrix} \quad (\text{A.13})$$

This last array is reshaped into a column vector which contains only the extreme values of the matrix D

$$D^{\text{extreme}} = \begin{bmatrix} d_1^{\max} \\ d_1^{\min} \\ d_2^{\max} \\ d_2^{\min} \\ \vdots \end{bmatrix} \quad (\text{A.14})$$

The code for the function `extreme` also contains some important housekeeping to ensure the original data vector can be reshaped into the $\frac{n}{2}$ by $\frac{2N}{n}$ matrix D . If the number of points, N , of the original data vector is not a multiple of $C = \frac{n}{2}$, D cannot be constructed. This problem was fixed by first finding the closest values of R and C that give a product only slightly larger than N . The definitions of R and C had to be modified to

$$C = \text{round}\left(\frac{n}{2}\right) \quad (\text{A.15})$$

and

$$R = \text{ceil}\left(\frac{N}{C}\right) \quad (\text{A.16})$$

where ‘round’ rounds up or down to the nearest integer and ‘ceil’ (ceiling) rounds up to the nearest larger integer. The ‘ceiling’ operation in Eq. (A.16) is used to ensure that $R \times C \geq N$. By adding $(R \times C) - N$ points, whose values are calculated from local averages, into the original data vector, at regular intervals, the length of the data vector is increased to be exactly $R \times C$.

To give an example of the operation of `extreme`, let us work through a simple numerical example. Let the vector x be all integers in the interval $[-8, 9]$, and then let $y = x^2$. We have $N = 18$ points and want to subsample down to $n = 12$. The

matrix D will have $C = \frac{n}{2} = \frac{12}{2} = 6$ columns and $R = \frac{2N}{n} = \frac{2 \times 18}{12} = 3$ rows.

$$D = \left(\begin{bmatrix} 64 \\ 49 \\ 36 \end{bmatrix} \begin{bmatrix} 25 \\ 16 \\ 9 \end{bmatrix} \begin{bmatrix} 4 \\ 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 4 \\ 9 \end{bmatrix} \begin{bmatrix} 16 \\ 25 \\ 36 \end{bmatrix} \begin{bmatrix} 49 \\ 64 \\ 81 \end{bmatrix} \right) \quad (\text{A.17})$$

The following D^{\max} and D^{\min} are:

$$D^{\max} = (64 \ 25 \ 4 \ 9 \ 36 \ 81) \quad (\text{A.18})$$

$$D^{\min} = (36 \ 9 \ 0 \ 1 \ 16 \ 49) \quad (\text{A.19})$$

The final D^{extreme} vector, is now

$$D^{\text{extreme}} = (64 \ 36 \ 25 \ 9 \ 4 \ 0 \ 9 \ 1 \ 36 \ 16 \ 81 \ 49) \quad (\text{A.20})$$

Notice how the values at the right-hand side of the vector begin to alternate high and low. Figure A.6 shows a graph of D^{extreme} and the original $y = x^2$ data. When plotting the new subsampled y values, the corresponding new x values are uniformly spaced starting from x_1 and ending and x_{18} this leads to some distortion in the x -axis. The extreme function is clearly not suitable for small-scale subsampling as shown here.




Figure A.6: A plot of the function $y = x^2$ and data that has been subsampled from 18 points to 12 using extreme

For a more realistic example, it is better to use real EEG. In Fig. A.7 I have plotted 250 seconds of EEG data ($f_s = 400$ Hz) that came from a lab rat. The raw

data were subsampled using `extreme` to give only 2000 data points. After zooming in along the x -axis, the subsampling begins to be noticeable. When there are twice as many data points than there are pixels on the screen, the subsampled graphs looks fine, but when the ratio of pixels to points gets below 1:1, the zig-zag structure of the subsampling becomes visible.

In functions where `extreme` was used, I usually set n to be about three or four thousand, as this was guaranteed to be greater than the number of pixels in the width of the screen (e.g., 1024, 1152 or 1280).

A.3 Data Smoothing

During the time series analysis, I found that the results often had a lot of fluctuation due to the small changes in the EEG amplitudes and frequencies. There were also spikes due to artefacts in the raw data. The results needed to be smoothed. To get a good idea of the gross changes,

A.3.1 Convolution Moving Average

The first smoothing technique I applied was a moving average. The most natural way to apply a moving average smoother is to take the average of all the data points from i to $i + n$ then another average for the points from $i + 1$ to $i + n + 1$ and so on. This method is extremely slow, especially if the data set is large. A better way is to *convolve* the data set with a small vector of all ones. The convolution of two functions f and g is denoted by $f * g$ and defined as:

$$h(t) = (f * g)(t) = \int f(\tau) g(t - \tau) d\tau \quad (\text{A.21})$$

The convolution operation produces a third function h , which is a measure of the similarity between f and g at an offset τ . When written in discretely Eq. (A.21) becomes:

$$h_k = (f * g)_k = \sum_{j=1}^N f_j g_{k+1-j} \quad (\text{A.22})$$

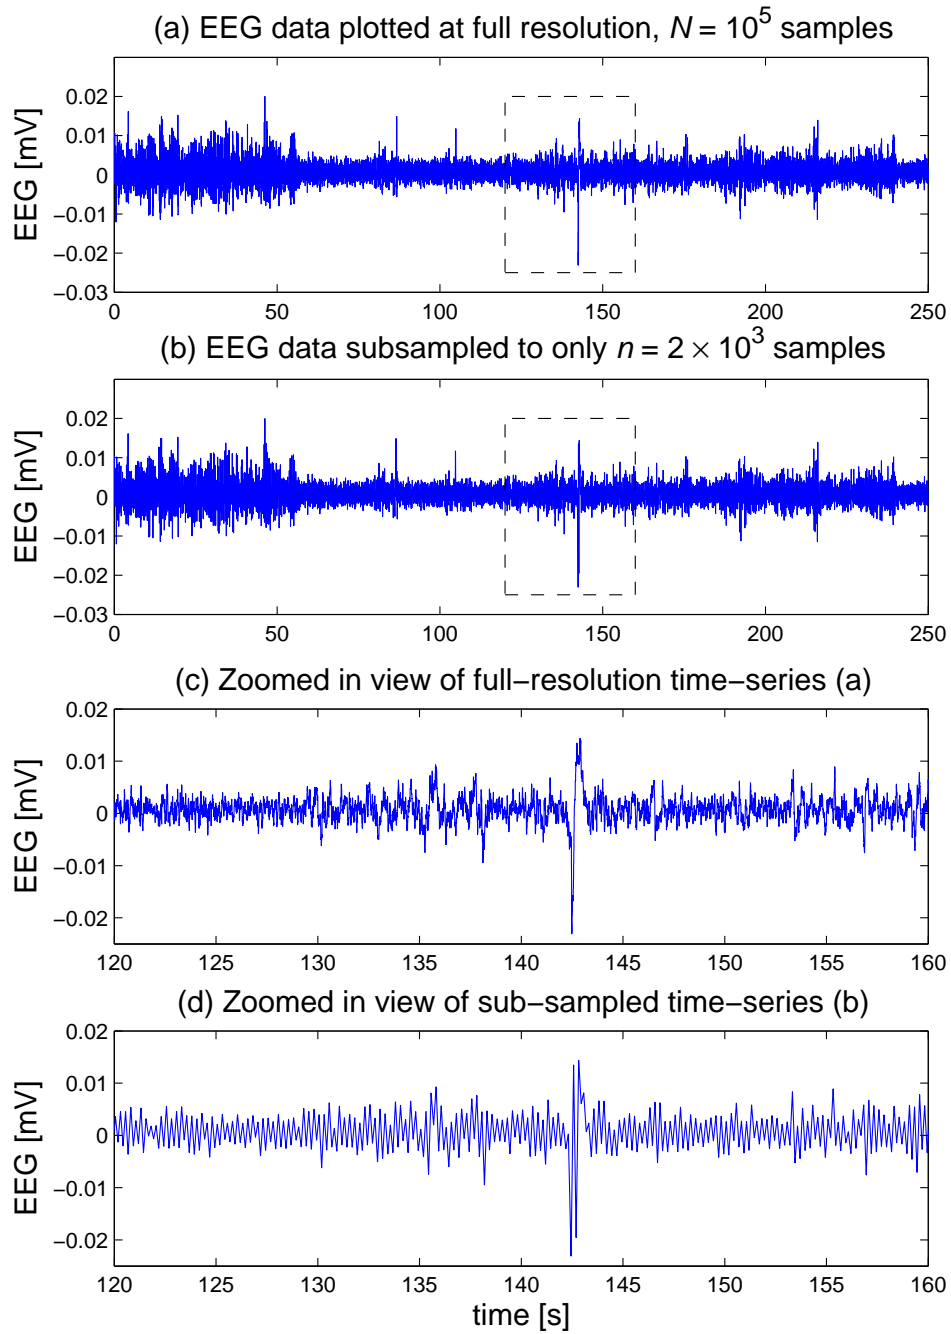


Figure A.7: Demonstration of the extreme function. Before plotting, the width of the x-axis was set to exactly 500 pixels on the monitor. (a) 100,000 samples of raw EEG data is plotted, giving 2,000 points plotted per pixel (the sample frequency f_s is 400 Hz). (b) using the extreme function the raw data was subsampled to only 2,000 samples. At this resolution, there are four data points to every pixel. To the eye, there is no difference in (a) and (b). (c) The raw data has been zoomed in along the x-axis to display only 16,000 samples, which is 32 samples per pixel. (d) The subsampled data is zoomed in to the same time scale as (c) and now only has 320 data points to be displayed in the 500 pixels.

where N is the number of elements in the first vector, f . For an example let $f = (1, 2, 3, 4)$ and $g = (1, 1, 1)$. Then the first term h_1 will be:

$$\begin{aligned}
 h_1 &= \sum_{j=1}^4 f_j g_{1+1-j} \\
 &= f_1 g_1 + f_2 g_0 + f_3 g_{-1} + f_4 g_{-2} \\
 &= 1 \times 1 + 2 \times 0 + 3 \times 0 + 4 \times 0 \\
 &= 1
 \end{aligned} \tag{A.23}$$

Notice that when the subscript index does not exist, as in g_0 or g_{-1} the missing value is replaced by zero. The other terms of h are computed similarly. To better illustrate the convolution procedure I have combined f and g into an array of the form $\begin{bmatrix} f \\ g \end{bmatrix}$. Now the next terms in the convolution process can be computed by taking the sums of the columnwise products. (The use here of the symbols \sum and \prod are not strictly correct)

$$\begin{aligned}
 h_2 &= \sum \left[\prod \begin{pmatrix} (1 & 2 & 3 & 4) \\ (1 & 1 & 1) \end{pmatrix} \right] \\
 &= \sum \begin{bmatrix} 0 & 1 & 2 & 0 & 0 \end{bmatrix} \\
 &= 3
 \end{aligned} \tag{A.24}$$

$$\begin{aligned}
 h_3 &= \sum \left[\prod \begin{pmatrix} (1 & 2 & 3 & 4) \\ (1 & 1 & 1) \end{pmatrix} \right] \\
 &= \sum \begin{bmatrix} 1 & 2 & 3 & 0 \end{bmatrix} \\
 &= 6
 \end{aligned} \tag{A.25}$$

$$\begin{aligned}
 h_4 &= \sum \left[\prod \begin{pmatrix} (1 & 2 & 3 & 4) \\ (1 & 1 & 1) \end{pmatrix} \right] \\
 &= \sum \begin{bmatrix} 0 & 2 & 3 & 4 \end{bmatrix} \\
 &= 9
 \end{aligned} \tag{A.26}$$

$$\begin{aligned}
h_5 &= \sum \left[\Pi \left(\begin{array}{cccc} (1 & 2 & 3 & 4) \\ & & (1 & 1 & 1) \end{array} \right) \right] \\
&= \sum \left[\begin{array}{ccccc} 0 & 0 & 3 & 4 & 0 \end{array} \right] \\
&= 7
\end{aligned} \tag{A.27}$$

$$\begin{aligned}
h_6 &= \sum \left[\Pi \left(\begin{array}{cccc} (1 & 2 & 3 & 4) \\ & & (1 & 1 & 1) \end{array} \right) \right] \\
&= \sum \left[\begin{array}{cccccc} 0 & 0 & 0 & 4 & 0 & 0 \end{array} \right] \\
&= 4
\end{aligned} \tag{A.28}$$

Thus for this example, $h = f * g$

$$h = (1 \ 3 \ 6 \ 9 \ 7 \ 4) \tag{A.29}$$

The first and last elements of h have very little meaning as only one non-zero term was used in their calculation. Whenever a convolution is evaluated close to the end points of a data vector, the resulting points lose their significance.

A convolution may not sound like a fast method, for performing something as trivial as a moving average. However, by applying the convolution theorem from mathematics (Spiegel, 1963), we can find an algorithm that is very quickly implemented numerically.

$$\mathcal{F}[f * g] = (\mathcal{F}[f] \cdot \mathcal{F}[g]) \tag{A.30}$$

where \mathcal{F} is the Fourier transform operator. By applying the inverse Fourier transform \mathcal{F}^{-1} we get

$$f * g = \mathcal{F}^{-1}[\mathcal{F}[f] \cdot \mathcal{F}[g]] \tag{A.31}$$

MATLAB has a built-in function, called `conv`, which evaluates convolutions using Eq. (A.31). It needs to do some zero-padding to keep the numerical calculation exact (Mathworks, 2002). If we have the two vectors f and g , then the convolution $f * g$ can be evaluated by:

```

F = fft([f zeros(1,length(g)-1)])
G = fft([g zeros(1,length(f)-1)])
then
conv(f,g) = ifft(F.*G).

```

For the moving average calculation, f is the data set and g is a small vector of length n containing only ones. The result of the convolution is not exactly a moving

average yet, it must be divided by the sum of g which in this case is simply n . For my implementation of the convolution moving average, I actually used another MATLAB function called `convn`, which is a multi-dimensional convolution. This was used for two reasons. Firstly, so that multiple channels could be smoothed simultaneously, and secondly, because `convn` provides an optional input argument to truncate the ends of $f * g$ so that it becomes the same size as f . Using `h = convn(f,g,'same')` to evaluate the example above, would return `h = [3 6 9 7]` which has the same number of elements as f , which has eliminated the most meaningless points from the ends.

To demonstrate the convolution smoothing I have plotted the spectral entropy (see section 4.4) of two minutes of rat EEG. In Fig. A.8, there are 20 points plotted per second. I set the convolution smoother to use a vector of 21 ones.




Figure A.8: Smoothed data using a convolution moving average smoother. Notice the flat, horizontal line at the edge of the data in (a) – this is an undesirable edge effect.

One problem with the moving average smoother is that the end points become very strongly pulled towards zero. This is because at the edges of the data vector,

less than n data points are summed but then they are still divided by n , making the average less than what it should have been. I got around this by simply replacing these distorted points with copies of the closest correct point. This left strange looking horizontal lines at the edges of the data set, but was better than the earlier problem, which is shown in Fig. A.9




Figure A.9: The same data set as Fig. A.8 except the edges are pulled down towards zero from the convolution process

A.3.2 Whittaker Difference Smoothing

```
z = whittaker(y,lambda)
z = whittaker(y,lambda,d)
```

Jamie Sleight, one of my colleagues, informed me of another smoothing technique. Eilers (1994) gives a smoothing technique that was first published by Whittaker in 1923, who originally called it “graduation”. Using today’s language, Eilers calls it “penalized least squares”.

Suppose we have a data series that consists of m points y_i sampled at equal intervals, to which we want to fit a smoothed series z_i . A “goodness of fit” measure is the sum of the squares of the residuals,

$$S_1 = \sum_{i=1}^m (y_i - z_i)^2. \quad (\text{A.32})$$

A similar equation can be written to describe the smoothness of z in terms of successive z -differences, of difference order d , $\Delta z_i = z_i - z_{i+d}$,

$$S_2 = \sum_{i=2}^m (\Delta z_i)^2. \quad (\text{A.33})$$

S_1 and S_2 can be combined to give a measure of both fit and smoothness

$$\begin{aligned} S &= S_1 + \lambda S_2 \\ S &= \sum_{i=1}^m (y_i - z_i)^2 + \lambda \sum_{i=2}^m (\Delta z_i)^2. \end{aligned} \quad (\text{A.34})$$

The parameter λ is used to trade between smoothness of z and fit to the data y . S now needs to be minimised in the the same fashion as a least-squares polynomial fit. In MATLAB this is best done using matrix arithmetic. In matrix notation Eq. (A.34) is

$$S = |y - z|^2 + \lambda |Dz|^2 \quad (\text{A.35})$$

where $Dz = \Delta z$, For example, when $m = 4$ and $d = 1$, D is

$$D = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{pmatrix} \quad (\text{A.36})$$

and the product of Dz is

$$Dz = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \\ z_3 \\ z_4 \end{pmatrix} = \begin{pmatrix} -z_1 + z_2 \\ -z_2 + z_3 \\ -z_3 + z_4 \end{pmatrix} \quad (\text{A.37})$$

The solution of minimising Eq. (A.35) is found by the system of linear equations

$$(I + \lambda D'D)z = y \quad (\text{A.38})$$

The equation can be solved for z in one line of MATLAB code using the ‘left matrix divide’ operator ‘\’

$$z = (I + \lambda D' * D) \setminus y;$$

Note that if there is a missing value in the data, eg. one of the $y_i = \text{NaN}$, then the above line of code will return z that contains only NaNs. Eilers (1994) gives a

solution to this problem by flagging all missing values and assigning weights, one for a true value and zero otherwise. The weights are arranged into a diagonal matrix W which replaces both the identity matrix I and replaces y with Wy in Eq. (A.38).

$$(W + \lambda D'D)z = Wy \quad (\text{A.39})$$

The Whittaker smoother is very fast and robust, and even better smoothing can be achieved with using second or third order differences, e.g., $d = 2 \Rightarrow \Delta_2 z = z_i - z_{i+2}$. One problem I have found, is that only experience can teach what to use for the value of λ . At first guess, I would try $\lambda = 10, 100$ or 1000 , but never less than 1, because then it only returns $z \approx y$.

In Fig. A.10 I have plotted the same spectral entropy data as in Fig. A.8, but this time smoothed using Whittaker's method. I needed to get the Whittaker smoother to produce similar results as the convolution smoother, but the parameters for the two algorithms mean completely different things. To compare like with like, I performed the Whittaker smoother with a range of λ until its output S_W was very similar to those of the convolution smoother S_C . I settled on the parameters $\lambda = 1000$ and the difference order $d = 2$. The rms difference between S_W and S_C was only 0.0023.

A.3.3 Comparison of Convolution and Whittaker Smoothing

To compare S_W and S_C , there are a few factors that can be considered. The first is the quality of the smoothing. Is the smoothed curve actually smooth? To help answer this question I have plotted the two smoothed curves from above again, and also their first (numerical) derivative in Fig. A.11.

It is clear that smoothing data using Whittaker's method gives much smoother results than a convolution moving average. There are also no bad edge effects with the Whittaker smoother.

Another question to test the quality of smoothing should be asked. Is the smooth curve true to the data? To help answer this, I have plotted the same spectral entropy data again and made new smooth curves S_C and S_W . This time the number of points used in the convolution was increased to 31. As before a new value of λ had to be chosen to give a minimum rms difference between S_C and S_W . At $\lambda = 6000$, the rms difference was 0.0027; Figure A.12 shows how the convolution moving average curve S_C sometimes changes in anti-phase with the data. This undesirable behavior is due to the equal weighting of all the data points used in calculating the moving average.

Perhaps a weighted moving average, that uses a Gaussian shape rather than the




Figure A.10: Smoothed data using a Whittaker smoother, $\lambda = 1000$ and $d = 2$. Notice that the curve follows the data points at the edge of the data in (a).

simple square shape I that have chosen, would give better smoothing results. However, The Whittaker smoothing technique performs so well that there is not much point in persisting with the convolution moving average.




Figure A.11: From looking at the first derivatives, the Whittaker smoother produces much smoother curves than the convolution smoother




Figure A.12: The convolution smoother is calculated with $n = 31$. Now the smooth line seems to go in anti-phase with the fluctuations of the data. The Whittaker does not have these problems.

Appendix B

Extra Rat Transition Results

The figures contained in this Appendix were calculated at the same time as Fig. 6.7 on Pg. 91. However, these results are all very similar and provide little extra information. If included in the main body they would only disrupt the flow of text. The figures show a consistent drop in power, rise in spectral entropy, rise in gamma frequency power and drop in spindle frequency power. For all The figures, the DC ECoG shows many fluctuations, and no consistent trend. engineering

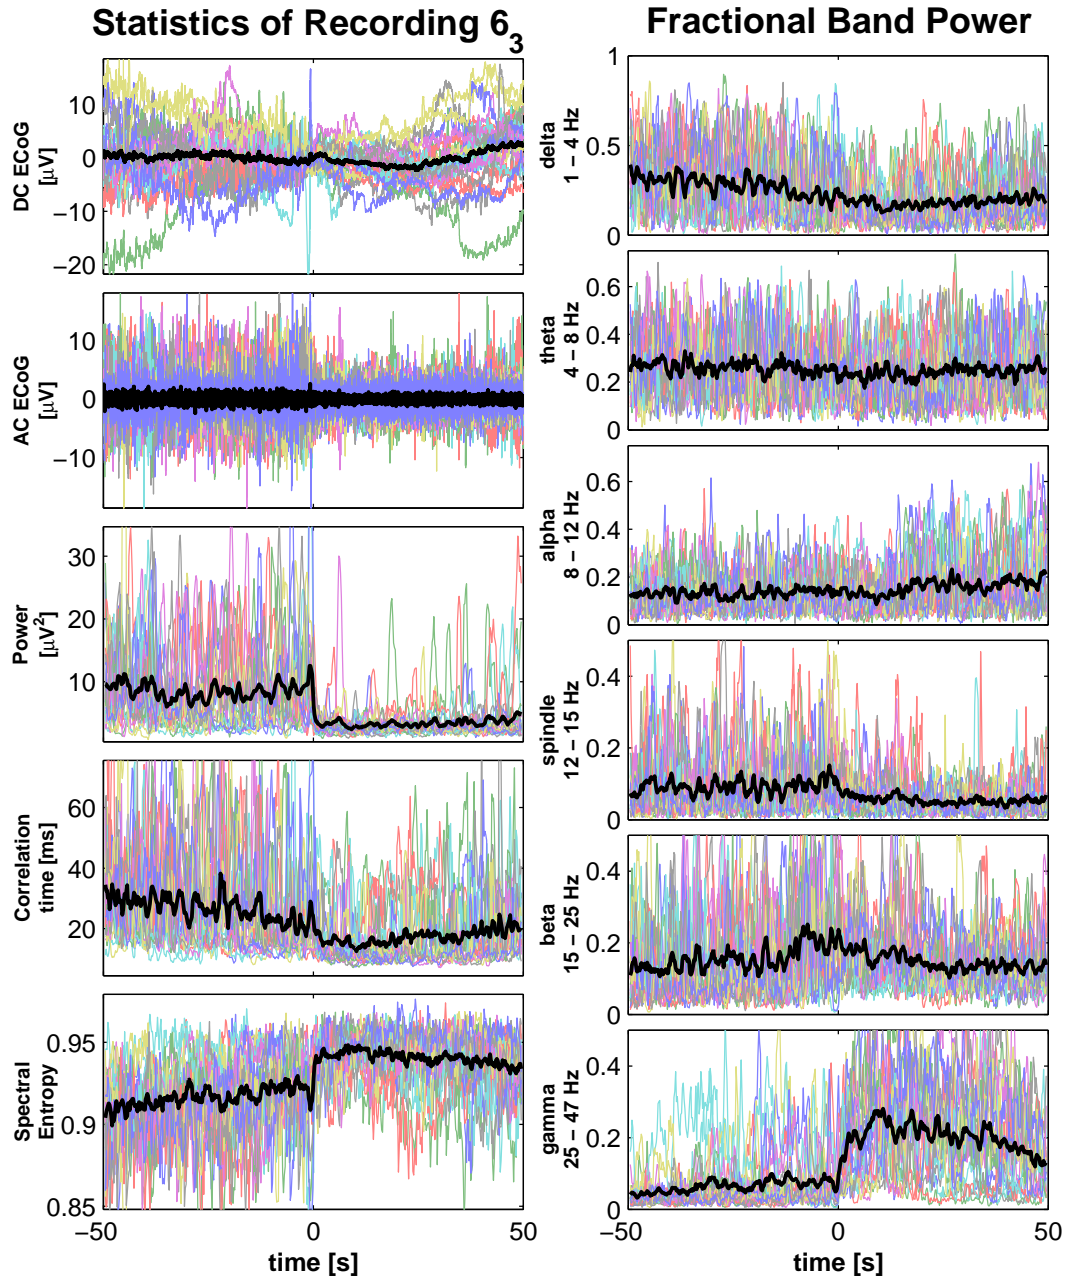


Figure B.1: The 20 time-aligned SWS to REM transitions for rat recording 6.3. The SWS to REM transition is again, most clear in the power, spectral entropy, spindle band and gamma band. The DC channel shows no clear positive jump in potential.

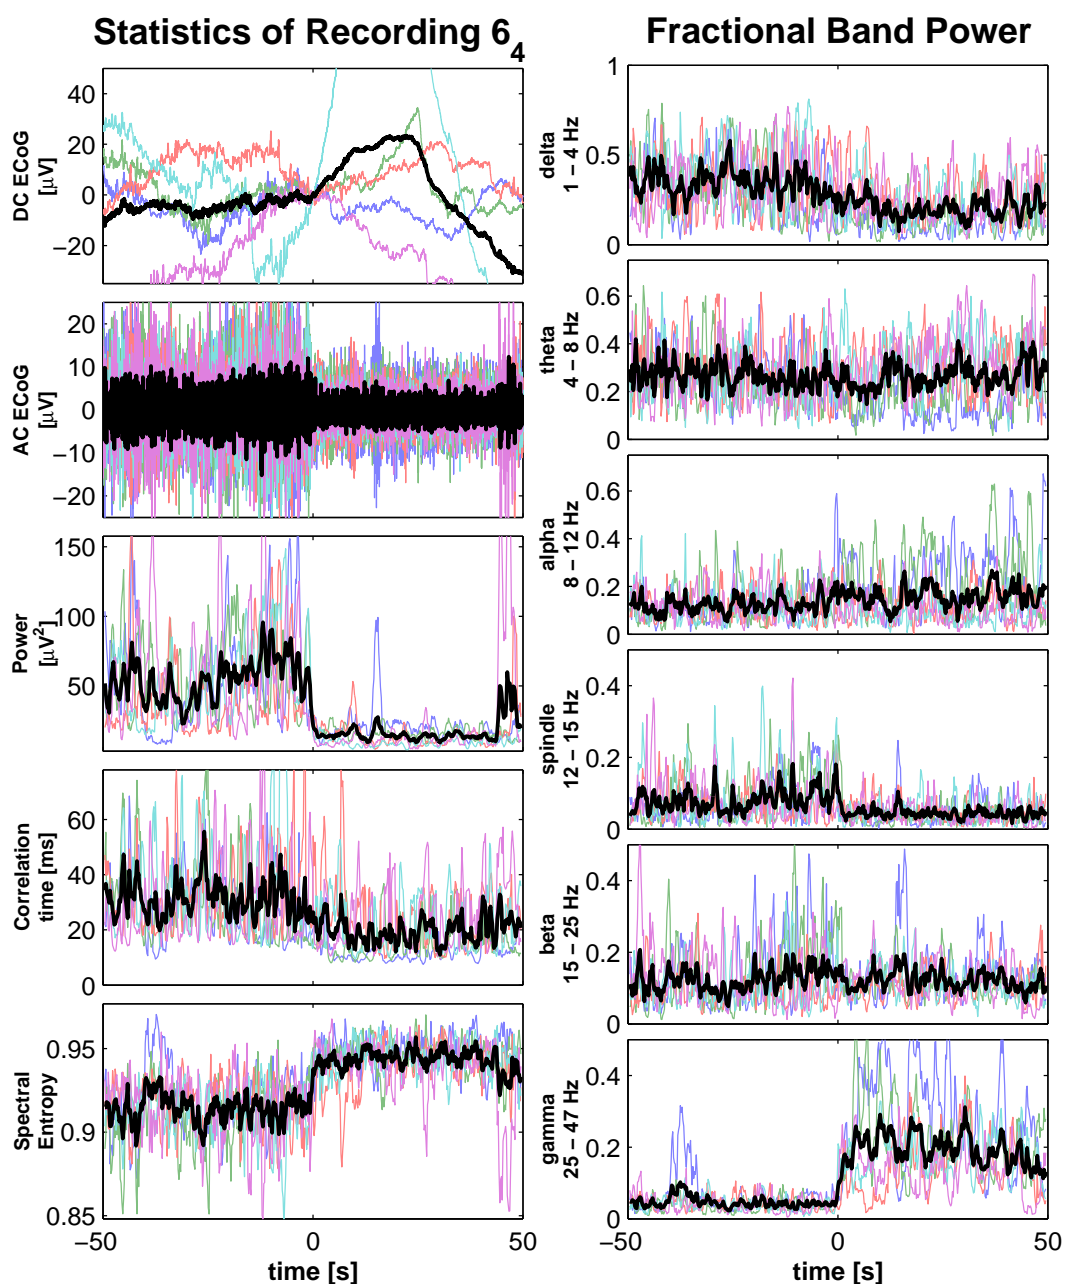


Figure B.2: The 5 time-aligned SWS to REM transitions for rat recording 6.4. Here the alpha frequency band shows a slight increase in the REM state. It is still not as distinctive as the gamma band, spindle bands, power or spectral entropy. There is a increase in the slope of the average DC potential, but our theory predicts a much quicker change than what is seen here.

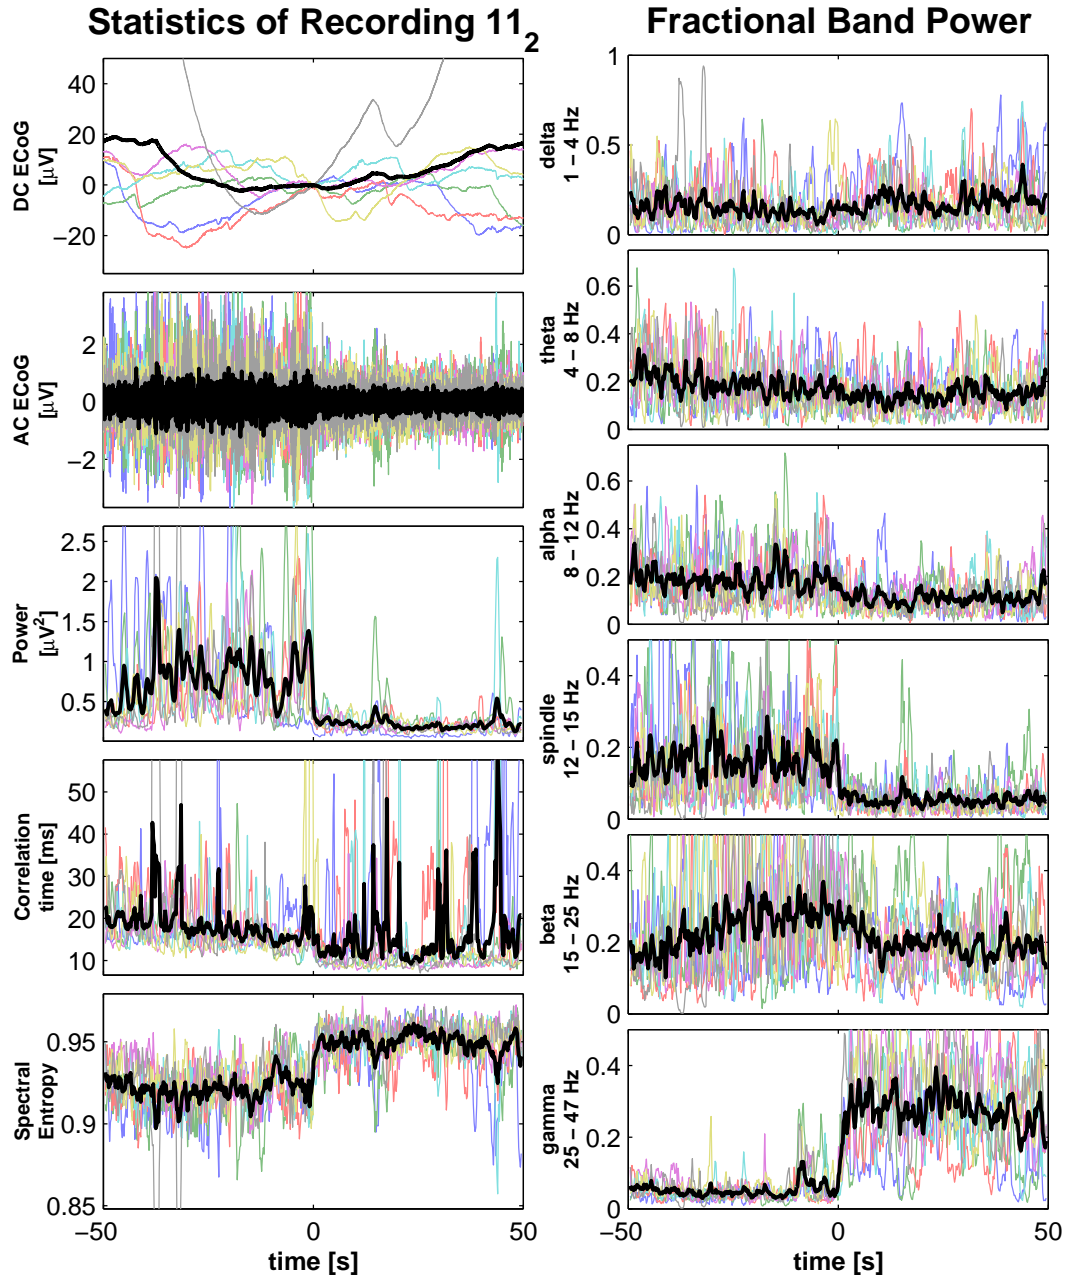


Figure B.3: The 7 time-aligned SWS to REM transitions for rat recording 11.2. The SWS to REM transition is clear in the time-series statistics. Contrary to rat recording 6.4 in Fig. B.2, the alpha frequency band decreases. Some of the DC time-series have positive slopes while others are negative and they cancel out leaving the average reasonably flat.




Figure B.4: 8 time-aligned SWS to REM transitions for rat recording 12.1. The changes in power, spectral entropy and spindle and gamma bands mark the transition at $t = 0$. The DC ECoG time-series show no consistent trend.

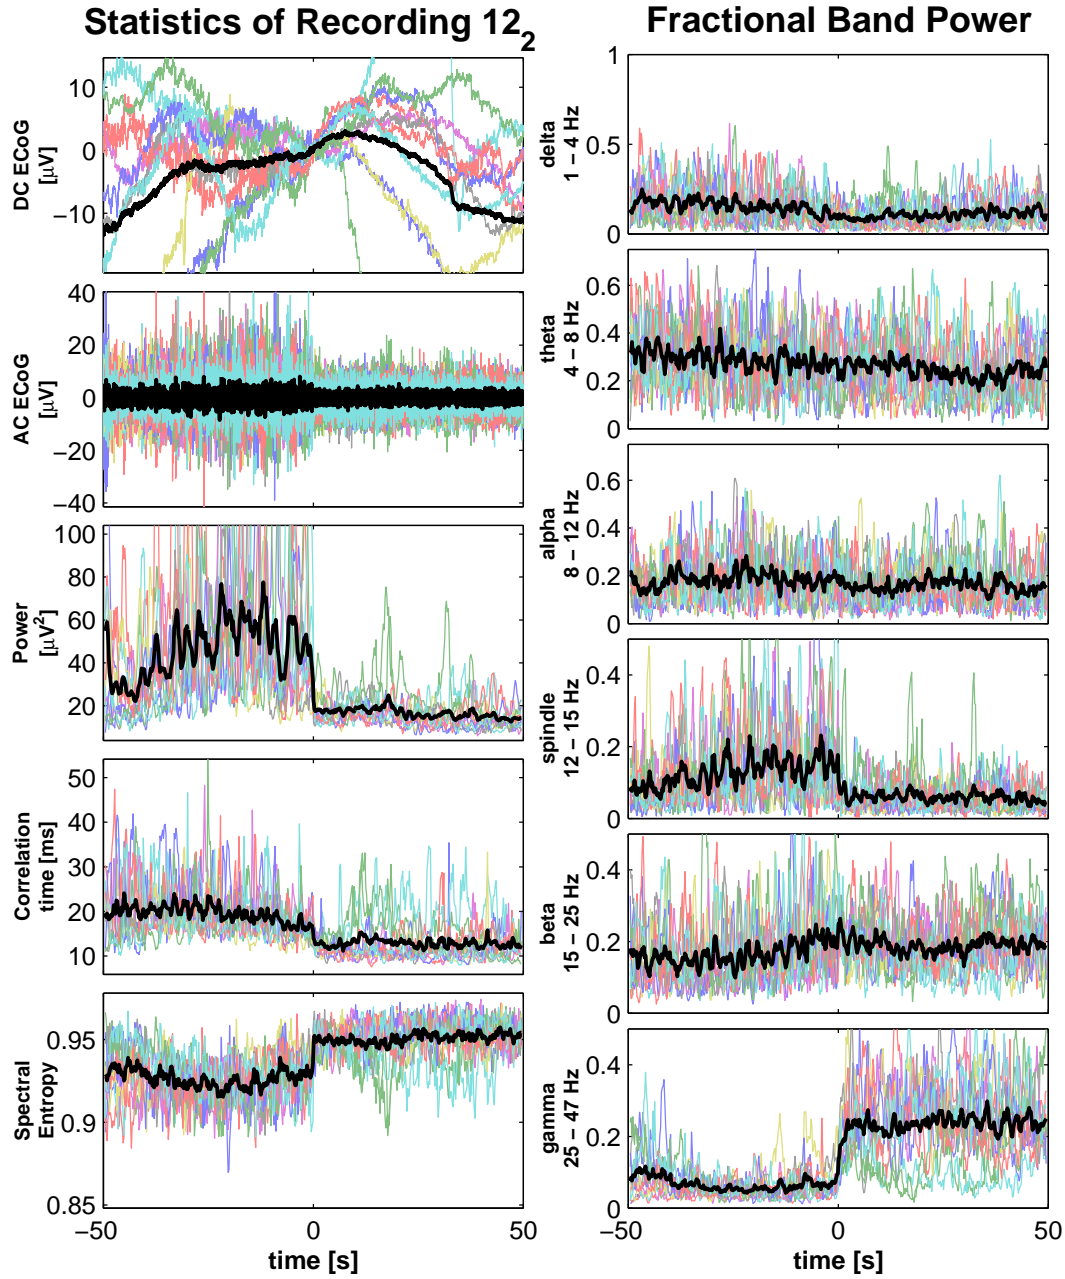


Figure B.5: 11 time-aligned SWS to REM transitions for rat recording 12.2. This is one of the few rat recordings where the correlation time gives a sharp change similar to the power and spectral entropy. The average DC potential seems to increase slightly at the point of transition, but it is not fast enough to call a DC shift.

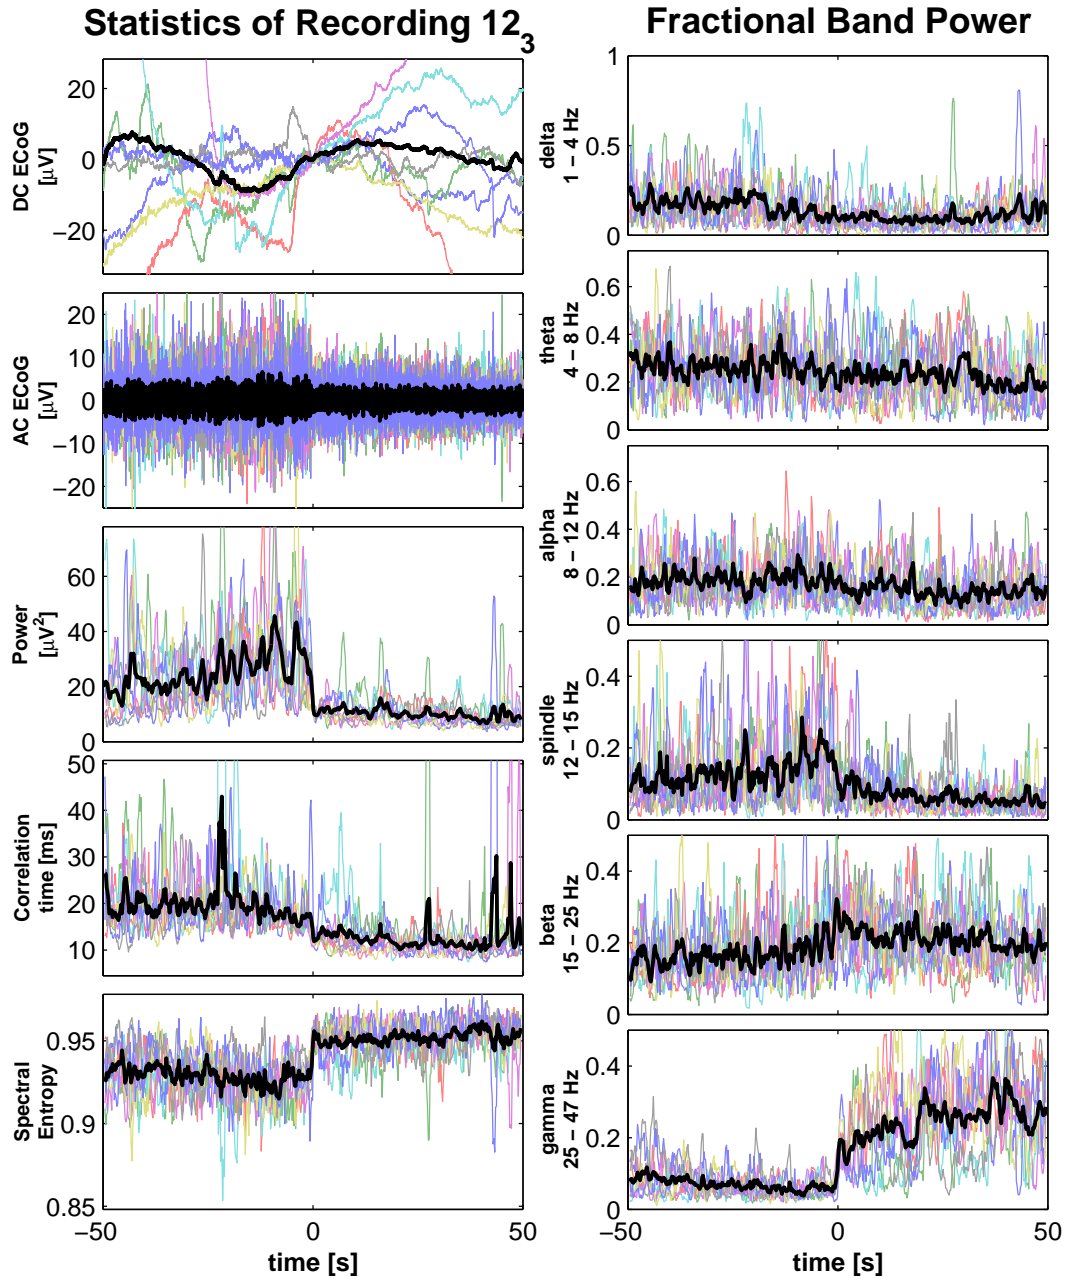


Figure B.6: 8 time-aligned SWS to REM transitions for rat recording 12.3. The DC average has a positive increase at about -5 seconds, it comes from one of the transitions with a large DC shift.

References

- Bavhenov, M., I. Timofeev, M. Steriade, and T. J. Sejnowski. Model of thalamo-cortical slow-wave sleep oscillations and transitions to activated states. *Journal of Neuroscience*, **22**, pp. 8691–8704 (2002).
- Blake, A. and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge Massachusetts (1987).
- Braitenberg, V. and A. Schüz. *Anatomy of a Cortex: Statistics and Geometry*. Springer-Verlag, Berlin (1991).
- Dement, W. and N. Kleitman. Cyclic variations in eeg during sleep and their relation to eye movements, body motility, and dreaming. *Electroencephalography and Clinical Neurophysiology*, **9**, pp. 673–690 (1957).
- Destexhe, A., D. Contreras, and M. Steriade. Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *Journal of Neuroscience*, **19(11)**, pp. 4595–4608 (1999).
- Eilers, P. H. C. Smoothing and interpolation with finite differences. In: *Graphic Gems*, edited by P. S. Heckbert, volume IV, pp. 241–250. Academic Press (1994).
- España, R. and T. E. Scammell. Sleep neurobiology for the clinician. *Sleep*, **27(4)**, pp. 811–820 (2004).
- Freeman, W. J. Predictions on neocortical dynamics derived from studies in paleo-cortex. In: *Induced Rhythms of the Brain*, edited by E. Basar and T. H. Bullock, chapter 9, pp. 183–199. Birkhaeuser, Boston (1992).
- Hammond, C. *Cellular and Molecular Neurobiology*. Academic Press, Orlando, FL (2001).
- Hasselmo, M. E. Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behavioural Brain Research*, **67**, pp. 1–27 (1995).

- Hill, S. and G. Tononi. Modeling sleep and wakefulness in the thalamocortical system. *Journal of Neurophysiology*, **93**, pp. 1671–1696 (2005).
- Liley, D. T. J., P. J. Cadusch, and J. J. Wright. A continuum theory of electro-cortical activity. *Neurocomputing*, **26–27**, pp. 795–800 (1999).
- Marshall, L., H. L. F. M. Mölle, and J. Born. Scalp recorded direct current brain potentials during human sleep. *European Journal of Neuroscience*, **10**, pp. 1167–1178 (1998).
- Marshall, L., M. Mölle, and J. Born. Spindle and slow wave rhythms at slow wave sleep transitions are linked to strong shifts in the cortical direct current. *Neuroscience*, **121**, pp. 1047–1053 (2003).
- Mathworks. *Matlab Functions Help*. The Mathworks, MA, USA (2002).
- Nunez, P. L. The brain wave function: A model for the EEG. *Mathematical Biosciences*, **21**, pp. 279–297 (1974).
- Rechtschaffen, A. and A. Kale. *A Manual of Standardized Terminology, Techniques, and Scoring System for Sleep Stages of Human Subjects*. U.S. Govt Printing Office, Washington DC (1968).
- Rennie, C. J., J. J. Wright, and P. A. Robinson. Mechanisms for cortical electrical activity and emergence of gamma rhythm. *Journal of Theoretical Biology*, **205**, pp. 17–35 (2000).
- Robinson, P. A., C. J. Rennie, and J. J. Wright. Propagation and stability of waves of electrical activity in the cerebral cortex. *Physical Review E*, **56**, pp. 826–840 (1997).
- Shannon, C. E. and W. Weaver. *The Mathematical Theory of Information*. University of Illinois Press, Urbana, Illinois (1949).
- Spiegel, M. R. *Theory and Problems of Advanced Calculus*. McGraw-Hill, Singapore (1963).
- Steriade, M. and R. W. McCarley. *Brain Control of Wakefulness and Sleep*, 2nd ed. Kluwer Academic/Plenum Publishers, New York (2005).
- Steyn-Ross, D. A. *Modelling the Anaestheto-Dynamic Phase Transition of the Cerebral Cortex*. Ph.D. thesis, University of Waikato, Hamilton, New Zealand (2002).

- Steyn-Ross, D. A., M. L. Steyn-Ross, J. W. Sleigh, M. T. Wilson, I. P. Gillies, and J. J. Wright. The sleep cycle modelled as a cortical phase transition. *Journal of Biological Physics*, **31**, pp. 547–569 (2005).
- Steyn-Ross, D. A., M. L. Steyn-Ross, L. C. Wilcocks, and J. W. Sleigh. Toward a theory of the general anaesthetic-induced phase transition of the cerebral cortex: II. Stochastic numerical simulations, spectral entropy, and correlations. *Physical Review E*, **64**, p. 011918 (2001a).
- Steyn-Ross, M. L., D. A. Steyn-Ross, and J. W. Sleigh. Modelling general anaesthesia as a first-order phase transition in the cortex. *Progress in Biophysics and Molecular Biology*, **85**, pp. 369–385 (2004).
- Steyn-Ross, M. L., D. A. Steyn-Ross, J. W. Sleigh, and D. T. J. Liley. Theoretical electroencephalogram stationary spectrum for a white-noise-driven cortex: Evidence for a general anaesthetic-induced phase transition. *Physical Review E*, **60**, pp. 7299–7311 (1999).
- Steyn-Ross, M. L., D. A. Steyn-Ross, J. W. Sleigh, and D. R. Whiting. Theoretical predictions for spatial covariance of the EEG signal during the anaesthetic-induced phase transition: Increased correlation length and emergence of self-organization. *Physical Review E*, **68**, p. 021902 (2003).
- Steyn-Ross, M. L., D. A. Steyn-Ross, J. W. Sleigh, and L. C. Wilcocks. Toward a theory of the general anaesthetic-induced phase transition of the cerebral cortex: I. A statistical mechanics analogy. *Physical Review E*, **64**, p. 011917 (2001b).
- Vanhatalo, S., J. M. Palava, M. D. Holmes, J. Miller, J. Voipio, and K. Kalia. Infralow oscillations modulate excitability and interictal epileptic activity in the human cortex during sleep. *Proceedings of the National Academy of Sciences*, **101**, pp. 5053–5057 (2004).
- Weiten, W. *Psychology Themes and Variations, second edition*. Brooks/Cole, Pacific Grove, CA (1992).
- Whiting, D. R. *Modelling the spatial effects of the anaesthetic-induced phase transition in the cerebral cortex*. Master's thesis, University of Waikato, Hamilton, New Zealand (2003).
- Wilson, M. T., M. L. Steyn-Ross, D. A. Steyn-Ross, and J. W. Sleigh. Predictions and simulations of cortical dynamics during natural sleep using a continuum approach. *Physical Review E*, **72**, pp. 051910–1–051910–14 (2005).

- Wright, J. J. and D. T. J. Liley. Dynamics of the brain at global and microscopic scales: Neural networks and the EEG. *Behavioral and Brain Science*, **19**, pp. 285–316 (1996).