

Epilepsy research using nonlinear signal
processing

Angus Wallace BSc, BEng (Biomedical Honours)

July 9, 2008

Abstract

This thesis applies several standard nonlinear quantifiers to EEG analysis to examine both human primary generalised epilepsy (PGE) and rat models of human epilepsy.

We analysed rat EEG, and then used the analysed data, in parallel with an impedance recording, to better understand the events during experiments. Next, the nonlinear analysis of EEG was used to attempt to model the behaviour of the impedance data. This modelling did not yield a useful predictive tool, so we recommend the continued recording of impedance data as a means of augmenting EEG recordings.

The analyses were also applied to human data, and showed differences between the PGE and control groups in apparently normal EEG. We then attempted to use these differences to detect the presence of PGE in an unclassified subject – a diagnostic tool. This was done using a feed-forward neural network. We found that the inter-group differences were exploitable and facilitated the diagnosis of PGE in previously unknown subjects. The extent to which this is useful as a diagnostic tool should be assessed by further trials.

Finally, the analyses were used to examine data from a paralysed human subject, in an attempt to identify the mental task being performed by that subject. This was not successful, suggesting that the same analyses that were useful in discriminating between PGE and control were not useful in detecting the mental state of the subject. It was also apparent that the presence of EMG (in an unparalysed state) assisted task-classification.

TABLE OF CONTENTS

1	Introduction	7
1.1	What will be in this thesis	7
1.2	Epilepsy and EEG analysis	9
1.3	How to analyse signals	12
2	Linear Systems	14
2.1	What is a linear system?	14
2.2	Linear analyses	15
2.2.1	Autocorrelation function	16
2.2.2	Fourier transform	16
2.2.3	Wavelet analysis	19
2.2.4	Eigenvalues/vectors	21
2.2.5	Principal component analysis	21
2.2.6	Independent components analysis	23
2.2.7	Autoregressive Analysis	26
2.2.8	Laplacian EEG analysis	28
2.3	Data as a representation of the system	31
2.3.1	Stationarity	31
2.3.1.1	Stationarity assessment	32
2.3.2	How many data are “enough data”?	33
2.3.2.1	Linear Perspective	33
2.3.2.2	Nonlinear Perspective	34
3	Nonlinear Systems	36
3.1	What is a nonlinear system?	37
3.2	What is a chaotic system?	37
3.3	Nonlinear analyses	40
3.3.1	Reconstruction of the Phase Space	40
3.3.1.1	Embedding	42
3.3.1.2	Choosing an embedding dimension	43
3.3.1.3	Choosing a Time Lag	44
3.3.2	Correlation dimension	45
3.3.3	Entropy	46
3.3.4	Mutual information	49
3.3.5	Nonlinear prediction	50
3.3.6	Nonlinear filtering	51
3.3.7	Lyapunov exponents	52
3.3.8	Data Compression	53
3.4	Requirements of nonlinear analysis	53

3.4.1	Bootstrapping	54
3.4.2	Surrogate Data	55
3.4.2.1	Non-stationarity	57
3.4.2.2	Application	57
4	Classification	60
4.1	Linear discriminant analysis	62
4.2	Artificial Neural Networks	65
4.2.1	The Neuron Model	65
4.2.1.1	The artificial neuron	65
4.2.2	Network Structure	67
4.2.2.1	Network memory	68
4.2.3	Network Training	69
4.2.3.1	Back-propagation	69
4.2.3.2	Supervised vs unsupervised training	70
4.2.3.3	Parallel and Series-Parallel Networks	71
5	Brain background	73
5.1	Anatomy	73
5.1.1	The neuron	74
5.1.2	Non-neuronal cells	77
5.1.3	Neural circuits	78
5.1.4	Brain Regions	79
5.2	EEG	80
5.2.1	Limitations of EEG	82
5.2.2	Patterns	83
5.3	Epilepsy	85
5.3.1	Epileptic seizure anticipation	87
5.3.2	Epilepsy diagnosis	88
5.3.3	Cell Swelling	88
5.3.4	Relations between cell-swelling and impedance	89
5.4	Nonlinear brain research	90
5.4.1	Mental task classification	91
5.4.2	Modelling of the brain as a system	92
5.4.3	Spreading depression	93
6	Data collection	95
6.1	Pharmacologically-induced rat epileptogenesis	95
6.1.1	Preparation	96
6.1.2	Experimental procedure	96
6.1.3	Pharmacological agents	97

6.1.3.1	Picrotoxin	98
6.1.3.2	Kainic Acid	98
6.1.3.3	Fluorocitrate	98
6.1.3.4	Citrate	99
6.1.4	Impedance recording	99
6.1.4.1	Cell-swelling theory	99
6.1.4.2	Impedance hardware	100
6.1.5	Recorded signals	104
6.1.6	Post processing of impedance data	106
6.2	Human data acquisition	109
6.2.1	Preparation	109
6.2.2	Experimental procedure	111
6.2.3	Data processing	112
6.2.4	Experiment design	113
6.3	Paralysed human data collection	113
6.3.1	Preparation	114
6.3.2	Experimental procedure	115
6.3.3	Data processing	116
7	Developed algorithms	117
7.1	Entropy	117
7.2	Mutual Information	118
7.3	Correlation Dimension	119
7.4	Linear Discriminant	120
7.5	Surrogate analysis	120
7.5.1	Testing	121
7.5.2	Analysis	123
7.5.3	Discussion	128
8	Experiment: Nonlinear analysis of rat epileptiform EEG	129
8.1	Objective	130
8.2	Data analysis	130
8.2.1	Display	132
8.3	Results	133
8.3.1	Analysis of EEG	140
8.3.2	EEG seizure activity	140
8.3.3	Inter-ictal periods and seizure anticipation	141
8.3.4	Cell swelling	141
8.3.5	Spreading depression	142
8.4	Discussion	142

8.4.1	Pre-ictal EEG	142
8.4.2	Ictal EEG	142
8.4.3	Seeking a predictable relationship between EEG and impedance . . .	143
8.4.4	Spreading depression	143
8.4.5	Surrogate data	144
8.5	Conclusion	144
9	Experiment: Rat impedance modelling during epileptiform events	146
9.1	Hypothesis	146
9.2	Introduction	147
9.3	Method	147
9.3.1	Data Analysis	147
9.3.2	Data assembly	147
9.3.3	Neural Network modelling	149
9.4	Results	150
9.5	Discussion	151
9.5.1	EEG vs Impedance	152
9.5.2	Further work	152
9.6	Conclusions	153
10	Experiment: Human data analysis	154
10.1	Hypothesis	154
10.2	Introduction	155
10.3	Methods	155
10.3.1	Data Analysis	155
10.3.2	Data display methods	155
10.4	Results	156
10.4.1	Local MI variations (adjacent electrode pairs)	156
10.4.2	Mutual information analysis by brain region	156
10.5	Discussion	158
10.5.1	Local MI variation	158
10.5.2	Regional MI variation	158
10.5.3	Group Classification (diagnosis of PGE)	160
10.5.4	Task-classification	160
10.6	Conclusion	160
11	Experiment: Diagnosis of PGE in human subjects	162
11.1	Hypothesis	163
11.2	Introduction	163
11.3	Method	163
11.3.1	Data collection	163

11.3.2	Data organisation	164
11.3.3	Classification	165
11.4	Results	167
11.4.1	Task-based breakdown of classification	168
11.4.2	Variation: Subject-based instances	169
11.5	Discussion	170
11.5.1	Which variables are useful for classification?	171
11.5.1.1	Laplacian analysis	171
11.5.1.2	Useful Tasks	171
11.5.2	Implications for the application of this method in a clinical setting	171
11.5.3	Task performance variations	172
11.5.4	Future directions	172
11.6	Conclusion	173
12	Experiment: Classification of mental task in paralysed human subject	174
12.1	Hypothesis	174
12.2	Introduction	175
12.3	Method	175
12.3.1	Data analysis	175
12.3.2	Data collation	175
12.3.3	Assembly of the training matrix	177
12.3.4	Classification	177
12.4	Results	178
12.5	Discussion	180
12.5.1	Limitations in the data	180
12.5.2	Statistical differences between the tasks	180
12.5.3	EMG artifact in EEG	181
12.6	Conclusion	181
13	Conclusions	182
13.1	Thesis summary	182
13.2	Comments on the analyses	184
13.3	Future work	185
13.3.1	Diagnosis of PGE from non-ictal EEG	185
13.3.2	Analysis of rat data	185
13.3.3	Classification of mental task in paralysed person	186
13.4	Main contributions	187
13.5	Conclusion	188

A Appendices	189
A.1 Appendix A: Additional Figures	190
A.1.1 Regional differences for entropy and correlation dimension analysis . .	190
A.1.2 Mutual information analysis by brain region, PGE vs control	191
A.2 Appendix B: Neural network choices	195
A.3 Appendix C: Human PGE diagnosis	196
A.3.1 Selection of useful variables	196
A.3.2 How many variables should be retained?	197
A.3.3 The importance of being earnest about data separation	198
A.3.4 Averaging multiple neural network simulations	199

List of Figures

1.1	Human multichannel EEG showing seizure	10
2.1	Difference between spectrographic analysis (left) and Wavelet analysis (right)	18
2.2	Wavelet analysis of single recording during seizure-like activity .	20
2.3	Electromagnetic field generated by a pyramidal cell	29
3.1	An example of a chaotic system exhibiting bifurcations	39
3.2	Time series data of Lorenz model	41
3.3	Phase-space representation of Lorenz model	42
3.4	Embedding dimension illustration	44
3.5	Entropy as a function of probability of outcome	48
3.6	A Venn diagram, illustrating Mutual Information	49
3.7	Rat EEG showing seizure	57
4.1	Linear discriminant analysis – example data	60
4.2	Two figures illustrating why there is more classification error as the number of classes increases	62
4.3	The structure of an artificial neuron	66
4.4	A small example circuit, showing the internal connections of a group of neurons	67
4.5	A series configuration for use of a trained neural network	71
4.6	The same circuit, but in a parallel configuration to facilitate network training	72

5.1	A typical representation of a neuron	74
5.2	An action potential	76
5.3	An electrical model for brain impedance	90
6.1	EEG electrodes and impedance drivers	100
6.2	The path of low-frequency current through neural tissue	101
6.3	Swollen cells, and their effect on low-frequency current passage	102
6.4	Path of high-frequency current	102
6.5	Impedance recording experimental setup	104
6.6	The rat head, experimental setup	105
6.7	An example of wrapped data – sampled at 2kHz	107
6.8	Wrapped data and unwrapped data – sampled at 2kHz	108
6.9	Real discontinuities	109
7.1	Schematic illustration of the brain, indicating the brain regions	119
7.2	Autocorrelation function of EEG and three different surrogates	121
7.3	FFT of three different surrogates, and EEG	122
7.4	Amplitude histogram of surrogates vs EEG	122
7.5	MI of surrogates (red; $\pm\sigma$ blue), vs MI of EEG (black).	124
7.6	Entropy of EEG (red; $\pm\sigma$ blue) vs surrogate (black)	125
7.7	Correlation dimension of EEG (red; $\pm\sigma$ blue) vs surrogate (black)	125
7.8	Innovation of MI of EEG with respect to MI of surrogates	126
7.9	Innovation of entropy	127
7.10	Innovation of correlation dimension	127
8.1	Rat display GUI, kainic acid experiment	132
8.2	SD98, kainic acid	134
8.3	SD64 kainic	135
8.4	SD 64 zoom kainic	136
8.5	SD94 - picrotoxin	137
8.6	SD79 - fluorocitrate	138

8.7 SD79 - fluorocitrate, detail	139
9.1 Neural net training data organisation	148
9.2 Estimate of impedance – familiar data	150
9.3 Predictive test of impedance estimate	151
10.1 Local MI variations during tasks, control (red) and PGE (blue) subjects.	157
10.2 Regional MI variations during tasks, left frontal vs other brain regions, control (red) and PGE (blue) subjects.	159
11.1 Summary of analysis of EEG leading to an estimate of the sub- jects' class	166
11.2 Subject classification based on all tasks. XY plot (left) and ROC plot (right)	167
11.3 Subject classification based on eyesclosed (first) task only, XY plot (left) and ROC plot (right)	168
11.4 Subject classification based on eyesclosed and eyesopen tasks only, XY plot (left) and ROC plot (right)	168
11.5 Subject classification based on first four tasks, XY plot (left) and ROC plot (right)	169
11.6 200 best variables	170
12.1 Graph of EEG spectral energy of paralysed and unparalysed subjects	176
A.1 Regional entropy analysis	190
A.2 Correlation dimension analysis	191
A.3 Task-based mean and SEM of MI of Right Frontal vs other brain regions, PGE vs control	191
A.4 Task-based mean and SEM of MI of Right Occipital vs other brain regions, PGE vs control	192

A.5	Task-based mean and SEM of MI of Left Occipital vs other brain regions, PGE vs control	192
A.6	Task-based mean and SEM of MI of Right Temporal vs other brain regions, PGE vs control	193
A.7	Task-based mean and SEM of MI of Left Temporal vs other brain regions, PGE vs control	193
A.8	Task-based mean and SEM of MI of Right Occipital vs other brain regions, PGE vs control	194
A.9	Task-based mean and SEM of MI of Left Occipital vs other brain regions, PGE vs control	194

Acknowledgements

I owe many people thanks.

Firstly, thank you Selina, for your support and encouragement throughout my degree. Thanks, mum and dad, for always indulging my curiosity and interest in science. Thank you Carla and Brian for your help and support. Your proof-reading was very valuable, Brian – cheers. To my friends for their help and support – not only in the context of this thesis, but in all things. In particular, thank you Jesse, Hannah, Owen, Martin and Gez.

I think there are few people who have been as lucky as I to have such fantastic supervisors, thank you Kenneth and John. Your advice and support has been fantastic, and it has been a pleasure to work and socialise with you. Similarly, the assistance and friendship provided by my colleagues in the human lab and the rat lab has been wonderful. Marita, Emma, Trent, Sean and Steve, your advice, help and humour has been great. Thanks also to my Swedish colleagues, especially Torsten and Michal for facilitating this research. Finally, thank you to all the volunteers who participated in our experiments – your time and effort are very much appreciated.

This work has been supported by a Flinders University Research Scholarship, a Wellcome grant, STINT and NHMRC grants.

Declaration

The work contained herein is my own, unless otherwise stated. This thesis is submitted in total fulfilment of the degree of Doctor of Philosophy at The Flinders University of South Australia.

Ethics approval was sought and obtained for all the data-acquisition experiments described herein.

Chapter 1

Introduction

1.1 What will be in this thesis

This thesis applies mathematical algorithms to the analysis of EEG data. The work is novel partly because of the data being analysed, and partly the approach. There are few groups applying a two-pronged human and animal-model EEG analysis approach that is comparative. Few also record data at the (high) sampling rate employed by us. Furthermore, although the analyses themselves are not novel, the signal-processing approaches are, as far as I have found, not attempted elsewhere. In particular the combination of non-linear quantification of EEG and neural network training and modelling is a little-investigated area.

Because this thesis is concerned with the application of well-known non-linear algorithms, it is useful to discuss what is available (the signal processor's *toolkit*), what was chosen, and why. It begins with an examination of signal processing. Linear analyses are described first, followed by nonlinear, as well as an explanation of the meaning of the terms *linear* and *nonlinear*. Both chapters 2 and 3 relate the analyses to EEG and epilepsy research, so that the practical aspects are apparent and the limitations clear.

Chapter 4 discusses classification, and examines several methods as well as their relative merits and disadvantages.

Chapter 5 introduces some physiological detail about the brain, including an examination of the important structures. It discusses the origin and limitations of EEG, and some means of analysing EEG data. It then discusses the current state of epilepsy research and has a brief examination of general signal-processing-based neuroscientific research.

The analyses contained in this thesis were performed on data collected in acquisition experiments in two epilepsy research laboratories in the Flinders Medical Centre, South Australia. I was not involved in the conception and goal-design stages of these acquisition experiments. However, I was involved with the execution of some of the rat experiments, and in the group discussing future developments and directions for the experiments. I was also involved in some of the fine-detail planning in the human experiments. The work in this thesis was performed on data derived from these experiments after the fact, and represents an extension of the original methodology around which the acquisition experiments were designed.

I have made a conceptual separation between the data acquisition experiments (where the data were acquired, but for a different purpose than that discussed in this document) and my signal-processing experiments (where I analysed the data and attempted to draw conclusions). For this reason, I have included a separate chapter on data acquisition. This helped prevent repetition (because I made multiple uses of the data from the data-acquisition experiments), but also helps to maintain the aforementioned conceptual separation between the data-acquisition and the data-analysis experiments. The three data acquisition experiments are discussed in the three sections of Chapter 6. Similarly, I used similar tools for the analysis of data in all of my experiments, and these tools are discussed in Chapter 7, which provides detail of those algorithms and processes.

Chapters 8 through 12 detail my analysis experiments. These are all de-

signed to explore the possibilities of using nonlinear analyses, modelling and classification to analyse EEG data.

1.2 Epilepsy and EEG analysis

Epilepsy is a poorly understood neurological condition that affects approximately 8 people in one thousand. Symptoms of the disease vary, but for many, epilepsy is a serious impingement on their quality of life. Poor understanding means that drugs prescribed to alleviate symptoms tend to be rather crude and ineffective – in fact, of the sufferers treated, 20% are not significantly assisted by drug therapy [48]. There are many different types of epilepsy, and all are characterised by seizures: periods of abnormal brain activity resulting in altered consciousness or loss of consciousness.

Electroencephalogram (EEG) is probably the most commonly applied neurological tool. EEG is a recording of electrical signals that the brain produces during its operation, and is thought to reflect the internal operation of the brain. The transition from normal EEG to a seizure, in a subject with primary generalised epilepsy, is shown in figure 1.1.

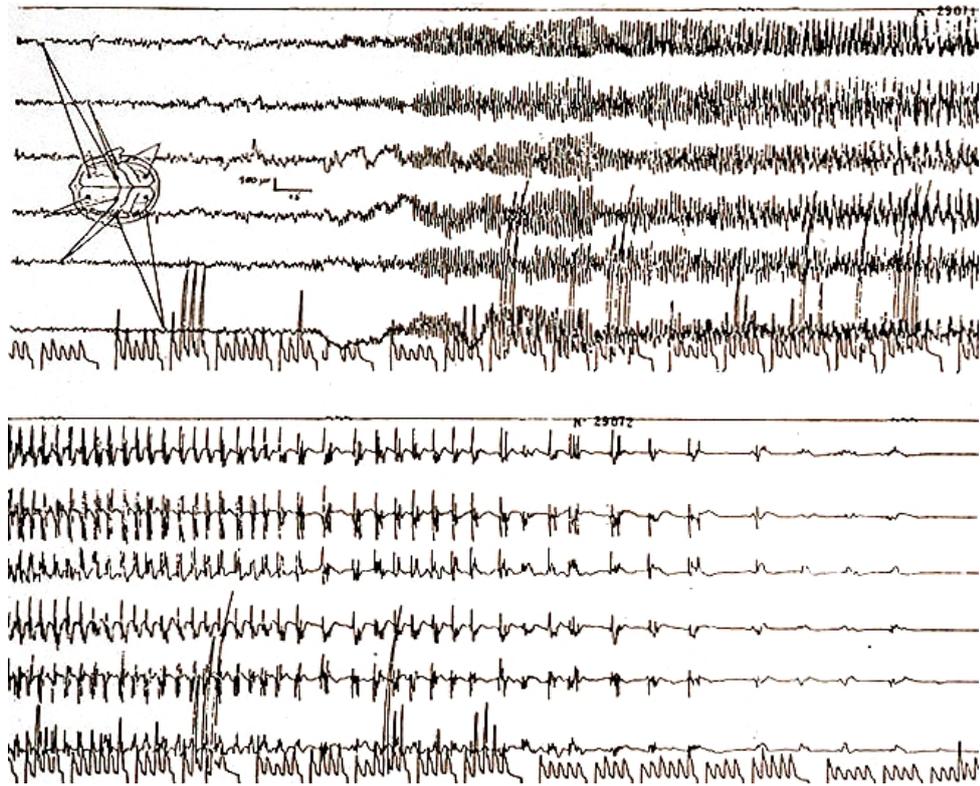


Figure 1.1: Human multichannel EEG showing seizure

This is an EEG trace, recorded by Gastaut and Broughton in 1976. It shows 6 channels of EEG in two rows. In the first row, there is a progression from relatively normal EEG through to a seizure, which calms in the second row.

The type of epilepsy shown in figure 1.1 is generalised, meaning that it emerges, simultaneously, across the entire brain¹. The causes of generalised seizures are not understood, nor are the processes by which they begin and end. We do know of a genetic predisposition to such seizures (specific genes are associated with ion channel abnormalities), but this does not completely describe why some people are susceptible to such seizures while others are not, nor why some people respond well to medication while others do not.

The examination of EEG by a qualified and experienced technician or clinician is a very useful practice, and one that is likely to continue for some time. Such examination is able to uncover previously unknown pathologies, or allow a specific diagnosis to be made. However, this process is subjective, and

¹There is another type of seizure, called a focal seizure, which looks similar, but affects only one area within the brain. A focal seizure may generalise (become a whole-brain event) at which point it looks similar to the generalised seizure shown in figure 1.1.

is not systematic nor necessarily repeatable [61]. Applying signal processing techniques (linear and nonlinear) allow the visualisation and quantification of different aspects of the EEG data. These are generally repeatable, automated and systematic, which are all advantages over the visual inspection of EEG.

For the purpose of better understanding epilepsy, it would be helpful to better understand the relationship between the disease, the symptoms, and the EEG. This is something that is needed because our understanding of the methods of brain function is rudimentary – the mechanisms of memory storage and retrieval, information processing, integration, decision making, high-level pattern recognition, and many others, are as yet poorly understood. What we do have is an understanding of many low-level mechanisms of the brain (the operation of neurons and simple neural circuits) and knowledge of the high-level function and operation of the *mind* (as studied in psychology, sociology, and other areas). What is lacking is a unifying theory by which we can move between these two worlds, the micro and the macro, and understand the systems' levels as greater and greater levels of complexity are reached.

One method of gaining insight into a complex system such as the brain is to examine it when its behaviour is dysfunctional, by highlighting aspects of behaviour that are otherwise occluded – it can act as an “extra data point.” Because of this, a comparison between subjects with and without epilepsy can potentially provide useful data about brain function.

The process of analysing the brain in this way is made difficult, because our tools for examining functions of the living brain are crude. The principal tool is EEG, and this has limited spatial resolution and specificity (sections 2.2.8 and 5.2.1). Also, our ability to understand and analyse complex systems in general is limited. Emergent characteristics of complex systems are poorly understood, so the application of reductionist analysis techniques to such systems tends to result in a partial and fragmented comprehension. This is partly because our tools for the analysis of nonlinear systems are somewhat rudimentary,

and the collaboration between experts in nonlinear theory and neuroscience is somewhat immature (although great strides have been made in the last decade). Despite these shortcomings, there has been much success in the analysis of EEG using both linear and nonlinear tools. This is largely because it seems as though EEG is reflective of brain state and activity, despite its coarse-grained nature.

There is a long and successful tradition, probably 50 years, of linear analysis of EEG. Increasingly, we are becoming aware of some of the limitations of linear analysis – these are chiefly associated with linear analyses’ inherent assumptions and resultant constraints. Nonlinear analysis makes fewer assumptions about the data and the system, and it also provides new tools and perspectives from which to approach signal processing.

1.3 How to analyse signals

Broadly speaking, there are two main groups of analyses. The first are the “linear” analyses. These are analyses which assume that the system under scrutiny obeys the rules of linear systems – a rule called *superposition*, which means that such a system can be expressed as the sum of its parts (more detail in chapter 2).

This is a useful principle, and allows the reductionism that is so important to many areas of scientific research. However, it is now recognised that most physical systems do not perfectly exhibit this characteristic, so that the employment of linear analyses means approximating systems as linear. Often, assumptions such as these do not greatly impact on the effectiveness – for example spectrographic analysis is commonly employed as a means to analyse systems known to exhibit nonlinear behaviour.

A nonlinear system is one which doesn’t exhibit superposition and hence cannot be described by linear equations. The result of this is that nonlinear systems are in general not solvable mathematically, except in particular cases

for particular states of the system. This was particularly the case prior to the modelling afforded by computers. Nonlinear analysis does not make the aforementioned assumptions of superposition, and the result is that the conclusions drawn from such analyses can be more generalisable.

Chapter 2

Linear Systems

This chapter provides an outline of what a linear system is, and how it can be analysed. I then discuss the meaning of data that are descriptive of a system, and some of the caveats for dealing with such data.

2.1 What is a linear system?

Linear systems exhibit the property of *superposition* – a combination of the two properties *additivity* and *homogeneity*¹ (equations 2.1 and 2.2).

$$f(x + y) = f(x) + f(y) \tag{2.1}$$

$$f(\alpha x) = \alpha f(x) \tag{2.2}$$

A system that possesses this quality is able to be disassembled into small pieces, analysed separately, and the understanding gained therefrom can then be reassembled from the various parts and understood holistically. This process is called *reductionism* and has been, perhaps, the keystone of science for the past three to four hundred years.

¹If I drive my car to buy some topsoil, a common practice is for my car to be weighed upon arrival and upon departure, so the quantity of soil is calculated as the difference in these weights. This uses the principle of additivity. The additional weight of the topsoil in my car will cause a performance decrease that is (approximately) proportional. That is the principle of homogeneity.

A good example of this is an ideal billiard table. If a red billiard ball is hit by the white ball, we can separate the momentum vector of the white ball into two orthogonal components, and calculate the effect of these separately. The resultant velocity vector for the red ball is simply the sum of two similar orthogonal vectors. Notice that I described this as an *ideal* billiard table – on an actual table this is only approximately true. This is because tiny non-idealities can affect the linearity of the system. For example:

- The balls do not behave as perfect springs – there is hysteresis in their elasticity
- Friction in the felt of the table
- Air currents over the table

These non-idealities are very small, and in a good table do not significantly affect the linearity of the system. However, billiards is a *very* carefully controlled system – the billiards environment is designed to be consistent: the balls are near-perfect spheres, the table is very even, the bumpers are very consistent when a ball collides. Contrast billiards with golf (a game using similar principles) – there are many more parameters that cannot be nicely approximated as linear. Despite this, if one hits a golf ball at a certain velocity, at a certain angle, in a certain direction, and there is a cross-wind of a certain speed then we can use Newton’s (linear) laws of motion to calculate where the ball will land *with a reasonable degree of accuracy*. This shows that even for systems that only *approximate* linearity, linear methods of analysis are very powerful and useful.

2.2 Linear analyses

This section discusses some commonly-used linear analyses, and relates them to the field of neuroscientific research.

2.2.1 Autocorrelation function

Equation 2.3 describes the autocorrelation function

$$\hat{\phi}_{xx}[m] = \frac{1}{Q} \sum_{n=0}^{Q-|m|-1} x[n] x[n+|m|] \quad (2.3)$$

The autocorrelation function is an estimate of how well a discrete-time signal matches itself at different time displacements. The autocorrelation function reflects periodicity in the signal, as well as sameness across the signal boundary (although only when the calculation is circular). Equation 2.3 shows a comparison of a signal with a delayed version of itself (auto-correlation), rather than two separate signals (cross-correlation). In the case of EEG analysis, both auto- and cross-correlation can be useful. Autocorrelation can reveal information about an electrode's patterns over time, and correlation can reveal shared patterns between electrodes. The autocorrelation function is related to the power spectral density by the Fourier transform (section 2.2.2) operator.

Larossa et al [38], used a combination of autocorrelation and spectral analysis (section 2.2.2) to identify cyclic behaviour involved in spreading depression (section 5.4.3) activity of Sprague-Dawley rats. They concluded that glial- and neuronal-glial interactions were involved with the synchronisation occurring during seizures and spreading depression. However, as the recordings used brain slices and micro-electrodes, the signal analysis tools used in the study are not directly applicable to our *in vivo* experiments. Autocorrelation has also been used to compare spike-wave discharges within a seizure [52],

2.2.2 Fourier transform

The Fourier transform finds an alternative base in which to represent a time-series signal. This base is referred to as the frequency domain, and allows one to view a time-series signal as the sum of its constituent frequency components and their strength across all time. To find the power contained in each fre-

quency we find the Fourier transform of the autocorrelation function. Strictly speaking, the evaluation of a Fourier transform of a non-stationary system (a system whose statistical properties change with time – see section 2.3.1) is meaningless, because of frequency smearing caused by the system’s changing properties.

One manner in which we can circumvent this is to use spectrograms (also called a “Short term Fourier transform”), which uses a sliding window to selectively perform Fourier analysis on small time segments. If a small number of samples are examined (representing a short duration EEG) then we can generally say that the signal is approximately stationary for this short duration. Thus, a spectrogram is really a map of how the power of the constituent frequencies change with time.

Fourier transforms and spectrographic analysis have been used extensively in the study of EEG, with good results. Our lab has been using spectrographic analysis since 1999, and has uncovered many new findings (eg. [47, 46, 45, 54]). As can be seen from the changes evident through the years, we have attempted to diversify our analysis of EEG and the experiments. Spectrographic analysis is, and will likely remain for some time, the default analysis. This is largely because it is well known, easy to evaluate and easy to interpret – it also correlates with EEG in a manner that is easy to grasp, allowing inferences made from examination of the frequency domain to be easily translated back to the time domain.

Analysis of EEG with a view to examining brain “chirps” [66] has also been undertaken. Brain chirps are brief periods of synchronised neuronal activity occurring at a fundamental frequency and also at higher harmonics. Schiff et al describe chirps as a “highly sensitive and specific marker for epileptic seizure activity” and propose seizure detection using filters matched to the chirp. There are several problems with the use of this procedure in a clinical setting – the primary issue being that it requires brain surface electrodes as

the chirps are not detectable through the skull.

Fourier analysis has limitations (figure 2.1). One issue is that the window to be analysed can only be made as “narrow” as the wavelength of the lowest frequency to be examined². However, a sample size of one (wavelength) means there is likely to be a significant error. This is because as we decrease the width of the windowing function (and thus give the spectra a finer time resolution), we have fewer data with which to calculate the spectra, and so the error of calculation increases. There is a trade-off between accuracy of analysis and temporal resolution [4], and our temporal resolution (of all frequencies) is limited by the length of the longest wavelength we wish to resolve.

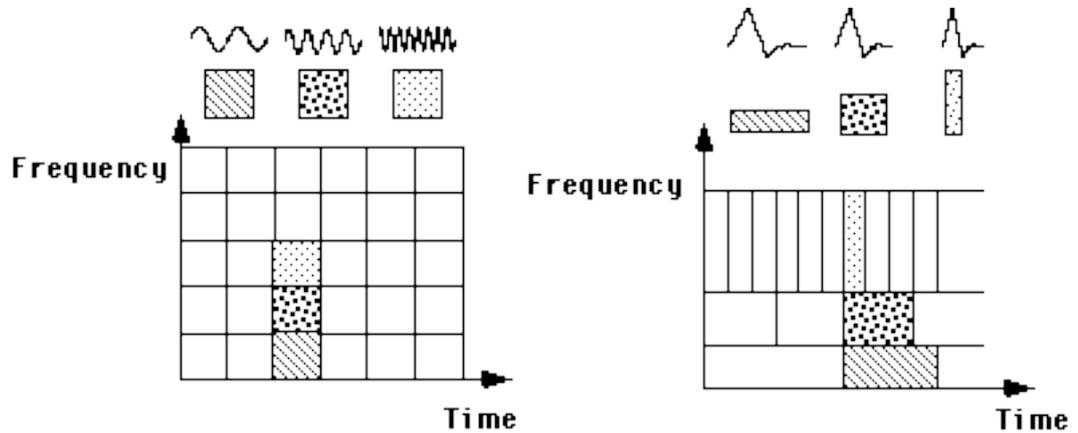


Figure 2.1: Difference between spectrographic analysis (left) and Wavelet analysis (right)

http://www.amara.com/IEEEwave/IW_wave-vs_four.html

The two diagrams in this figure show the differences between the fast Fourier Transform (FFT) and the discrete wavelet transform (DWT). The size of the windows in the FFT is constant across frequencies, so that the window-size-requirements imposed by the lowest frequency (the window needs to be at least as long as the longest wavelength) are imposed on every higher frequency as well. Contrast this with the wavelet analysis (right) where the smallest required window is used for each frequency band, thus maximising temporal resolution of the analysis.

²To resolve a frequency, we must analyse at least a whole wavelength at that frequency.

2.2.3 Wavelet analysis

Wavelet transforms represent another form of time-frequency analysis, and allow a representation of a time-series signal as a summation of wavelets. This is similar to a Fourier transform, except that wavelets are localised in both time and frequency, whereas Fourier transforms are specific in frequency only [4]. What this amounts to is that a wavelet transform provides information about frequency components at specific times in the evolution of the signal – in some ways this is similar to a spectrogram.

The Heisenberg-like uncertainty by which spectrographic analysis is constrained also exists for wavelets – but where Fourier analysis has a fixed window width across all frequencies, wavelets allow the uncertainty to be minimised for each frequency band (the temporal resolution of the wavelet transform varies across the frequency spectrum, figure 2.1). In particular, when the analysed signal is non-stationary, wavelet analysis allows better control of the trade-off between the temporal resolution and frequency domain accuracy [53].

There are many different shaped wavelets; Latka et al [39] use a "Mexican Hat"-style wavelet, which they describe as being "particularly suitable for studying epileptic events". Using this, they are able to differentiate between spikes in the EEG, and noise artifacts.

In practice, wavelet transforms are often used in a manner quite similar to spectrograms. A good example of this is Saab et al [62], who use a 5-level Daubechies wavelet frequency analysis, followed by a Bayesian estimation of the seizure likelihood. The seizure and non-seizure data were manually classified, and each sample was characterised by 3 parameters, the distributions of which represented the a priori probabilities. The likelihood that a given segment of EEG is seizure was then estimated using Bayes' formula (equation 2.4)

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (2.4)$$

Thus, this analysis is very computationally light – the only analysis that

occurs is the wavelet analysis, since the *a priori* probabilities (the $P(B|A)$ component of Bayes' equation) are calculated before the experiment. The inherent simplicity results in a limit to the sensitivity of the analysis, since any relationship too complicated to be expressed in a relative 3-dimensional histogram will be ignored.

Nyikos et al [53] have extensively employed wavelet analysis to produce time-frequency visualisations of EEG. Epileptiform activity was induced in brain slices (from hippocampus, entorhinal and perirhinal cortex) by perfusion with an elevated K^+ and absent Mg^{2+} ionic bath. This activity was recorded via micro and clamp electrodes for offline analysis.

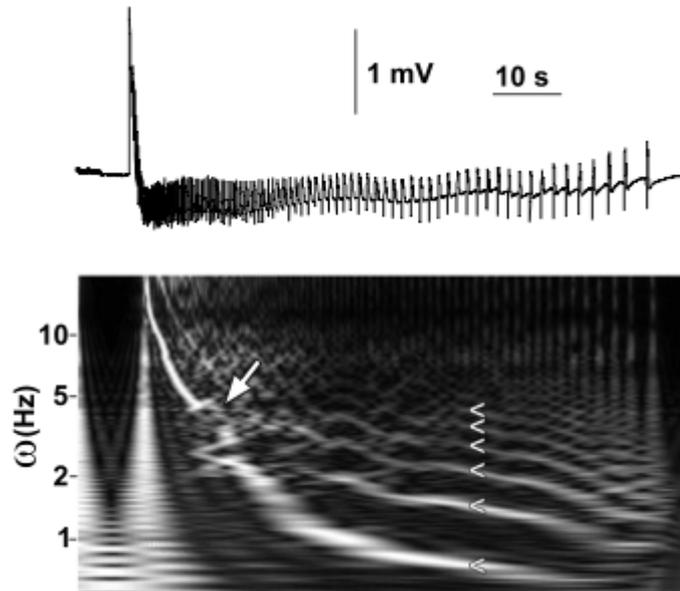


Figure 2.2: Wavelet analysis of single recording during seizure-like activity. These data were recorded from a single cell in a slice of brain tissue exhibiting epileptiform activity due to immersion in an ionic bath (elevated K^+ and absent Mg^{2+}). The image was produced by Nyikos, et al, and sourced from their article [53].

Using time-frequency wavelet analysis, they were able to examine the process of the seizure (shown in figure 2.2), and concluded that initial “precise temporal synchrony is gradually destroyed during ictal events”. The figure shows the greater temporal resolution at the higher frequencies, which would not be the case using spectrographic analysis – making such data much more

difficult to visualise.

2.2.4 Eigenvalues/vectors

Eigenvalues and eigenvectors are another means of representing a matrix in a new form. If we look at a matrix, \mathbf{A} , then an eigenvector of \mathbf{A} , \bar{x} , has the property that $A\bar{x} = c\bar{x}$ where c is the eigenvalue of \mathbf{A} corresponding to the particular eigenvector \bar{x} [35]. They provide a way of specifying one set of basis functions for a system (described in a matrix). A set of basis functions completely represent a system, and this relates to transform like Fourier and wavelet. For example, a signal is able to be represented as a summation of many frequencies (using a Fourier transform). In this case, the matrix of eigenvectors is the Fourier transform matrix, which is merely sampled sinusoids at different frequencies. This transforms the signal from the time basis into the frequency basis (domain). Alternatively simply putting a wavelet into each row of the eigenvector matrix (i.e. wavelets with different scales at different time displacements) will implement a wavelet transform [82]. This property is not unique to the wavelet or Fourier transform – any basis set can be concatenated into a matrix to linearly transform the data into a new domain.

2.2.5 Principal component analysis

Principal component analysis (PCA) is one method of *blind signal separation* (BSS), a process of estimating a set of original signals from a set of signals mixed through a linear combination.

When talking about BSS, it is common to refer to a cocktail party, where there are several people speaking simultaneously, and at least as many microphones scattered around the room recording the sound. Each microphone will therefore record a different mix of the six voices, depending on their relative proximity to the speakers. Using ICA and BSS, it is possible to use these recorded channels of mixed data to produce channels of unmixed data, where

each channel contains one speaker only. This process becomes more difficult as the recorded channels become more similar (the microphones and speakers are closer together).

PCA finds common signal components in discrete signals. It can be used to detect structural similarities between signals [1], but it also allows one to construct sets of data that are uncorrelated. Often this is a means of reducing the volume of data, by reducing or eliminating redundancy between channels. It differs from Fourier analysis in that Fourier analysis decomposes the signal into a series of pure sinusoids, whereas PCA reduces it into a set such that the first principal component describes as much of the signal as possible, the second describes as much of the remaining signal as possible, etc. In terms of the phase space, this means that all the vectors are orthogonal, which is different to Independent Components Analysis (see 2.2.6), which is concerned with finding independent components.

The method of finding successive principal components is to find eigenvectors. These correspond to the direction of new axes, onto which we will project the n -dimensional data. The eigenvectors are mutually orthogonal. The eigenvalues that correspond to these eigenvectors indicate how powerful each particular vector, as a direction for a new axis, is in describing the behaviour of the system. Each successive vector describes less and less of the signals (because the first vector is that which accounts for the maximum possible variation in all signals). Thus, we can decide when we are satisfied with the *approximation* of the signals that this method yields [1]. This can be viewed as a series approximation.

PCA is useful for providing a data *summary*, particularly when two data sets describe the same (or a similar) system *and there are a different number of data* in each set. Such a scenario could make comparisons difficult, but by comparing the n most important principal components, an absolute comparison can still be made. This is the procedure used by Kayser and Tenke

[31] to examine the differences between high- and low-density skull electrode configurations and their effect on surface Laplacian estimation (section 2.2.8). They recorded EEG from 17 adults during odd-ball (an example is in section 6.2.2) tasks using 129 electrodes, and then sub-sampled the electrodes to produce an alternate data set with 31 electrodes. They estimated the current source density (CSD) using surface Laplacian (section 2.2.8) for both data sets (high- and low-density electrode configurations), and used PCA to provide a summary of the CSD estimate for each data set for comparison. Using this technique, they were able to show that there was a strong correlation between the CSD estimate for the two data sets during evoked response tasks (ERPs), meaning that there is only a small gain to be made by increasing the electrode density of the EEG recording. They are, however, careful to clarify that these results may vary even using other ERP paradigms, because it “involves rather subtle and highly-specific topographic effects” [31].

2.2.6 Independent components analysis

Independent components analysis (ICA) is another form of blind signal separation. ICA differs from PCA in that PCA attempts to form a new basis using as few components as possible, whereas ICA attempts to identify the source signals from which the mixed data were linearly combined. In other words, ICA seeks a new basis comprised of independent signals, whereas PCA seeks orthogonal signals. In the context of EEG, the mixed data are the electrical signals measured on the skull. These are a mixed and smeared version of the source signals produced by radial (section 2.2.8) dipoles within the brain.

ICA operates under the assumption that the source signals are linearly mixed summations of statistically independent, discrete sources. Generally, one of the many solutions³ to ICA can produce the original signals, but practically (in the context of EEG) what is found is not a compact brain region

³There are several ambiguities inherent in the separation of signals using ICA. This includes source specificity – i.e. which source belongs where.

acting as a source, but rather a distributed network responsible for a particular signal component. In order to be able to separate source signals, one must assume that they are statistically independent, and that they are non-Gaussian. Thus separation becomes a process of minimising mutual information between channels.

Mathematically, BSS can be thought of as follows [17]:

$$\mathbf{x} = A\mathbf{s}$$

where \mathbf{x} are the EEG data (of size m channels by n samples), A is the mixing matrix (of size m by p) and \mathbf{s} are the source signals (of size p sources by n samples). Since the only observable part of this is \mathbf{x} , then we need to estimate A so that we can find $\mathbf{s} = A^{-1}\mathbf{x}$. It is the process of estimating A that varies between ICA/BSS implementations.

Using blind signal separation to analyse EEG could pose problems because of source ambiguity. This means that, because the unmixing matrix estimate is based only on the goal of independent source vectors, there is no spatial information retained and one cannot know the location of the source vector. In the above equation, ICA will not actually estimate A^{-1} , but a permutation and rescaling of it. Source ambiguity can affect calculations and conclusions, because we won't know which part of the brain a particular signal comes from.

Similar to ICA is *blind deconvolution*, which can be used to undo the effects of an unknown filter. If one assumes that the original signal was white (the signal's probability density function⁴ is uniform), then it becomes a process of removing correlations across time (i.e. whitening the signal). This process, as well as that of ICA, can be viewed as a redundancy reduction (in the case of BSS, reducing mutual information between outputs and in the case of blind

⁴The calculation of a histogram of a random variable's amplitudes across all time will provide an indication of the likelihood of observing certain values. If the limit of the histogram is taken so that the bins become arbitrarily small, one obtains the probability density function (PDF).

deconvolution, whitening) [6].

Bell, et al [6] have refined and developed some techniques of blind separation and blind deconvolution, specifically for application to EEG analysis. For the separation of sources, they use a criterion of information maximisation. They claim that by maximising the entropy of the outputs, the mutual information between the outputs and the inputs is also maximised. They pass the input through a nonlinear function (typically some variant of a sigmoidal function) which levels the probability density function of the signal. This means that the probability of obtaining any value is more even – it normalises the probability density function, which increases the information transfer (refer to [70] for more detail).

A common application of ICA in EEG analysis is the removal of muscle and ocular artifact. This is an important step, as it is increasingly thought that muscle artifact (EMG⁵) represents a much larger corruption of EEG than previously thought [83]. Such is the approach taken by Frank and Frishkoff [17], when attempting to remove eye blinks and other ocular activity from EEG. A blink template was constructed, which was tailored to each subject. There were three criteria by which blinks were identified. The first was that ocular activity would have a polarity inversion at the level of the eye, the second was that the signal would correlate to the blink template at a threshold greater than 0.85, and the third was a correlation between the spectrum of the component and a known blink sample. They tested their algorithm on EEG with simulated blinks introduced, and concluded that it was, in general, successful in removing the blink without removing non-ocular EEG from the data set. They found that when “blink-splitting” (where the activity from a blink was split into two source components) then they would tend to also remove non-ocular EEG – lowering the correlation between the filtered and original (uncorrupted) data.

⁵Electromyogram (EMG) data are a recording of muscle electrical activity, in the same way as an EEG is a recording of brain electrical activity. EMG signals are very powerful, relative to EEG signals, and so represent a common artifact in EEG recordings.

On real data, they found that often blinks would occupy multiple components, and that removal of only one was insufficient, whereas removal of two would improve EEG in frontal areas, but degrade it in posterior areas. This clearly shows that there are many practical issues associated with the application of ICA to real-world data.

2.2.7 Autoregressive Analysis

When examining stochastic signals, we are making guesses about the underlying system that produces the signal. Generally, this system is treated as a black-box. An effective way of learning about the system is to model it, and autoregressive (AR) analysis involves the use of an adapting model of the system. The state space (generally interchangeable with "phase space," see 3.3.1) variables of the system are approximated, and only the noisy component (the *innovation* – the signal component not described by the model) is stored. This is also used to refine the attributes of the model to more accurately describe the system.

The basic theory of AR analysis is that the best way to predict what a system will do in the future is to look at what the system did in the past. Hence, we make a prediction of future samples as a weighted sum of previous samples.

If $x[1..n]$ is a series of n samples measured from a system with an unknown transfer function, then we can make a prediction of $x[n+1]$, as shown in equation 2.5

$$x[n+1] = \sum_{k=1}^p a_k x[n-k] + e[n] \quad (2.5)$$

This formula states that the next sample we measure, $x[n+1]$, will be approximately equal to the weighted sum of previous samples. The number of previous samples, p , we consider is called the *order* of our AR analysis, and will be reflected in the accuracy of our prediction. The term $e[n]$ refers to the discrepancy between our predicted value and the real value of $x[n+1]$, and is

called the innovation of the system: that part of the system's behaviour that we don't understand and cannot predict.

The analysis discussed here makes the assumption that linear modelling can describe our system. One can also perform nonlinear AR (NLAR or NAR) where we consider other relationships between historical samples, in addition to the weighted sum. Such a model is described by equation 2.6.

$$x[n+1] = \sum_{k=1}^p a_k x[n-k] + b_k x[n-k-u] x[n-k-v] + e[n] \quad (2.6)$$

Note that this only examines a single nonlinear relationship – namely a weighted sum of the product of two points, separated in time by a fixed amount ($u - v$). If we wish to consider all possible relationships (many values for u and v), up to second order⁶ (quadratic), then we would have 3 sigmas (one for the linear relations, one for sums of products, and one for the squared components), and many more parameters to each summation. As the order of the analysis increases, it is clear that the complexity of the calculation increases exponentially. Because of this, nonlinear terms are generally added sparingly to AR models.

AR models have been used extensively for EEG analysis, with several objectives. One use has been to estimate the direction of flow of influence of rhythms between regions of the brain [50], while modelling using nonlinear autoregressive analysis has been used to model seizure mechanisms in an attempt to identify common structural elements between different seizure types [65]. Some studies have also used AR analysis to estimate the power spectrum [3].

Probably the biggest limitation of using AR analysis occurs when attempting to model a complex, poorly-understood system. This is because the experimenter cannot know, prior to the modelling, which parameters to include in the model, and which to exclude. In an attempt to circumvent this, con-

⁶Order can mean two things in this context. Firstly it can mean the number of AR parameters, and secondly it can refer to the mathematical order of the relationship that the parameters describe. In this case, I mean the latter.

figurations such as the NLAR (discussed above) have been made. Further generalisations (attempts to make the models robust to varied kinds of system behaviour) have also been introduced – such as autoregressive integrated moving average (ARIMA) and many others. However the fundamental postulation of AR (that a sufficiently sized model can approximate any system to any degree of accuracy) is only relevant if such a model can be practically constructed.

2.2.8 Laplacian EEG analysis

It is common for EEG to be analysed by examining not lead potentials, but differences between lead potentials (eg. [71]). This is referred to as the Laplacian, and means that the overlap in measured data between electrodes is minimised. An assessment of the location of sources can be made, and the results adjusted to minimise overlap between channels.

More specifically, a Laplacian analysis will examine a given electrode, and attempt to represent the data from that electrode as a linear combination of the surrounding nearby electrodes. The component of that signal that cannot be estimated from the surrounding electrodes is called the *innovation* (this concept is similar to that of innovation in auto-regressive analysis). It is this signal that is then used to replace the signal from that particular electrode. Laplacian analysis yields EEG with a better spatial resolution, and can be thought of as the application of a high-pass spatial filter [75].

One of the great benefits of Laplacian analysis is that the data it produces are reference-free. EEG is typically recorded relative to a reference electrode (since potentials can only be measured *between* locations). This means that any effects at the site of the reference electrode are imparted to every other electrode in the EEG – an undesirable effect. Because Laplacian analysis finds the component of a signal that is unexplained by nearby electrodes, we are necessarily removing the effect of the reference (because that is the same in

every electrode).

The biophysical theory behind Laplacian analysis is as follows: the largest contributor to EEG are electromagnetic signals caused by radial currents from pyramidal cells (figure 2.3 and section 5.2). The current passes radially from the cell, which can be considered as a radial dipole, and reaches the scalp to air interface. At this point, the current spreads throughout the scalp, perpendicular to the source dipole.

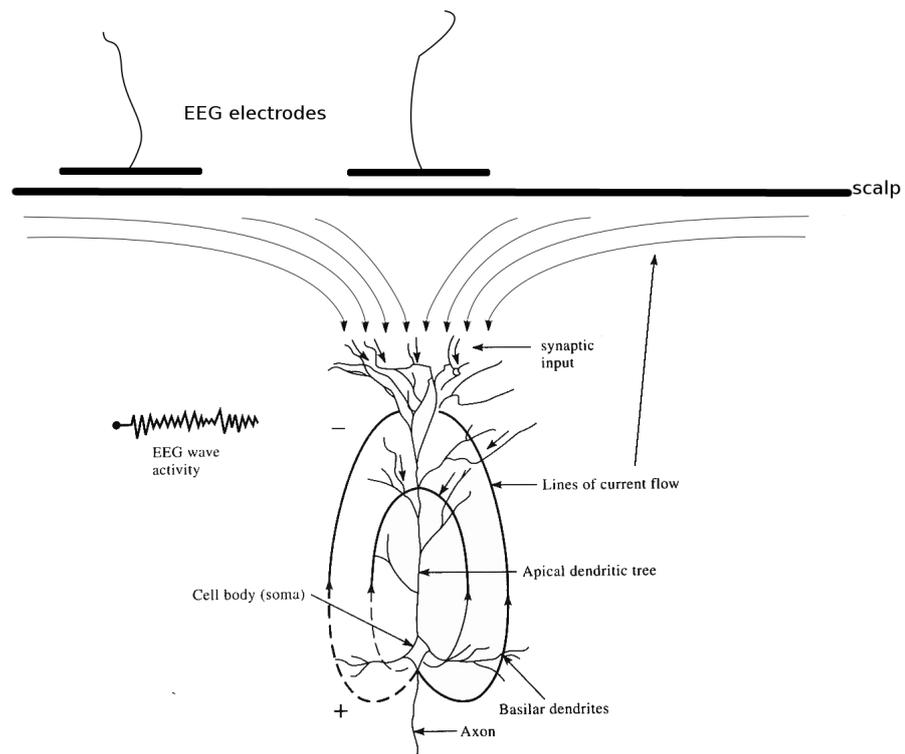


Figure 2.3: Electromagnetic field generated by a pyramidal cell

This figure illustrates the manner in which current from a single cell can affect the input of multiple EEG electrodes. This is a reminder that EEG recordings are coarse-grained. The current passes radially from the pyramidal cell, and reaches the scalp to air interface. At this point, the current spreads throughout the scalp, perpendicular to the source dipole. In this way, the current from each pyramidal cell can affect multiple EEG channels.

This figure was adapted from figure 4.26, page 163, Medical Instrumentation, [81]. (Note that the various elements in this figure are not to scale)

By looking at nearby electrodes, it is possible to estimate the component of the signal that is radial (i.e the pure signal propagating radially – a good approximation for the EEG at the dura [16]) and the component of the signal

that is transverse (i.e the signal that propagates along the surface of the scalp, from where the radial signal reaches the scalp to air boundary). The analysis uses splines to interpolate between electrodes. It is possible to do this to various degrees of accuracy, depending on whether actual electrode coordinates are used (individually tailored), or a set of generic scalp coordinates are used (un-tailored) or a set of spherical coordinates are used. The use of these approximating coordinates allows a great simplification of the mathematics of calculating the splines⁷.

EEG is a crude montage of many signals (section 5.2). The CSD⁸ is defined as the strength of the local, radial currents that sum to produce the EEG [76]. If we examine a potential at a point on the surface of the scalp (approximating as spherical), then we can estimate the potential [15] using equations 2.7 and 2.8

$$V(\mathbf{r}) = c_0 + \sum_{i=1}^N c_i g_m(\mathbf{r}, \mathbf{r}_i) \quad (2.7)$$

where

$$g_m(x) = \frac{1}{4\pi} \sum_{n=1}^{\infty} \frac{2n+1}{(n(n+1))^m} P_n(x) \quad (2.8)$$

Where the function $P_n(x)$ refers to an n th-order Legendre polynomial which form a set of basis functions for the spherical surface. From these equations, we obtain a matrix which describes the way in which the potential that we measure at each electrode site is a linear combination of potentials at other, nearby, electrode sites. This matrix can then be inverted, and used to calculate the scalp surface Laplacian (the estimate of the contribution to the measured potential that is made by the radial currents only).

An important question to examine is: for the purpose of estimating a surface Laplacian of the skull, how do various calculation methods compare? This is tested by Tadonnet et al [75], who examine several different Laplacian estimates. Firstly Hjorth's method, which makes a local estimate of the surface

⁷The process is called the "Laplacian" because of the work done in spherical coordinate systems by Laplace.

⁸Current source density: introduced in section 2.2.8.

Laplacian by calculating the unaccounted-for component of the average potential of the surrounding electrodes. This is compared with a global estimate, using surface spline interpolation – which is the analysis described earlier in this section. Such a global estimate amounts to finding a new basis set for the EEG data. They found that the two methods resulted in equivalent results – the main difference being that Hjorth’s method requires equidistantly spaced adjacent electrodes. This means that it requires a careful electrode layout and cannot be used to estimate the surface Laplacian at the edge of the EEG cap.

2.3 Data as a representation of the system

When examining data, it is important to remember that it is only *the data* that we are examining, and not *the system*. The data will reflect only certain aspects of the system (in a Platonic sense), and these may be combinations of several conceptually different components. Also, the signal that we measure is corrupted by noise – both external noise, and noise introduced by quantisation and sampling.

It is very important to be aware of the assumptions inherent in the various analyses that we perform – if these assumptions are incorrect, then incorrect interpretation will follow.

2.3.1 Stationarity

This idea is especially important when examining the validity of the conclusions drawn from data analysis. A stationary system is one whose statistical quantities (e.g. its probability density function) are constant over time [28]. Qualitatively, changes in stationarity can manifest as “changes in structure of the time series” or changes in the baseline of the time series [78]. For systems in which this is not the case, stationarity can often be approximated by con-

sidering a system as stationary only over a short period⁹, and then conducting analyses only over this shorter interval [34]. If we consider a physical system such as a pendulum, then for the system to remain stationary its parameters must remain constant: the length of the string, the mass of the ball, the air pressure, wind, etc – anything that would affect the behaviour of the system. This is not completely an accurate analogy to the brain because the behaviour of the brain, as a system, is emergent¹⁰. Also, the way in which neurons affect each other (excitatory, inhibitory, the strength of the response) will affect stationarity - although this is much slower to change than the neuron state.

The brain is stationary only locally and only for short periods [42, 39], and this is an important consideration when implementing signal analyses. To implement a short-term analysis, we can window the data. Note that generally, windowing the data will introduce artifacts into the data. The strength of these artifacts depends on the type of window used, and the analysis conducted.

Weak stationarity is another real-world compromise, and refers to a system that is stationary only to second order statistics. This is insufficient for a nonlinear approach – because in order for such analyses to be meaningful for nonlinear quantities, they must consider higher-order statistics [28].

2.3.1.1 Stationarity assessment

Using linear measures to assess stationarity can be problematic when one is going to examine the system using nonlinear tools. This is primarily because linear methods examine statistics up to and including second order, whereas nonlinear analyses also consider higher orders.

A good method to assess nonlinear stationarity is to use nonlinear predic-

⁹This is called short-term stationarity, and implies that the system is approximately stationary over short periods of time.

¹⁰This means that because of its complexity, the brain cannot be understood by understanding the rules and properties of the neurons – the system behaves in ways that are not obvious, and behaviour can change with time. In a sense, this means that although the physical properties of the neurons (conduction time, membrane permeability, extracellular fluid composition, etc) *are* relevant to stationarity there are other properties that are important. Emergence is discussed in section 5.2.2.

tive mechanisms. This involves a comparison of prediction errors for different segments of the time series. These errors should be consistent if the system is stationary [28].

For testing linear stationarity, it will make more sense to use linear analyses as part of our stationarity assessment. This forms part of a more general approach to testing stationarity, which is to use the same functions to test and analyse the data wherever possible.

2.3.2 How many data are “enough data”?

When performing signal processing, the number of data we have is limited by the available time we have for collection. Thus, a higher sampling rate will yield more data in a given time.

Because of this, we must tailor our sampling rate to our knowledge of the system’s stationarity, and also the number of data we require for our analyses. This varies depending on the type of analysis we wish to perform.

2.3.2.1 Linear Perspective

When examining a time series from an unknown source, the minimum duration of the data for analysis should be the inverse of the lowest constituent frequency when testing the degree of stationarity. This criterion simply means that we need enough data to be able to resolve the lowest frequency (at least one wavelength thereof) [4] without a short-term-stationary system changing state. As we approach the limit of our stationarity approximation (in a non-stationary system), lengthening this time series adds temporal smearing, because the system is evolving.

Increasing the sampling rate allows us to obtain more data for a given time, but not necessarily more information. This means that we are able to accommodate the recording duration limit that is required to approximate stationarity, but still obtain enough data to perform useful analyses. A higher

sampling rate also allows the resolution of higher frequencies. The Nyquist frequency is the highest frequency in a signal that we are able to resolve during sampling. In order to create an unambiguous sampled signal, free of aliasing, we need to sample at least double the Nyquist frequency.

Note that both these “rules” require a recording of a system that is free from noise. So while it is theoretically true that a frequency is resolvable when sampling at twice that frequency, this is not always true in practice. Thus, in a noiseless system, if we sample at double the highest frequency we wish to resolve, then we completely characterise that system and sampling at a higher rate, though yielding more data, will provide no additional information.

2.3.2.2 Nonlinear Perspective

Generally, nonlinear analysis does not resolve frequency components in the same manner as does Fourier analysis and other linear analyses. When analysing a system where we have no cause to expect a specific upper frequency then the higher the sampling rate the better. In fact, even if we are only desirous of resolving frequencies to 300 Hz, then it is still reasonable to sample to a much greater rate than that suggested by the Nyquist frequency. This is because it assists us in the removal of noise by processes such as nonlinear filtering (see 3.3.6). This is for several reasons.

The first is that nonlinear systems typically exhibit sensitivity to initial conditions, so being able to precisely measure the system’s output is important. In the real world, where we have a limited amount of data, corrupted by noise, then sampled and quantised, we are necessarily limiting our ability to precisely measure the desired information. Increasing the sampling rate can allow nonlinear modelling to try to extract the behaviour of the system from the acquired data.

Some studies have shown that the sampling rate can have an effect on the results of nonlinear analyses. Jing and Takigawa [27] showed that, by varying

the sampling rate of EEG, different results were obtained when analysing the correlation dimension of the data. Similarly, Martinerie et al [48] said that “to improve the time window resolution, higher sampling rates would be needed,” meaning that by sampling at a higher rate, they could perform their analysis on data of shorter duration, while maintaining the same *number* of samples.

Chapter 3

Nonlinear Systems

In chapter 2, I introduced linear systems, and described what they were and their relative scarcity in the real world. A point I made, was that many nonlinear systems approximate linearity and that for this reason the application of linear analysis to such systems is a useful enterprise. However, it has limitations. Recall my examples of approximately-linear systems – while it is possible to make predictions of outcome (eg. a struck billiard-ball) using linear analyses, this method of analysis would not work as well if we needed to extrapolate the calculations over a longer duration, or more object interactions. Consider a break in a game of billiards – the white ball hits a cluster of billiard balls which scatter: rebounding from the sides of the table and each other. Because of the large number of interactions, nonlinear effects will accumulate. Using Newton’s (linear) laws to attempt to forecast the outcome of such a scenario is unlikely to yield accurate results.

This chapter introduces nonlinear theory, beginning with a description of what the words *nonlinear* and *chaotic* mean, and contrasting them with the linear systems previously discussed. It then describes some nonlinear analyses, and finally discusses some of the requirements and caveats of nonlinear analysis.

3.1 What is a nonlinear system?

As mentioned in section 1.3, a nonlinear system is one which does not exhibit the property of *superposition* – the combination of two properties *additivity* and *homogeneity* (equations 3.1 and 3.2).

$$f(x + y) = f(x) + f(y) \tag{3.1}$$

$$f(\alpha x) = \alpha f(x) \tag{3.2}$$

In actuality, *most* physical systems are nonlinear, leading Stanislaw Ulam to say that talking about nonlinear science is “like talking about non-elephant zoology”. The inapplicability of superposition means that common scientific methods, such as reductionism, cannot be applied to nonlinear systems – or can at best be used to approximate the behaviour of the system.

Nonlinear systems are typically described by differential equations which are often difficult or impossible to solve. Historically, this has made the analysis of nonlinear systems problematic and is only now changing due to the increasing sophistication of computer-based modelling.

3.2 What is a chaotic system?

For a system to be chaotic, it must be nonlinear. This is a necessary, but insufficient condition. The defining characteristic of a chaotic system is an extreme sensitivity to initial conditions. This means that the future behaviour of such a system ranges widely, depending on the system’s state at present or, to put it another way, that small perturbations in the system can cause large changes in the system’s future behaviour. This concept is also known as the butterfly effect (“a butterfly flaps its wings in Peking, and the weather in Tokyo is different”).

The aim of this section is to illustrate the behaviour of a chaotic system.

For simplicity we will look (as does Kelso [33]) at a simple system of a body of oil being heated from underneath.

While there is a small heat differential between the cooler top and the warmer bottom of the body of oil, the heat moves by conduction – the quasi-random movement of individual atoms. For a level of heating below a certain threshold, this will produce a stable system.

If the heat differential is made a little larger, then a convection system is established, whereby hotter, less dense oil from the bottom is displaced by cooler, denser oil from the top. This forms a steady current that flows in a closed loop (the *direction* of flow is random, and is impossible to predict). This occurrence is called a bifurcation, because a previous steady state of the system (that of heat conduction at lower temperatures) is (at higher temperatures) no longer stable, and has been replaced by one of two new stable states (by a quasi-random decision) [20].

When the temperature differential is increased still further, there is another set of bifurcations, then another, and another, until a chaotic state is reached. This is called an emergent property of the system, and it is increasingly thought that this is relevant to brain operation.

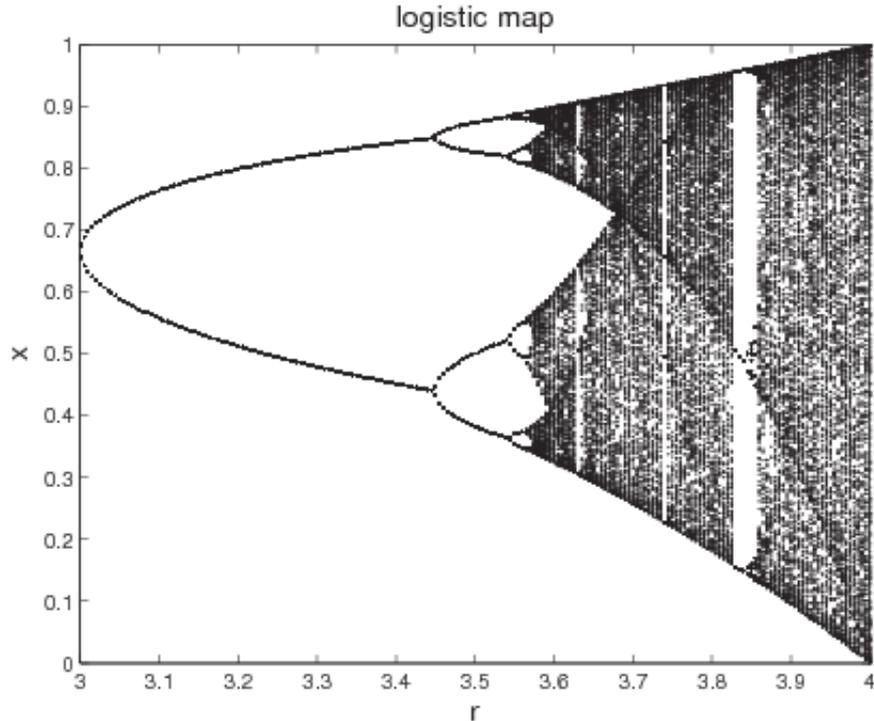


Figure 3.1: An example of a chaotic system exhibiting bifurcations

This figure is a bifurcation diagram of the logistic map. It was created using the equation $x_{n+1} = rx_n(1 - x_n)$, and shows the likelihood of various final values for x depending on values of r . As can be seen, while $r < 3$, there is one outcome for x . However, when $3 < r < 3.45$ there are two potential values. The transitions at $r = 3$ and $r = 3.45$ are called bifurcations, and represent a “threshold of change” in the system. Note that the bifurcations come more and more often as r increases until, when we reach the following Feigenbaum constant (the apparently transcendental ratios between successive bifurcation intervals), the system becomes chaotic (there are many probable values for x). If we relate this figure to the heated oil experiment discussed in section 3.2, then the values of $r < 3$ correspond to the oil being gently heated so that heat is distributed throughout the oil by conduction. Values of $3 < r < 3.45$ correspond to the temperature being increased slightly, so that a convection system is established. The two states refer to the direction of oil flow. Above $r > 3.45$ (as the rate of application of heat increases) progressively more states become possible for the convection pattern (resulting in more complicated flows), until chaotic behaviour emerges at the first Feigenbaum constant.

3.3 Nonlinear analyses

Generally, successful approaches to signal analysis assume that a system is either linear stochastic¹ with an ignorable nonlinear component, or that the stochastic component is small and the system is nonlinear deterministic. This is for several main reasons [19]:

1. The future state of a chaotic system is sensitively dependent on the current state, so that any noise leads to an inability to predict the future of the system [28, 20]. This is an important aspect of many signal analyses, such as autocorrelation and nonlinear phase-space analysis.
2. Because of the sensitivity of nonlinear systems, they often appear to behave in quasi-random ways, making prediction very difficult, so it can be unclear which component of a recorded signal is noise, and which is true data reflecting the behaviour of the system.
3. Sampling (temporal discretisation) and quantisation (signal level discretisation) *always* introduce noise [59], so it can become difficult to distinguish between noise introduced during measurement, noise as a result of sampling and quantisation, and the actual evolution of the system.

3.3.1 Reconstruction of the Phase Space

A representation of a signal in phase space means forming a "complete" description of the signal at every point in time. The phase space variables capture a full and clear view of the dynamics of the signal, and can reveal information that is invisible in the time series [74]. There are two main ways in which phase space variables can be obtained from time series data. The first is by combining information from multichannel, simultaneous recordings, and the second

¹A stochastic system is one that is non-deterministic, and whose behaviour is described by probability density functions (PDFs).

is to embed a single channel of time series data into multiple phase space dimensions. A combination of the two can also be used (ie. forming a phase space vector from the embedding of multiple, spatially separated, simultaneous recordings).

Let us briefly look at an example system to illustrate this. The Lorenz equation was defined while trying to model weather systems [20], and contains three system parameters in three differential equations. The three equations are shown in equation 3.3.

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x) \\ \frac{dy}{dt} &= x(\rho - z) - y \\ \frac{dz}{dt} &= xy - \beta z\end{aligned}\tag{3.3}$$

We model the system by choosing values for σ , ρ and β , as well as setting initial values such that $x = y = z = 0.1$, and iterating the equations 9000 times with a time increment of 0.001 units per iteration. This process yields time series data which are displayed in figure 3.2.

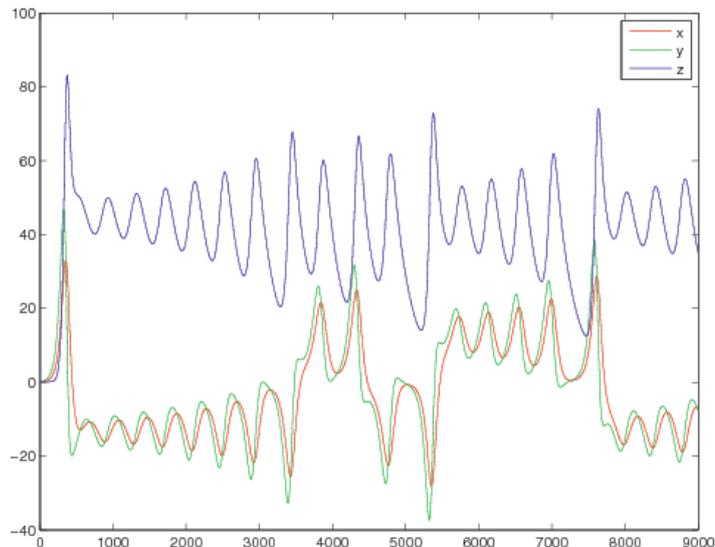


Figure 3.2: Time series data of Lorenz model

The Lorenz equations are a simple set of three nonlinear differential equations that were designed to model a simple weather system. The behaviour that emerged from the system exhibited extreme sensitivity to initial conditions. This graph illustrates the behaviour for starting conditions of $x = y = z = 0.1$.

Although this is useful in examining the system, it is perhaps not quite as useful as figure 3.3, which contains the same data, embedded into three dimensions:

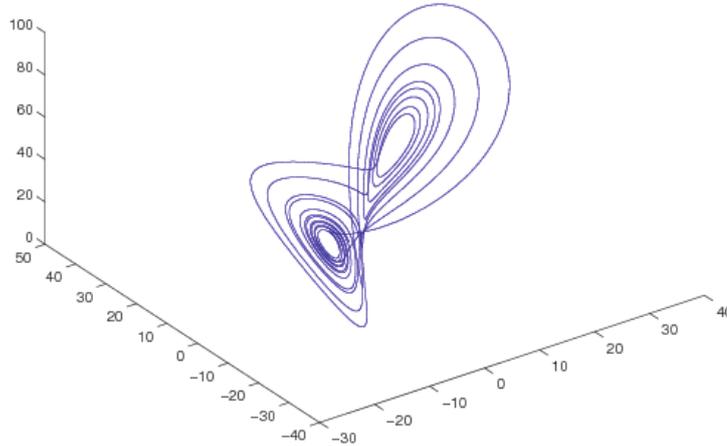


Figure 3.3: Phase-space representation of Lorenz model

By embedding the data (shown in figure 3.2) into a three dimensional phase space, we are able to more clearly see the behaviour of the system. The process of embedding is described in section 3.3.1.1.

3.3.1.1 Embedding

Embedding is the process by which the phase space evolution of a system is derived from its time series. The embedding theorem was developed in the early 1980s [55].

Embedding involves taking samples from the time-series data, vectorising them into n dimensions, waiting time T_s (referred to as the lag) and then repeating the process [30, 19]. Thus if we had a segment of data $x(t) = [1, 2, 3, 4, 5, 6, 7]$, and we decided to embed it in three dimensions with a lag of 1, then our first three embedded points would be

$$X1 = [1, 2, 3], X2 = [2, 3, 4], \text{ and } X3 = [3, 4, 5]$$

Here is a dilemma, especially for an unknown system, because a choice must be made as to what dimensionality the embedding process should impose on the data. If our only knowledge of the system we had was time-series data, how

could we know that the embedding dimension should be three dimensional (as shown), and not two- or four-dimensional? For instance, if two dimensional embedding was used, with a lag of one we'd have:

$$X1 = [1, 2], X2 = [2, 3], \text{ and } X3 = [3, 4]$$

which is quite different. This is the first problem.

Even if we retain three dimensional embedding, but use a lag of 2, we obtain

$$X1 = [1, 3, 5], X2 = [2, 4, 6], \text{ and } X3 = [3, 5, 7]$$

which is different again. How is the lag chosen? This is the second problem.

The two problems are solved in quite distinct ways.

3.3.1.2 Choosing an embedding dimension

If some time-series data are embedded at a lower dimension than they ought, a form of aliasing occurs. This results in the mapping (or folding) of points from a higher dimension to an incorrect place in a lower dimension – this is called “projection”. Imagine a room full of objects with a distant light shining from one side onto a wall on the other. The shadows cast by the objects are essentially a projection – they lack three-dimensional information.

Projection produces a situation where false nearest neighbours occur [2]. These are points that are erroneously placed nearby in the phase space, and this is an artifact of the projection process. Clearly, this is an undesirable situation. If, on the other hand, time series data are embedded in a dimension that is too high, then there is a subspace that contains no useful information [19]. This isn't really a problem, however it makes any calculations performed within this space more computationally intensive. The embedding dimension is generally found by computing a quantity like the fractional dimension, and then ensuring that the embedding dimension exceeds this. In practice it can involve a computation of false-nearest-neighbours (figure 3.4) for a range of dimensions, and the embedding dimension is chosen to set the number of false

nearest neighbours to be sufficiently small.

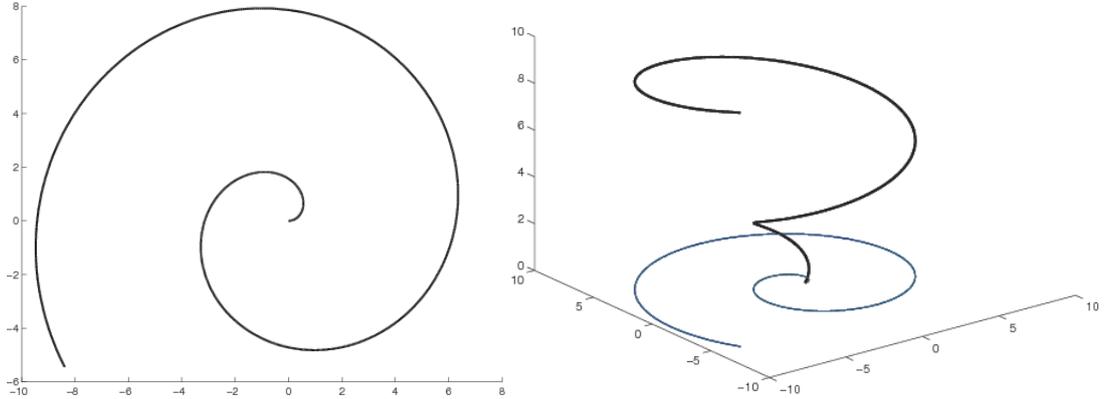


Figure 3.4: Embedding dimension illustration

The same data, represented in a two-dimensional phase space (left) and a three-dimensional phase space (right). The 3D trajectory casts a blue shadow, on the x - y plane, from a distant light source directly above the trajectory. Notice that the shadow corresponds exactly to the trajectory in the two-dimensional phase space. This is a “projection”.

If we examine these phase-spaces for nearest neighbours, it is clear that we will obtain different results, because the two-dimensional space results in different spatial relationships between trajectories, causing points be false-neighbours because of incorrect mapping. Clearly, information is lost when representing these data in two dimensions.

3.3.1.3 Choosing a Time Lag

The embedding theorem is not useful in determining an appropriate time lag, and this is a quantity that must be established using other means. Abarbanel [2] considers that there are three things to consider when choosing an appropriate time delay:

- Because we are using sampled data, T_s should be an integer multiple of the sampling period.
- If the time delay is too short, the system will not have changed. This amounts to the two phase-space points lacking sufficient independence [19], or that the new points contain little new information. This is akin to over-sampling, and merely results in huge amounts of data, containing

little information (however, this lack of new information can be used to assist in noise reduction).

- As mentioned, chaotic systems are unpredictable past certain temporal horizons. If T_s is too long, then the two points are completely unrelated and we will obtain a string of random points – which will be useless.

Thus, in calculating an appropriate time delay, we examine the mutual information (see 3.3.4) of the signal at different time lags. Clearly at a time lag of zero, the mutual information is very high and as the lag increases the mutual information drops. Generally speaking, there are local minima and maxima in the mutual information as T_s changes. If we chose the first local minimum as the lag, then we have points which have a reasonable level of interdependence, but are not so close in time as to be redundant. The time lag will affect what is visible in the phase space, i.e. it controls the obviousness of particular frequency components. In their paper on nonlinear noise reduction, Kantz, et al. [29] discuss the effect of embedding at different time lags. This controls the visibility of different frequency components of a signal. For instance, embedding with a small lag means that high frequency signal components will be prominent in the phase space, whereas a large lag will result in low frequency components being more obvious.

There are papers that deal specifically with using multichannel simultaneous recordings of the system [12] which can help in the analysis by using the mutual information between channels.

3.3.2 Correlation dimension

Correlation dimension is an estimate of the number of degrees of freedom of a signal [34]. It is seriously affected by non-stationarity, which tends to reduce the estimated dimension [29], while drift tends to increase the estimated dimension. Thus, the correlation dimension can also be used to assess stationarity. There are other measures of dimension (eg. box-counting dimension,

Hausdorff dimension) but the correlation dimension has the virtue of being relatively straightforward to calculate.

Correlation dimension is estimated from the correlation integral, which measures the likelihood that randomly chosen pairs of points are separated by a distance less than ϵ ,

$$C_x(\epsilon) = \Pr(D_x < \epsilon),$$

where D_x is the distance between randomly chosen points of the embedded signal. The correlation integral is approximately proportional to ϵ raised to the power of the correlation dimension, and so correlation dimension can be estimated from the slope of the plot of correlation integral against ϵ .

$$CD = \lim_{\epsilon \rightarrow 0} \frac{\log(C(\epsilon))}{\log(\epsilon)}.$$

In the application to real data that are short, noisy and non-stationary, one might expect that analyses results would be unreliable. However there is evidence that, if carefully applied and interpreted, nonlinear analyses such as correlation dimension can yield information superior to linear time- and frequency-based analyses [41]. A good example is [78], where EEG in patients with temporal lobe epilepsy were analysed using a correlation dimension and a correlation entropy. They used implanted electrodes to record EEG from 5 patients and compared 20 s stationary pre- and post-ictal EEG samples, although they assessed EEG stationarity by eye only. Using these two groups, they compared their analyses with a visual inspection of the same epochs of EEG, and found that the nonlinear quantifiers could distinguish between ictal and non-ictal EEG.

3.3.3 Entropy

The concept of entropy in signal processing was developed by Shannon [70] who defined it as shown in equation 3.4.

$$H = -K \sum_{i=1}^n p(x_i) \log p(x_i), \quad (3.4)$$

where $p(x_i) = Pr(X = x_i)$ is the probability of an event x_i in the data set X , and K is a constant that determines the unit of the output. Shannon's definition of entropy relates to the Maxwell-Boltzmann-Gibbs entropy that is relevant to thermodynamics. This provides a measure of the number of ways in which the micro-states of a system can be arranged to produce the given, observed macro-state – which gives an indication of the complexity of the system.

The preceding equation, when applied to a binary system, states that maximum entropy occurs when we have the greatest uncertainty about the outcome of our measurement, as shown in figure 3.5. The figure illustrates the entropy of new data from a binary system and shows that entropy describes the amount of information in the data, which is related to our uncertainty regarding new data. If we consider $P(X)$ as the probability of a new sample being 1, then the entropy can be viewed as how much information we learn as $Pr(X)$ varies. If $Pr(X) = 1$, then we know what to expect (the next sample will be 1), and we don't learn anything new – hence the value of the entropy is 0. The same is true if $Pr(X) = 0$, since we then are sure that the next sample will be 0. As we become less certain of the value of the next sample, the entropy increases until, at the point of maximum uncertainty (where we are equally likely to receive a 1 or a 0), the entropy attains a maximal value.

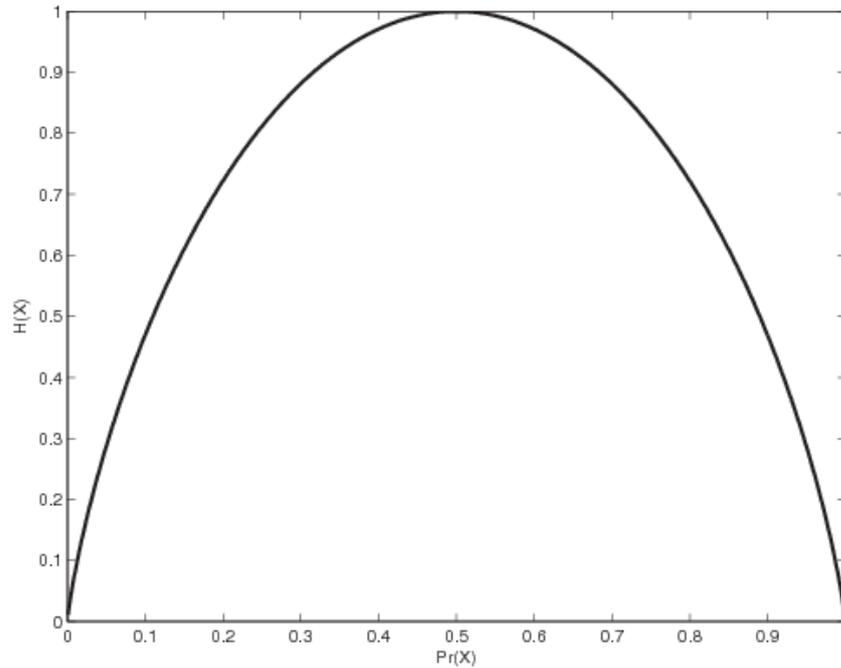


Figure 3.5: Entropy as a function of probability of outcome

This figure was produced using the function $H = -\bar{x} \cdot \log(P_i(\bar{x})) - x \cdot \log(P_i(x))$, where \bar{x} denotes $\text{NOT}(x)$. It illustrates the way in which the entropy (or the information content of a signal) is determined by the likelihood of the learned information. If we are told something unlikely, (such as “the sun will **not** rise tomorrow”) then we learn more than if we’re told something we know already (such as “the sun **will** rise tomorrow”). In this figure, incoming data is a bit, so maximum information is learned when there is equal chance of a 1 or 0, and there is less information when we’re more likely to receive a 1 than a 0, or visa versa.

Shannon was the first to define the information content of data, and his work has come to be known as *information theory*². If we are receiving serial binary data, and the calculated entropy is less than 1 bit per bit of received

²Shannon’s seminal paper [70] discusses many interesting ideas, such as a “series approximation to English”, where a stochastic process produces random letters. It has knowledge of English words and how letters and words relate to each other (i.e letter and word patterns). It is able to produce very English-like statements, although they ultimately have no meaningful content. Shannon showed that typical written English text has an entropy of between 0.6 and 1.3 bits per character. In English, the letters *e*, *a*, *t* and *o* are common, whereas the letters *j*, *q*, *x* and *z* are rare. If English had a more uniform frequency distribution of letters (for example, if the letters *e*, *a*, *t* and *o* were used less often, while *j*, *q*, *x* and *z* were used more often) then it would be more difficult to predict the next letter, and the average entropy of English would increase. In fact, this is precisely what was done in early mono-alphabetic substitution cryptography to help defeat cryptanalysis based on the frequency distribution of letters. It was common to use both the letters *y* or *z* to represent *e*.

data, then there is redundancy in the data. Redundancy means that the same information could be conveyed to the receiver using fewer transmitted bits, if a more optimised encoding scheme was used. It is this principle that is applied in data compression software.

3.3.4 Mutual information

Mutual information (MI) is a statistical analysis of the common signal components between channels. If one visualises the sets of possible outputs for two signals (see figure 3.6), then the mutual information is the amount of overlap between these two sets.

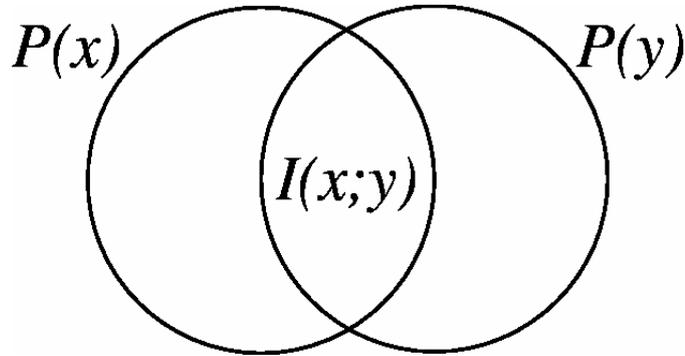


Figure 3.6: A Venn diagram, illustrating Mutual Information

This Venn diagram shows the information contained in two sets of data, with the overlapping component being the information that is mutual to both sets.

Mutual information between two random variables is defined as shown in equation 3.5.

$$I(X; Y) = \sum_{y \in Y} \sum_{x \in X} p(x, y) \log_2 \frac{p(x, y)}{p(x)p(y)}, \quad (3.5)$$

where $p(x, y)$ is the joint probability of X and Y , and $p(x)$ and $p(y)$ are their marginal probabilities. Because this equation takes the log in base-2, our MI calculation yields an estimation in bits³.

Mutual information is calculated by ranking the data. These rankings are compared between signals and, broadly speaking, MI is the amount that can be deduced about one signal if the other is known. This is a way of measuring

³Alternatively, were we to use a log in base-10, our result would be measured in dits.

dependence between two data sets: A high level of mutual information implies the data are dependent, and there is little information gained by learning the second data set, whereas a lower level implies independence, and that additional information can be gained from the second data set. This tool can be used in several ways.

Firstly, a signal can be compared with itself (by using a delay), or with a portion of itself (ie. for establishing an appropriate embedding lag)⁴. This is similar to an autocorrelation and can give an indication of long term changes in the system, by highlighting temporal dependencies.

Also, the mutual information can be found between channels. This is similar to a cross-correlation. This will give an indication of how the different areas within the brain relate to each other. If there is a high level of mutual information, one might conclude that the two areas are involved in the same or similar activity, or that there is some other phenomenon (such as seizure) occurring. Less MI might imply that the two areas are not related.

3.3.5 Nonlinear prediction

When deciding on parameters to embed time series data, we are ascertaining the state of the system by measuring variables that are controlled by the state of the system – not measuring its state directly. Thus there is an error component in our evaluation of exactly what the state of the system is (both now and then). There are several methods of reducing this error⁵ but an interesting method uses a basic phase space prediction algorithm.

Examine a small population of states that are close to the current state, see how they evolved (to where in the phase space does the system progress after a certain time lag?), and then perform some averaging (eg. mean) as to their next movements – this yields an average behaviour of the system for states

⁴A single channel could be analysed by sliding it along itself (wrapping around the ends) and calculating the mutual information.

⁵as mentioned, it can be difficult to ascertain what is error due to the embedding, what is measurement noise and what is a signal component

close to the current state. This represents our prediction of the evolution of the system from the current state [28].

When quantifying the effectiveness of a predictor, one must ensure that the analysis is actually making a prediction that has not yet happened (out-of-sample) rather than trying to account for past events (in-sample) [29]. Thus, it is essential that training data and testing data are separated.

Prediction can be used to test for the presence of, or a change in, nonlinear quantifiers. For instance, one region of data is used as a "training ground" for a nonlinear predictive algorithm (such as described above) and its effectiveness is tested in another region of data. This will tend to highlight changes in the data (if the accuracy of the prediction decreases then the training data is not a good model of the test data). Such changes could suggest that the system is not stationary across the prediction time and we need to use a temporally smaller window in our analyses.

The application of the nonlinear predictor to unfamiliar data will always result in worse performance, so it is more meaningful to compare the relative prediction accuracy between two sets of testing data, than to just compare testing and training data sets.

3.3.6 Nonlinear filtering

The concept of prediction can be used in the construction of a nonlinear filter. By looking in the phase space, and examining how a particular sample differs from the expected (based upon the evolution of n nearest neighbours), a correction can be made. Look at a point, and then look at its future position in the phase space. If this differs from the average of the futures of the n closest points, then consider it an error and replace it with the average (weighted mean). Kantz talks about this in his paper nonlinear noise reduction [29], as well as his book with Schreiber [28]. Alternatively, we could examine just

the component of the signal that is not predicted (perhaps calling this the "chaotic" component), in a manner similar to AR analysis.

3.3.7 Lyapunov exponents

It is worth emphasising that if a *deterministic* system is in *exactly* the same state that it was previously, then its future behaviour will mirror *exactly* what it did then. Based on this we might expect that if the system is in *nearly* the same state, then it will behave in a similar way. However, an important attribute of chaotic systems is their sensitivity to initial conditions⁶. This means that two points in phase space can be quite close and yet the evolution of the system will proceed in a completely different manner [20, 19]. In fact, the rate of divergence is typically exponential, and this has been called a *necessary* condition for a system to be chaotic⁷ [28]. It is possible to calculate a "prediction horizon", by quantifying the rate of divergence of nearby states, such as those just mentioned. This is called a Lyapunov exponent, and is denoted by λ . Note that an n-dimensional system has n values of λ that describe the rate of divergence.

The largest Lyapunov exponent, λ_{max} , is often taken to represent the "amount of chaos" in the system. A positive value of λ_{max} means exponential divergence of nearby trajectories and hence the system is chaotic – that is why it is common to talk only about λ_{max} . It is worth emphasising that the Lyapunov exponent is so-named because the rate of divergence of nearby trajectories in a chaotic system is exponential. Lyapunov exponents are rarely used for real-world EEG analysis, because of their instability and unreliability when there are insufficient data [7] which is often the case when utilising a moving-window analysis.

⁶As stated by Kantz [28], "...our everyday experience, <<similar causes have similar effects>>, is invalid for chaotic systems."

⁷Exponential divergence is a result of sensitivity to initial conditions.

3.3.8 Data Compression

A problem commonly encountered with EEG analysis is that of handling the large number of data. One way to circumvent this is to perform analysis of the zero-crossings of EEG data. The actual analysis in such a study (eg. [79]) is concerned with the times between zero-crossings, rather than the actual data recorded – effectively an implementation of data compression. This means, however, that much useful data is lost. For example, if there is a low amplitude high frequency that is added to a large amplitude low frequency, then the high frequency component may not produce zero crossings causing information loss.

Van Puttan et al's [79] experiment records multiple channels and attempts to examine their temporal relations. This is similar to other studies, eg. Mackenzie [45]. But since Mackenzie et al used raw EEG data, a finer temporal resolution is afforded. The trade-off is that there are many more data, and therefore calculations take longer and are more involved. Van Puttan's research details phase-locking and link rates in the brain. They speculate that, together, these form transient oscillations between neuronal networks, which allow coding for individual percepts to represent or operate sensory and cognitive functions [79]. New theories of the brain indicate that thoughts originate in short-lived patterns of time-dependent synchrony, which links separate areas of the brain. The researchers believe that synchrony binding between the neuronal subsystems is reflected in transient phase-locking events.

3.4 Requirements of nonlinear analysis

This heading might better be phrased “Is it reasonable to analyse EEG using nonlinear measures?” and the answer is not trivial. It has been implicitly assumed by many researchers that the operation of the brain must be nonlinear. Even if this is the case, it does not necessarily follow that the system perceived in the EEG is also nonlinear. Section 5.2 introduces the idea that EEG is

produced from the summative behaviour of billions of neurons and that, in the case of human subjects, it has been filtered by its passage through the dura and the skull. Even if the behaviour of the neurons, or single cell recordings, yield data containing nonlinearities, it does not necessarily mean that EEG, recorded from scalp electrodes, will also contain nonlinearities.

Despite the maturation of nonlinear theory, there are many critics. For example, Lehnertz et al [42] state that "it is now commonly accepted that the existence of a deterministic and even chaotic structure underlying neuronal dynamics is difficult or even impossible to prove." However, they proceed to say that "there is converging evidence that nonlinear approaches to the analysis of brain systems are able to generate new clinical measures as well as new ways of interpreting brain function, particularly with regard to epileptic brain states."

3.4.1 Bootstrapping

This technique can be used when there are insufficient data to analyse, and to reduce the variance in analysis outcomes.

If we have a data set as shown in 3.6

$$X = [x_1 x_2 \dots x_m] \tag{3.6}$$

and a statistic of that set, as in 3.7.

$$\theta = E[x] \tag{3.7}$$

We can take $X_i \in X$, a subset of X at random (and with replacement), and from this, calculate $S_i = E[X_i]$, an estimate of θ , based on X_i . This procedure can be repeated with i being arbitrarily large. We can make a distribution of S in order to find $\hat{\theta}$, an approximation for θ based upon the probability density function $[S_1, S_2, S_3 \dots S_i]$. Because we are taking an average of many calculations, bootstrap theory tells us that $\hat{\theta}$ will have a smaller variance [13]

than θ , although this will depend on the specifics of the data.

When bootstrapping, we examine the data and perform statistical analyses on it. We then ask, how sensitive are these analyses to changes in our data set. If we hadn't recorded a particular subset of the samples, how would our statistical conclusions differ? This is based on the premise that we have imperfect data, and that there are *samples we could have recorded, but didn't*. How could these missing samples affect our conclusions? That is the question that bootstrapping attempts to answer.

3.4.2 Surrogate Data

One method of validating the choice to use nonlinear analyses is by the use of surrogate data [30, 67]. This is a well established method of verifying that observed phenomena are due to variations in nonlinear quantities, and not those of a linear system with an added stochastic component.

In order to implement this, a null hypothesis is made. This generally reads in a manner similar to: "the data are generated by a Gaussian (linear) process undergoing a possibly nonlinear static transform [37]." If nonlinear dynamics are able to be detected from the time series, then the null hypothesis can be rejected. However, failure to reject the null hypothesis does not imply the absence of nonlinear dynamics, as there are other causes for this observation.

Kugiumtzis [37] discusses this in detail and suggests that one possibility is that noise masks the nonlinearity, or that there are insufficient data, or even that the data does not accurately reflect the behaviour of the system (see section 2.3). He also suggests that because different nonlinear quantifiers examine different aspects of the system, it is naive to examine a single analysis of one data set and form a conclusion. What this amounts to is that failure to reject the null-hypothesis implies either that it is true, or that there is a lack of evidence.

Essentially, the aim of the surrogate test is to reduce the possibility of

false-positive results – we aim to prevent the incorrect identification of chaos. This is somewhat problematic, because there are several legitimate nonlinear analyses that are also sensitive to linear dynamics. One such analysis is that of entropy, which is affected by the amplitude of the analysed signal. This can be corrected by scaling the output by the inverse of the amplitude of the input (section 7.1).

As mentioned, surrogate data provide a means of testing the validity of analysis of data using a particular methodology. In this case, we test whether the changes we see in our nonlinear quantifiers can be accounted for without them reflecting changes in brain states. For example, let us suppose that EEG is a linear stochastic process, and that our act of measuring this introduces nonlinearities – the nonlinearity that is apparent in our analysis. This is equivalent to equation 3.8

$$D_{meas} = f(s_n), \tag{3.8}$$

where s_n is a stochastic process, and f is a nonlinear function. It is possible to test for (and hopefully reject) this possibility using surrogate data tests, thus helping to validate our analysis techniques.

It is impossible to completely reject the null hypothesis, we can only do so to a particular level of significance. However, we can show that, despite trying, we were unable to account for the observed changes using simpler explanations, and therefore consider it reasonable to attribute these changes to nonlinearities in EEG. This is really a manifestation of Occam's Razor: we only accept the more complicated explanation for observed data (that the appearance of nonlinear variation in the EEG is a reflection of changes in brain state) if we are able to reject the simpler explanation (that the appearance is due to our measurement or recording method, or some other artifact). Practically, this amounts to modelling multiple data surrogates, using different algorithms, and testing these with our nonlinear quantifiers.

3.4.2.1 Non-stationarity

Non-stationarity must be considered in the use of surrogate data, since the statistical properties of the system change with time. If this is not considered, and the possibility of non-stationarity is not part of the null hypothesis, then the null hypothesis could be (incorrectly) rejected because of the non-stationarity, rather than nonlinearity. Hence, to allow our data to be non-stationary while still testing for nonlinearity, we must consciously include this in our surrogate [67]. One method of doing this is by designing our surrogate time series to mimic the “instantaneous” mean and variance of the original time series.

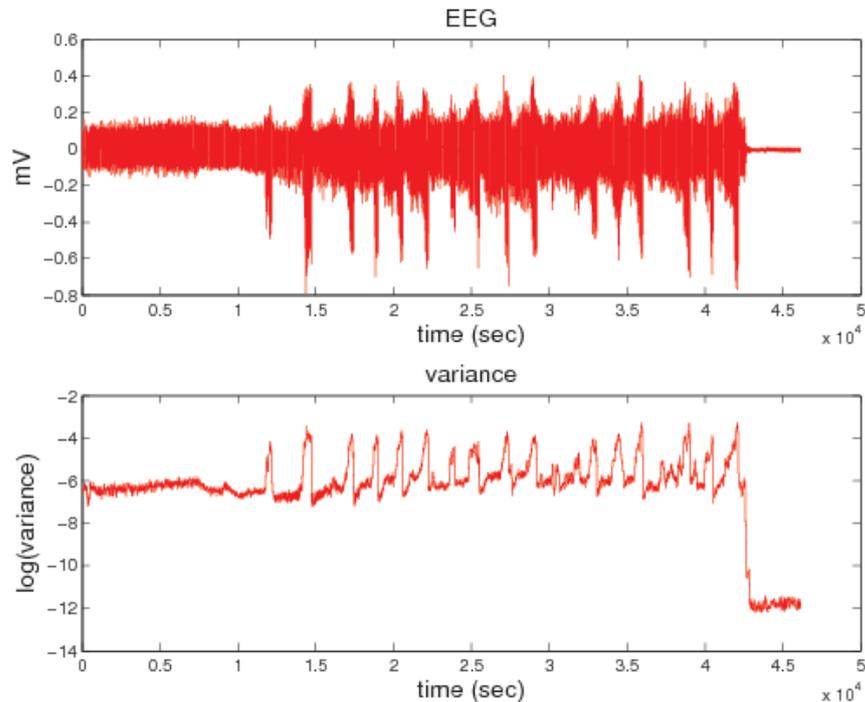


Figure 3.7: Rat EEG showing seizure

This figure illustrates that EEG is non-stationary. During the seizure there is an increase in variance in the baseline EEG, due to increases in amplitude, which represents a change in the statistics of the system.

3.4.2.2 Application

Thus, a precursor to EEG analysis using nonlinear tools is testing of those tools using surrogate data. There were several aspects to this, as there are

many different surrogates (and null-hypotheses) to test.

1. The first surrogate tests the null hypothesis that the data were produced by a linear stochastic process. It was constructed by taking the Fourier Transform of the data, multiplying it by random phases and then applying the inverse Fourier Transform to yield the surrogate time series. This can be viewed as an attempt to make an AR model of the data, but incorporating some 'flexibility' into the AR parameters – essentially attempting to match against multiple AR models. So, if

$$|S_k|^2 = \left| \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} s_n e^{i2\Pi kn/N} \right|^2 \quad (3.9)$$

is the discrete Fourier Transform of the original data, then the surrogate data are found as shown in equation 3.10.

$$S_{surrogate} = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{i\alpha k} |S_k| e^{-i2\Pi kn/N} \quad (3.10)$$

where $0 \leq \alpha_k \leq 2\Pi$ are random numbers [67].

2. Let us imagine that inside the brain is an ideal EEG source, which is a Gaussian linear process. However, attenuation by the dura, skull and scalp, as well as the transfer function of the measurement apparatus is a nonlinear function. This is the null hypothesis that the second surrogate is designed to test. The algorithm operates as follows. First, make Gaussian-distributed data of the same length as the EEG data. The Gaussian data and the EEG data are both sorted by magnitude. The sorted EEG data are now replaced with the sorted Gaussian data, and multiplied by a random phase. Then the data are un-sorted, to “reconstruct” the EEG data. What we have effectively done here is approximately preserved the shape, and sample-to-sample relationships within the EEG, but have rescaled it so that it now has a Gaussian distribution

function. This method of surrogate generation is called the amplitude adjusted Fourier transform (AAFT) [77].

A problem facing this method of surrogate generation is as follows: Because we are rescaling a finite number of points, we have only a limited subset of a Gaussian distribution function to scale to – so the process isn't exact. The difference between the ideal scaling and that which is realised is independent from sample to sample, and hence is white noise. This additive white noise process causes a flattening of the spectra of the surrogate relative to the data. Clearly then, this surrogate will have differing low-order statistics to the original data, which must be repaired before analysis.

Because of this limitation, the process was refined by Schreiber and Schmitz [68] to improve its performance. Their method is called iterative AAFT (iAAFT), and makes successive improvements to the spectrum of the surrogate so that it matches that of the data as closely as possible. iAAFT works by shuffling the real data $\{x\}$ to produce $\{z^{(i)}\}$ (or just using white-noise), and then iteratively performing the following two steps, where i refers to the i^{th} iteration. [36, 67]:

- (a) make a power spectrum (but not the phase) of $\{z^{(i)}\}$ equal to that of $\{x\}$. Define $\{y^{(i)}\}$ as the time series of $\{z^{(i)}\}$.
- (b) reorder $\{x\}$ to the same rank structure as $\{y^{(i)}\}$, making their autocorrelation identical.

which, after a finite number of iterations, results in a surrogate with an abs(FFT) and an autocorrelation close to that of the original data, but lacking higher-order statistical similarities with the EEG.

This method, due to the iterations, is substantially more computationally intensive than the AAFT surrogate.

Chapter 4

Classification

This chapter introduces some of the concepts involved in the classification (*automatic identification*) of an unknown sample. This is an intersection of several fields including information theory, statistics and machine learning. The chapter then discusses several methods of classification, including their method of implementation and some of their limitations.

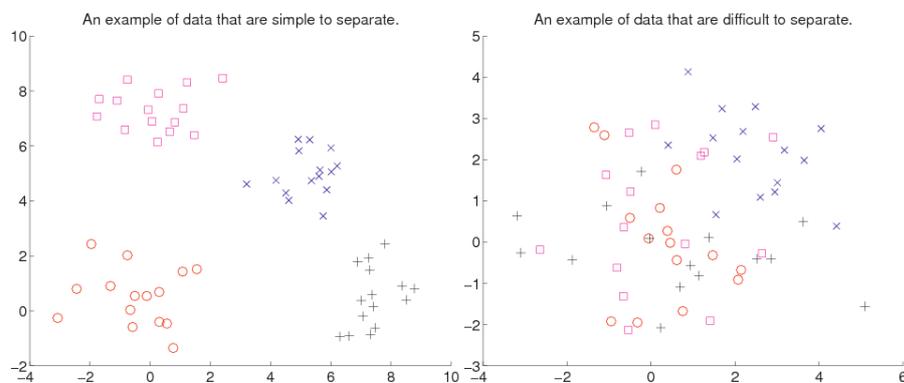


Figure 4.1: Linear discriminant analysis – example data

These figures show two fictional experiments, each with 4 classes of data, where each data point has two measured variables (represented on the axes). The left example shows data that have much more clearly defined groups than the data in the right figure. Let us imagine that we obtain data from an object from an unknown group. We wish to identify the group to which it belongs, so we compare the variables from our unknown object to the graphs showing the distribution of the groups. Clearly, the task of classifying the object will be easier with the left data than the right data (which have a much greater degree of overlap).

In general, this problem is extended to as many dimensions as there are variables describing the data.

Figure 4.1 shows two fictional experiments. Each experiment consists of two-dimensional data recorded from objects, of which there are four distinct types. One experiment shows objects that are simple to distinguish by these measurements, and the other shows objects that are difficult to distinguish. It may be that we record a new sample, and wish to decide to which group it belongs. This will be a much simpler process in the case of easily-distinguishable objects.

Classification is the process of using a statistical analysis to group items based on quantitative data reflecting the state of the items. These measured data are often called a *variable*, *trait* or *feature* of the item (and correspond to the axes in 4.1). Usually, the classifier is taught to differentiate between groups (the *types* or *classes* of objects) using a *training set* of data, that has already been classified. This allows the classifier to learn the relevant component of the data, and its application in discriminating between groups.

Another way of viewing this is to consider data as being placed into an n -dimensional space, where n is the number of variables that are used to describe each item. The goal of the classifier is to find a hyper-surface¹ that bisects (in the case where there are two groups) the training data into groups. This same surface is then applied to the testing data, to classify the data.

Another method is to approach classification as a probabilistic *estimation* of the group. This means saying

$$P(\text{class} \mid \vec{x}) = f(\vec{x}; \vec{\Theta})$$

where $\vec{\Theta}$ is a vector of parameters – the probability of the class given data \vec{x} can be found by using the data and parameters, $\vec{\Theta}$. This is very similar to the previous paragraph, if we realise that the parameters in $\vec{\Theta}$ are equivalent to the variables in the n -dimensional space, and the function f is the hyper-surface.

¹A hyper-surface is a surface in a higher-dimensional space.

The difference is that the hyper-plane produces a binary output, whereas our estimation yields a likelihood that the item is in a particular class. This can give information about the certainty of the classification process.

Communications theory and pattern recognition

If we consider classification as an attempt to recognise a noisy code, then it is clear that the greater the number of unique symbols being transmitted, the greater the likelihood of incorrectly classifying a symbol. This idea is illustrated in figure 4.2, and suggests that as the number of classes increases, classification errors increase.

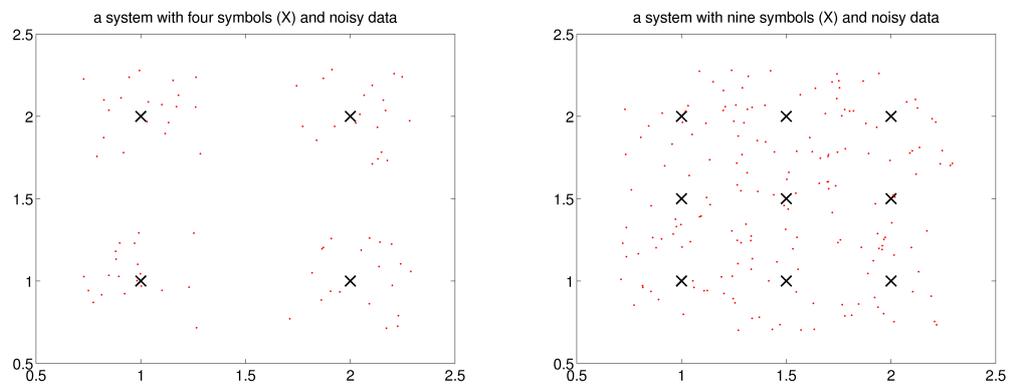


Figure 4.2: Two figures illustrating why there is more classification error as the number of classes increases

Let us imagine a system where there are two variables, from which we classify a datum. In the figures, the idealised class foci are shown by an X and sample class estimates (measured data) are shown as red dots. The left figure shows a four-class system and the right figure shows a nine-class system. It's clear that the nine-class system results in class foci being closer together, and that this will increase the rate of incorrect classification.

4.1 Linear discriminant analysis

Linear discriminant analysis, or *linear classification* is something of a blanket term, used to describe many linear classification algorithms – algorithms used to find a combination of attributes from which we can best distinguish between multiple *classes* (or categories) of data. The goal is to choose combinations of attributes so that the within-group variance of the data is minimised, and the

inter-group variance is maximised [60]. Another way of phrasing this is that we map these multidimensional data to a single dimension such that the number of standard deviations between the means of the groups is maximised, and the variance within groups is minimised. This will yield a one-dimensional group estimate that (hopefully) exhibits clusters of samples that reflect the groups.

It is clear from figure 4.1 that some data are easier to classify than others. Imagine that a random sample from an unknown group is placed upon each figure. If we wish to decide to which group this new sample should belong we can clearly do this more easily, and with greater likelihood of success, with disparate groups (left) than overlapping groups (right). This is an optimisation problem and, as is common with such problems, there are various methods by which to find the optimal solution. These methods have varying robustness, which are affected by properties of the data. The ratio of the number of samples to the dimensionality of the data is important. For data lacking sufficient independence, it is common to be analysing an under-determined matrix, which can cause many optimisation methods to fail [26].

Linear discriminant analysis (LDA) is a tool to help discriminate between groups of data. If we record a set of observations X_{train} and their class G_{train} , we can call this the training set. If we now record observations X_{test} from an unknown object, then our problem is to estimate the unknown class G_{test} . LDA makes assumptions about the statistics of the data, including that the pdfs within each group are normal.

The Fisher linear discriminant (FLD) was one of the first such algorithms to be developed. It was initially restricted to separation between two classes, but was later extended to n classes [60]. The FLD finds a linear combination of the variables of the measurands that maximises inter-group variance relative to intra-group variance. The data are arranged into an n by p matrix, which is of dimensions *observations* by *variables* (in the parlance of our experiments, this might be *subjects* by *electrodes*). We have g groups into which we classify

data (for our purposes, let $g = 2$). The intra-group variance is described in equation 4.1, where $m_{[i]}$ is an estimation of the mean of group i . W is the covariance matrix of the observations and y_i is a vector of parameters for each sample in group i . G is an n by g matrix, which is all zeros, except that $g_{ij} = 1$ iff² the i^{th} observation is in the j^{th} group.

$$W_y = \frac{\sum_i (y_i - m_{[i]})^2}{n - g} = \frac{\|\mathbf{y} - G\mathbf{m}\|^2}{n - g} \quad (4.1)$$

and the inter-group variance in equation 4.2.

$$B_y = \frac{\sum_i (m_{[i]} - y_i)^2}{g - 1} = \frac{\|G\mathbf{m} - y\mathbf{1}\|^2}{g - 1} \quad (4.2)$$

Linear discriminant analysis is a linear process. It seeks the best projection (linear combination of variables) that will transform multi-dimensional data into 1-dimensional data in the manner that best separates the groups. However, as explained above, it does this by producing a vector of weights and finding the inner vector product against the vector of parameters describing an item. This produces a scalar value that estimates the group to which the item belongs. This process is repeated for every item to be tested. However, the output is merely a linear combination of the input, hence there are limitations to relationships between parameters that LDA can use. Because LDA is limited to first order relationships it can fit only a hyper-plane (as opposed to a hyper-surface) to separate the classes. If we desire a more sophisticated discrimination between groups (such as two areas corresponding to the same group), then we need a more sophisticated classifier.

²iff means “if and only if”

4.2 Artificial Neural Networks

An artificial neural network (ANN) represents an attempt to mimic the behaviour of organic neural networks (such as the brain) for the purpose of performing a computation. It consists of a series of interconnected neurons (a simplified, artificial model of an organic neuron).

4.2.1 The Neuron Model

As suggested in figure 5.1, neurons in the human brain have a quite sophisticated internal structure. The neuron in a neural network (figure 4.3) is analogous to a neuron in the brain³, albeit simpler. In summary, the brain is very plastic and connections between neurons are constantly reinforced or depreciated depending on use. This is thought to be the manner in which the brain *learns* [22], although it is increasingly recognised that other mechanisms are at work [21].

4.2.1.1 The artificial neuron

Compared with an organic neuron (figure 5.1), an artificial neuron is very simple.

³See section 5.1 for more information regarding the structure of the brain

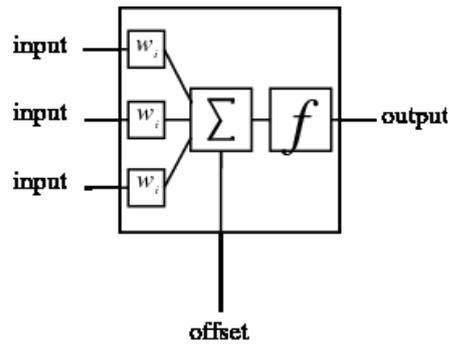


Figure 4.3: The structure of an artificial neuron

This figure illustrates the mechanics of the operation of an artificial neuron, as is used in an artificial neural network. Multiple inputs to the network are weighted and summed, and the result of that summation is passed through a nonlinear function (typically an inverse-tan, sigmoidal, or piecewise function), and outputted. The offset has the effect of controlling the tendency of the neuron to fire.

Each input to the artificial neuron is multiplied by a weight, and these are summed with an offset (the offset controls the excitability of the neuron). The summation produces a score that passes through a function to determine the output of the neuron. There are various output functions that can be used to determine the properties of the neurons, for example a hard limiter (producing a 1 if the summation is above a threshold, and 0 otherwise⁴), logarithm and a tan function (these produce values that tend towards 1 or 0, but without the hard cutoff produced by the Heavyside function).

If we are to compare this artificial neuron to an organic neuron, then the weights decide whether an input is excitatory or inhibitory, as well as how strong an affect it has, and the offset is analogous to basal membrane potential. Note that in the artificial neuron above, there is nothing akin to the temporal averaging exhibited by organic neurons (figure 5.2). For this behaviour to be included, a multi-tap delay would need to be introduced, either at each input, or feeding back from the output of the summation stage to its input.

⁴Also known as the Heavyside function

4.2.2 Network Structure

An ANN is a collection of interconnected neurons. Each neuron has an input and an output, and produces its output as a linear combination of the inputs passed through a (possibly nonlinear) function. Figure 4.4 shows a simple neural network.

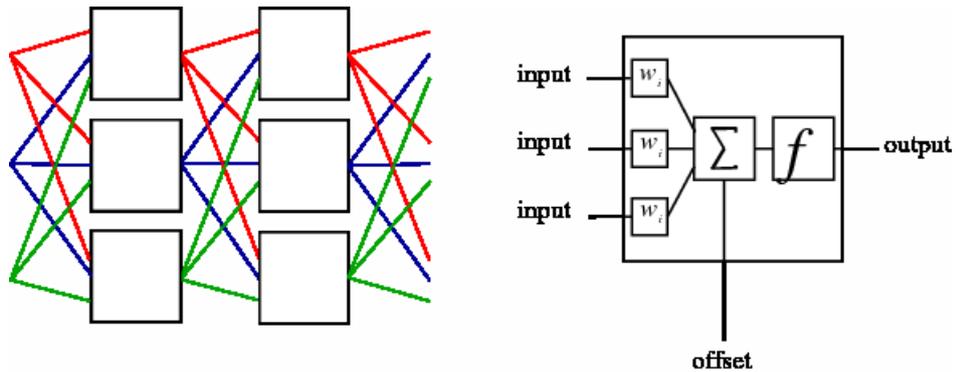


Figure 4.4: A small example circuit, showing the internal connections of a group of neurons

This figure illustrates a manner by which neurons can be connected to form a network. The example neurons have three inputs (because there are three neurons in each layer), but in practice this may vary. Each column of three neurons is a layer – the neurons in each layer receive input from the previous (left) layer, process the data, and pass output to neurons in the following (right) layer. In this way, the structure of a neural network is much more rigidly defined than an organic network, where connections may go (almost) anywhere.

Information typically moves from left to right, with inputs on the left and outputs on the right. The structure of an ANN is prescribed: there are successive *layers* of neurons (columns in the figure), where each neuron in each layer obtains input from all neurons in the previous (left) layer, and passes output to all neurons in the following (right) layer. This is a much simpler structure than that which occurs in the living brain.

The network is “trained” prior to use – a process during which the weights are iteratively adjusted to values that allow the network to “best” fulfil the assigned task. Thus, despite the fact that a neuron is connected to all neurons in the previous and successive layer, only useful connections will retain weights and contribute to the network. The process of training a network involves

presenting it with example inputs, and desired outputs, and adjusting the network's neurons' weights so that the *actual* outputs resemble the desired outputs. The network can then have the actual inputs applied, and its output observed.

4.2.2.1 Network memory

If an ANN contains a series of neurons, and there is no delay (or memory), then the output of the neural network is a function solely of the input, and the state of the weights. Once the network is trained, the values in the weights are fixed, so a given input, repeatedly applied, will always produce the same output. Such a network is described as *static*. There is another form of network where delays and/or feedback are present in the circuit. This allows the circuit to have a memory, and means that the output of the circuit is dependent on the weights, the inputs, *and historical inputs and outputs to and from the circuit*. These networks are referred to as *dynamic*. There are situations where such a circuit is preferable to a static circuit (such as modelling a temporal system, like simple harmonic motion).

A simple method by which to introduce memory into the circuit is simply by adding a tapped delay at the start of the circuit. Such a node will introduce historical inputs repeatedly. The tradeoff here is that the circuit becomes substantially more complex. If we imagine a simple circuit such as the one above, except with 10 input variables, then our weights matrix is of size $[10, 1]$ if we have only a single layer of neurons. If we introduce a tapped delay that has a memory of 5 samples, then at every instant in time, we introduce the present sample plus the previous 4 samples into the circuit. Thus, the weight matrix gains an extra dimension and occupies five times more memory, requiring many more calculations in the training process than our simpler memoryless system⁵. This would be more complicated still if we had feedback from output to input,

⁵This will vary depending on the training method, but there is often an exponential relationship between the number of network parameters, and the number of training operations.

and multiple layers of neurons with internal memory. Memory can also be incorporated into networks in many ways, including feedback loops at the neuron, layer and network level, or at the input or output of individual neurons.

4.2.3 Network Training

Training is the process of taking example input data, associating it with an example output data, and configuring the network weights so that the output of the network, given the example input data, best matches the example output data. Training is a complicated and poorly understood aspect of ANNs. There are many different methods of training, and different methods can be optimised for different outcomes. The different methods also have different computational and memory requirements.

The process of training is very closely related to that of control of an autonomous vehicle to seek the lowest point in a terrain. We want the vehicle to find the lowest point in the terrain, but we want it to take the most direct route there *and use the fewest processing cycles possible*, and the vehicle cannot “see” the terrain – it can only “feel” the slope. The algorithm is trying to optimise the weights of our network, and we would like to optimise our algorithm to achieve this.

4.2.3.1 Back-propagation

Back propagation is a training algorithm that uses gradient descent. A simple implementation is to adjust the weights in the direction that the performance function decreases the most. This will tend to minimise the performance function, but it also tends to become “stuck” at a local minimum – which may not be the global minimum (the *best* that the network can model the system). This occurs because, at each iteration, the simple gradient descent algorithm adjusts the weights in the negative of the gradient of the performance function. At a local minimum, there is no gradient, so there is no adjustment that

is made to the weights, and the training process stops. It is quite possible that the local minimum is nowhere near the global minimum. To avoid this problem, there are several alternatives:

The most common is to apply a concept of momentum. This means that if the weights of the network are changing in a particular direction, then they will tend to keep changing in that direction. This can allow a network to keep moving past a local minimum toward a global minimum. This behaviour is somewhat akin to a low-pass (moving average) filter, and will tend to result in better convergence of the network.

4.2.3.2 Supervised vs unsupervised training

Supervised training involves presenting the network with an input, as well as the desired output. The network processes the input, compares its output with the desired output and then alters the weight matrices so that the network's output more closely matches the desired output. This process is then repeated many times, with the network's approximation of the desired output iteratively improving. It is referred to as "supervised" because there is a teacher – the teacher being the goal data (desired output).

Unsupervised training occurs when there are no goal data. The network can respond to this in two main ways.

The first is for the network to self-tune to the environment. Once this is accomplished, the network can group the data into classes by making associations between the data. This process is also known as *clustering*. Training such as this often uses a layer of neurons that compete with each other, to see which best models the data.

The second method can also be known as reinforcement learning [22] and involves the use of a critic to produce an analysis of the environment data. From the input and the output, the system produces a performance metric, which it aims to minimise. It does this across time, so there is an inherent

delay to the learning. A cost function, derived from the performance metric, is calculated over time. As the system makes changes to its internal parameters, it reexamines the cost function and the effect these changes have, and then aims to iteratively adjust the parameters to produce a more favourable cost function. This is difficult to do because there is no teacher, and also because it can be difficult to know which parameter adjustment has had which particular effect on the cost function. Despite this, there have been several implementations of such algorithms.

4.2.3.3 Parallel and Series-Parallel Networks

As stated, a dynamic network has delayed feedback from the output to the input, although many dynamic networks also have internal tapped delays. Also mentioned was that dynamic networks have a tendency to be more complicated than static networks. If we make a dynamic network that has only delayed feedback from output to input (without the internal delays), then what we really have is a static feed-forward network, but with a delay, as shown in figure 4.5.

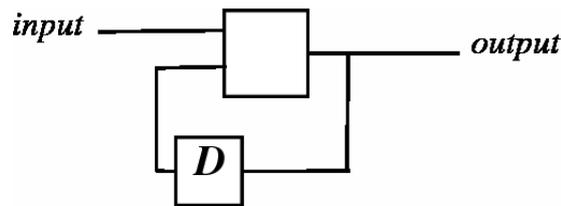


Figure 4.5: A series configuration for use of a trained neural network. This configuration is used when simulating the network. The outputs are passed back to the input, and in this way the system has memory.

When training, however, we are able to make use of a trick: *we already know what the output of the network **should** be*. Hence, during training we can simply remove the feedback from output to input, and insert the goal data into the 'feedback' input, as shown in figure 4.6.

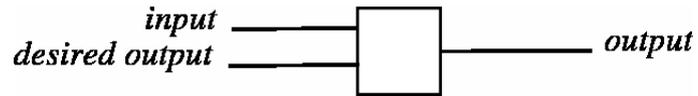


Figure 4.6: The same circuit, but in a parallel configuration to facilitate network training

For the purpose of training, the second input (to which the output will be routed during simulation) is passed the actual goal data.

Doing this results in more effective training, because we're training with the *real* output data. Also, the network is simpler to train because it is feed-forward only.

Chapter 5

Brain background

This chapter introduces the brain. Firstly, in section 5.1, there is a description of what the brain *is*, what it is made of, and how it is constructed. In section 5.2 I discuss the origin of EEG, the conclusions that can be drawn from EEG, and some of the limitations of these conclusions. I introduce epilepsy in section 5.3, describe what it is, the limitations of our current understanding, and several aspects that are currently under investigation by the scientific community. Also introduced is the concept of cell-swelling and brain impedance (Z) and the relationship of this quantity to EEG and epilepsy (section 5.3.3). I conclude the chapter (section 5.4) with an examination of general analysis of EEG data, which can provide assistance when attempting to understand and elucidate the mechanisms at work in the epileptic brain.

5.1 Anatomy

The human brain is quite possibly the most complicated system we know of. It is composed of many billions of neurons, intercommunicating by means of chemical neurotransmitters. Aside from neurons, there are hundreds of different cell types in the brain, most of which are not thought to be involved with cognition. Despite the myriad cell types, it is neurons that are the principal contributors to EEG (section 5.2).

5.1.1 The neuron

A neuron is commonly regarded as a building block of the brain's processing ability. It consists of a cell body or *soma*, dendrites and an axon.

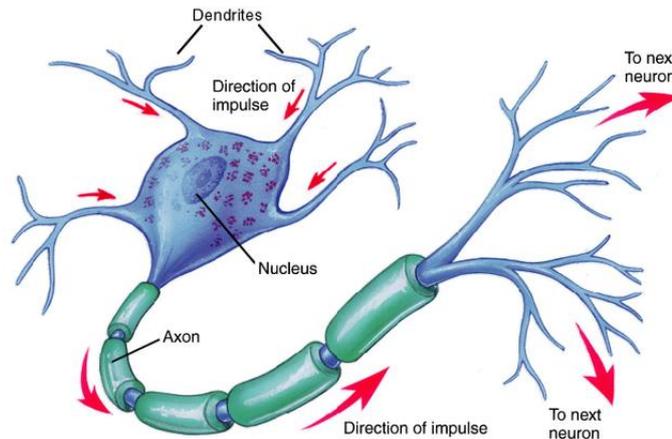


Figure 5.1: A typical representation of a neuron

<http://www.mhhe.com/socscience/intro/ibank/ibank/0002.jpg>

This image is an artist's representation of a neuron. It shows: the cell body, where most of the signal processing occurs; the dendrites, where signals are acquired; and the axon, where the outputs signals are conducted to other neurons. Notice the myelin sheaths on the axon – these are present to increase the conduction velocity through the axon.

Dendrites acquire input to the cell, and axons express the cell's output. In the brain, input and output are usually from and to other neurons. Most neurons acquire input from other neurons – they usually attach at the neuron's dendrite, and obtain input via their axon. Neurons maintain an ionic potential across their cell membrane by means of ion pumps – this is referred to as its resting potential. The osmotic gradient is regulated so that there is a high concentration of sodium ions (Na^{2+}) within the cell, and a high concentration of potassium ions (K^+) outside. A neuron's function is to *depolarise* when supplied with appropriate stimuli via its dendrites, and “transmit” the depolarisation as a stimulus to other neurons who then may, if conditions are appropriate, depolarise also.

The action potential

The resting state of the electric potential across the cell membrane is such that there is an excess of negative charge on the interior of the cell. Thus, there is a negative electric potential across the cell membrane. Incoming neurotransmitters locally affect this potential. They can either depolarise the cell (increasing the potential towards a positive potential) or hyper-polarise it (decreasing the potential towards a stronger negative potential).

The cell membrane of a neuron contains voltage sensitive ion channels whose permeability to their specific ion (sodium or potassium) is related to the electric potential across the cell membrane. At the resting potential (the default ionic imbalance across the membrane) these gates are closed, but if the potential increases above a threshold they will open, allowing ions to pass through the cell membrane. Sodium channels respond most rapidly to these changes in electric potential. They open, and allow sodium ions to enter the cell. This further increases the electric potential resulting in more sodium channels being opened – a chain-reaction. There is a rapid change in the cell's membrane potential – a process is called *depolarisation*, and this propagates across the cell soma as more channels open.

There are also potassium channels whose permeability is related to the electric potential. They operate in the same way as the sodium channels, except that their opening is slower than that of the sodium channels. The process of depolarisation occurs before the potassium channels open, but when they do, potassium ions flow out of the cell, reducing the electric potential back towards the original level. This process is called *repolarisation*, and includes the time taken for the cell to re-establish the resting ionic balances, during which it can not depolarise again. The combination of rapid depolarisation and repolarisation is referred to as an action potential (see figure 5.2), and the complex manner determining whether a neuron will depolarise is thought to be the manner in which neurons process information.

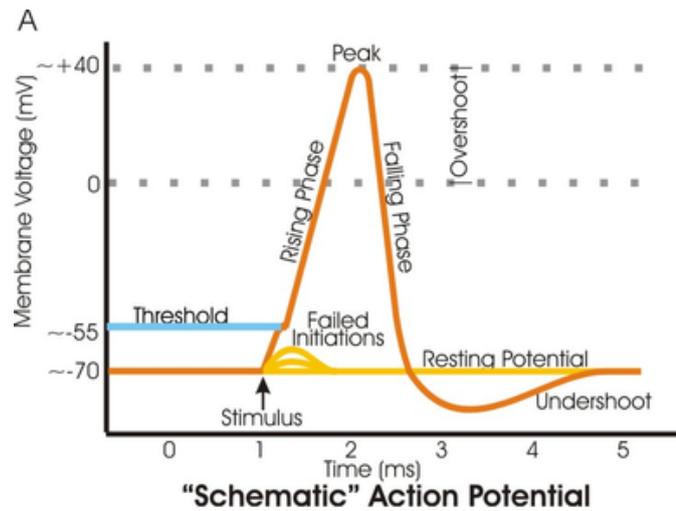


Figure 5.2: An action potential

http://en.wikipedia.org/wiki/Image:Action_potential_vert.png

This figure shows an idealised action potential. When the membrane voltage increases past the threshold voltage, an action potential occurs. When the threshold is not passed, the potential returns to the resting level. The rising phase is caused by the influx of sodium ions, and the falling phase is caused by the exodus of potassium ions (through the slower-opening voltage-gated potassium channels).

Whether or not a cell depolarises depends on whether its electric potential crosses a threshold. A cell can have inputs from many other cells, and the neuronal process is summative across cell inputs. It is also summative over time, because the electric potential will not instantaneously return to a resting level in the absence of input. Also, the cell has different sensitivity to input at different input locations. Finally, recall that a given cell can be encouraged or discouraged from depolarisation by an input, because inputs may raise or lower the membrane potential. The dendrites of a neuron can receive input from thousands of other cells [5], and each input can affect other inputs. Hopefully this conveys the idea that whether or not a cell will produce an action potential is a complex process, dependent on the number and type of stimuli reaching the cell, as well as their proximity and temporal relationships, and whether the cell has recently depolarised.

When an action potential is produced, it propagates down the cell's axon – a long conduit used to route the signals to other neurons. Axons have evolved

so as to maximise the rate of transmission of the action potential. To increase this rate, the axon is wrapped in a myelin sheath, which provides electrical and magnetic insulation. The action potential propagates along the axon until it reaches an axon terminal, where it triggers other membrane-potential sensitive gates. These gates release neurotransmitters into the extracellular space, where they diffuse across the small gap between the neurons (a *synapse*), and affect the electric potential of the next cell. If the polarity of the next cell is sufficiently affected, then the next cell will depolarise also.

5.1.2 Non-neuronal cells

There are three main types of glial cells. The first are astroglia and are important to neuronal function. There are also oligodendroglia, which produce myelin, and microglia that are associated with immune function.

Astrocytes are the most numerous type of glial cell, and are known to perform a range of functions, including metabolic regulation and ion concentration regulation – astrocytes help remove potassium ions, glutamate, and other neurotransmitters from the Extra-cellular fluid (ECF) [86]. Astrocytes are also responsible for vasomodulation, myelin sheath production and structural support. Astrocytes are ten times more numerous than neurons, although the extent to which they outnumber neurons varies across species. One function of glia is to encircle neurons – providing a supporting role to neuronal function. Astrocytes prevent the escape of neurotransmitter from synaptic junctions, and actively remove neurotransmitters from the extracellular fluid. For many years they were thought to be “support cells to neurons” but without any cognitive input. This theory has been questioned, and it now seems likely that astrocytes can modulate inter-neuron communication [24] and affect the behaviour of neurons. Astrocytes and neurons are now thought to communicate via the release of glutamate, ATP and calcium. For example, it has been found that astrocytic cell membranes have neurotransmitter receptors that can affect the

interior of the cell and that they can release substances that excite or inhibit neurons [86].

It is increasingly thought that astrocytes are important in the pathophysiology of epilepsy. It has been found that astrocytes in epileptic foci exhibit functional changes relative to other astrocytes. Another example is Lanerolle et al [10], who investigated temporal lobe epilepsy (TLE), which (in the specific case of mesial temporal lobe epilepsy) is characterised by focal seizures originating in the hippocampus. Because of this, a common successful treatment to control the seizures in TLE is the surgical removal of parts of the hippocampus and amygdala, which not only results in seizure control but also improves metabolism in other areas of the brain. Thus, it has been proposed that astrocytes may increase excitability of neurons in the hippocampal focus, by changes in their membrane ionic potentials, and increased sensitivity to glutamate. It has also been shown that in the sclerotic regions, astrocytes are less efficient at converting glutamate (a by-product of depolarisation) to glutamine [14] – a process that occurs only in astrocytes.

Sclerotic tissue in people with TLE has reduced numbers of neurons and additional astrocytes, when compared to normal tissue - the astrocytes are essentially forming scar tissue. This relates to the idea that astrocytes provide the structure in the brain [5]. Astrocytes also form networks connected via gap-junctions, that regulate astrocytic behaviour. It has been proposed that the astrocytic networks and their connections with neural networks might influence epileptic events [86].

5.1.3 Neural circuits

The human cortex has approximately 10 billion neurons, each of which are connected to approximately 6000 other neurons via synapses. The connections, and the strength of the effect they have on the neurons, are quite plastic – this allows the brain to adapt its structure to the demands of its environment,

making it very flexible and adaptable. To aid our understanding of the brain, it can be divided into regions that specialise in various tasks. For example, there are areas of neurons that are specialised for controlling motor actions (the motor cortex) and areas responsible for sensory perception (the sensory cortex).

The largest concentration of neurons in the cerebrum, and the area that performs most function in the brain, is the cerebral cortex [81]. This is a thin layer on the surface of the brain, usually less than 4mm thick. It is thought that the convoluted shape of the brain exists to maximise the surface area of the brain, and hence maximise the area (and computational power) of the cerebral cortex. In general, the more intelligent an animal, the greater the number of folds in the cerebral cortex, although the density of the neurons there remain relatively constant between species. This is somewhat of a simplification, since the cerebral cortex in mammals is referred to as the *neocortex*, and contains more layers of cells than the cerebral cortex found in more ancient animals. A more accurate statement is that, over the course of mammalian evolution, the amount of cortex has changed, “but its basic structure has not” [5].

There are at least 6 identifiable layers of cells in the cortex, each of which shows an identifiable structure. There is a cell that is larger, and forms long-distance connections between disparate layers. This type of cell is called the pyramidal cell (sections 2.2.8 and 5.2 describe the contribution of pyramidal cells to EEG), because of the shape of its soma. Pyramidal cells are arranged radially within the neocortex.

5.1.4 Brain Regions

This chapter began by looking at the components of the brain, and then at the low-level structures that these components form. We now examine some of the higher level structures in the brain. It was recognised in the early 20th century that the brain is organised around functional groups.

For the purpose of the work in this thesis, there is no need to examine low level brain regions. For this reason, we briefly describe the four main regions of the brain.

Frontal lobe: This area can be viewed as the area of cognition. It is involved with judgement, memory, problem solving and planning.

Parietal lobe: This area integrates sensory information.

Temporal lobe: The temporal lobes process auditory information, resolving memories, and is involved with establishing meaning.

Occipital lobe: This area is responsible for visual processing.

These four areas are present in both hemispheres of the brain, giving rise to eight main brain regions (see figure 7.1).

5.2 EEG

An electroencephalogram, or EEG, is the recording of electromagnetic fields radiated by the brain during the course of its operation. Neurons operate by the movement of charged ions within the brain and when these dipoles rapidly change (as they do during an action potential) an electromagnetic field is produced. Since the orientation of the neurons will affect the orientation of the electric field they produce, and observation from a distance will be a summation across many cells, random orientations in the cells will tend to cancel out the summative electric field to nothing. However, there is a class of neurons that are all oriented in the same direction relative to the cortical surface – these are known as pyramidal cells (figure 2.3). Because of their uniform orientation, a summation of the electrical field over many such cells will reinforce, and for this reason EEG is composed, almost exclusively, by signals generated by pyramidal cells.

There are several main methods of obtaining EEG recordings. In humans, the most common method is to use surface electrodes, which reside on the surface of the scalp. A second method is to remove a portion of the skull, or

drill a hole, and place the electrode on the surface of the brain. This can rest on the dura (a protective membrane between the brain and the skull) or it can penetrate the dura and rest directly on the surface of the brain itself. A recording made using this method is often referred to as an electrocorticogram (ECoG) [81]. A third method is to insert needle electrodes into the brain itself; these are usually called depth electrodes.

The only way to measure the activity within a single cell is to use a micro-electrode, and insert it into the cell. As a general rule, data obtained from surface electrodes are broader (meaning that it is a summation, a more global measure, of a greater number of cells) whereas needle electrodes provide finer spatial resolution (they record from fewer neurons).

So in general, EEG is a measure of a summation of electrical activity amongst many cells, and is mostly attributable to pyramidal cells because they are similarly oriented [81]. Because neurons aren't necessarily doing the same thing, most of the information is "lost in the noise" – much like a single spectator's yells at a football game. The structures that we can perceive in EEG are the result of synchronisation between many neurons – the behaviour of an individual neuron is imperceptible.

Because of the extremely complicated nature of the brain, it could be favourable to adopt a reductionist approach, however it is likely that this would only provide highly specific information, and that for general studies a more holistic approach would be required. This would mean examining the function of the brain at all levels of operation, from EEG to individual neurons – which would mean using needle electrodes on a per-neuron basis. In the future such an approach could be possible – for example, an advanced fMRI¹ might be able to resolve action potentials in individual neurons in real time.

¹Magnetic resonance imaging (MRI) uses electro-magnets to create images of tissue. fMRI is a variant of this that is designed to image neuronal activity in the brain.

5.2.1 Limitations of EEG

That EEG is a summation of the electrical activity of many neurons is somewhat misleading, because it is primarily the action of the neurotransmitter reception that is measurable – the conduction of the action potential is almost undetectable. Clearly, the EEG is not *the state of the brain*. However “decades of empirical observation indicate... ..statistical attributes derived from EEG reflect and track the underlying state of the brain” [58]. This means that a non-invasive recording and analysis technique is often able to provide information about the underlying, and apparently invisible, operation of the brain.

Electrodes placed outside the skull experience significant attenuation, due to the presence of the skull and other membranes, so dural or sub-dural recordings will tend to produce better results. Let us, for a moment, consider an ideal case where we have a number of EEG electrodes that approaches the number of neurons, and each electrode records EEG data only from the neurons directly below it. If we contrast this to the setup we are likely to have, the difference is marked. Firstly, each electrode will be recording data from many many neurons, and secondly there will be significant overlap between the “field of view” of various electrodes (figure 2.3). To rephrase this, there is much smearing of data between electrodes, which gives data recorded in adjacent electrodes a coherence that does not necessarily reflect the activity of the brain, but rather is a measurement artifact. There are several methods by which we can (at least partially) circumvent these issues. One such method is the Laplacian (section 2.2.8).

One big limitation of EEG, is that it does not record data from glial cells. This is because glial cells are randomly oriented, unlike pyramidal cells, so that even coherent behaviour across many glial cells is likely to be smeared to nothingness in EEG. For a long time, it was assumed that glial cells did not contribute to information processing because it was known that they played a supportive role to neurons. This is starting to be questioned, however because

of difficulties in recording electric signals from glial cells, significant progress in understanding the roles of glial cells is yet to be made.

5.2.2 Patterns

It should be emphasised that what all analyses (both linear and nonlinear) are really trying to achieve is the identification of patterns in EEG that correspond to macroscopic, comprehensible occurrences in brain activity. This is based upon the (implicit) assumption that there are detectable patterns within the EEG that can provide information about the underlying operation of the brain.

There are some interesting ideas *emerging* in the examination of patterns. This is necessarily a very broad field, but it is also, by its nature, broadly relevant. In this review, I have included it in the context of the operation of the brain – although it could have easily been included with analysis algorithms.

One relatively new idea is that complex systems are able to self-organise, and that the brain is likely to be such a system [33]. Other studies have suggested that chaotic patterns in the brain are normal, and Klonowski et al. [34] even went so far as to say that “it is healthy to be chaotic”. This alludes to the increasingly held opinion that the operation of the brain is inherently chaotic, and that any deviation from this is pathological. The ability of a system to self-organise is dependent on the rules governing the system, however for complex systems, behaviour of the emergent properties is not necessarily obvious. A general concept is that for such behaviour to occur, the system needs to be in an unstable state away from equilibrium – a state universal in living, biological organisms.

Other new work involves the construction of a systematic method of developing, organising, documenting and interfacing software patterns. In one sense this can be seen as a more general form of object-oriented code. One of the principal advocates of this paradigm is Coplien [9]. He often discusses the more universally applicable aspects of this theory, and his paper has many

references to real-world applications of pattern theory – such as bridge-road-environment networks. An interesting idea is that the concept of a pattern containing smaller, more specific patterns and being contained within a larger, broader pattern is similar to the fractal structures often exhibited by chaotic systems. This concept of repeating patterns at differing scales is related to the idea of *fractals*.

To illustrate the relevance of patterns and fractals, let us consider a flock of birds. Birds, when flocking, obey simple rules from which complex flocking behaviour emerges. Rules like:

- Don't make sudden moves
- Try to maintain a constant distance between neighbours

Imagine that several birds on the left-side of a flock see some food. No other birds in the flock see it. These birds turn (slightly) towards the food, which causes their neighbours to also turn slightly. The neighbours then see the food also, and thus turn even more towards it, causing their neighbours to turn towards it, etc. In this way, the emergent property of flocking is that there is a sharing of information across the flock – very few birds need to perceive something (threat or goal) for *the flock to perceive it* – the flock can almost be seen as a single organism! But, notice that this idea of “information sharing” is not codified into the rules that the individual birds are obeying – this is *emergent behaviour*.

In the same way, the neurons in each bird's brain are governed by simple rules, but the outcome of those rules is a complex system whose behaviour cannot be described or predicted from the simple rules – it is emergent behaviour. Notice, also, the parallel between the relationship of the bird to the flock, and the relationship of the neuron to the brain – it is a repeating *pattern*. Concepts such as this are discussed by Hofstadter [25].

5.3 Epilepsy

Epilepsy is a neurological disease that is often detrimental to quality of life. Its main manifestation is as seizures, which are periods of abnormal and dysfunctional brain activity during which cells fire with an abnormally high level of synchronicity. Epilepsy is an umbrella word used to describe a group of pathological conditions where the brain periodically expresses such abnormal function. The causes of epilepsy are not always well understood. It can be caused by damage to the brain (such as head trauma, stroke, infection) or it can appear with no obvious cause. Recent research has shown that some people possess a genetic predisposition as well [73, 63], and there are hypothesised links between epilepsy and migraines. A person who has epilepsy has their life affected in several ways, perhaps the greatest being the risk of injury each time they have a seizure involving loss of consciousness. This means that, in general, they cannot drive cars or operate machinery.

Approximately 0.8% of people exhibit symptoms of epilepsy and of these, 20% are not significantly assisted by drug therapy [48]. This is mainly because many of the drugs currently prescribed are not selective for the problem and have the effect of reducing action within the brain as a whole, so that the incidence of seizure is reduced. Particularly for serious epileptics, there can be unwanted side effects, such as an overall dulling of awareness. Drugs which worked on the actual problem would be better targeted, more efficient, and would have fewer general side effects. Alternatively, if a person with epilepsy had 1-2 minutes of warning before the onset of a seizure, then they would be able to drive a car, quite aside from any of the plethora of other safety and therapeutic benefits having this type of prescience would convey.

There are several different kinds of seizure, and there are several different ways in which they can be classified. If we divide them according to their appearance to an external observer:

1. Absence seizures (historically called *petit mal*) are characterised by dis-

traction of the sufferer. They will pause, and possibly stare, and will later resume as if nothing had occurred (consciousness is affected). These are generalised seizures, meaning that they occur across the whole brain.

2. Tonic-clonic seizures (historically called *grand mal*) are generalised, and have an initial muscle contraction phase, followed by rhythmic muscle contractions (consciousness is affected).
3. Partial seizures cause aura and other focal symptoms, which can be perceived in many ways, and occur when a partial seizure affects a motor or sensory area of the brain. Partial seizures affect only one area in the brain, and can cause the person to perceive a range of sensations (including smells, taste, movement, light, sound) or make uncontrolled muscle movements of regions of the body.

Primary generalised seizures involve the whole brain and are a dramatic illustration of the behaviour of a complex system. The brain can be operating, apparently normally, and can suddenly “switch” into a seizure – this can happen across the entire brain, simultaneously. In systems terminology, this could be referred to as a phase boundary, and it means that when the parameters of the system cross a threshold, the behaviour of the system can change dramatically.

Partial seizures occur only in one part of the brain, and will generally impair the function of that part of the brain. Hence, a partial seizure that is occurring in the visual cortex can cause the person to experience visual hallucinations, whereas a partial seizure in the motor cortex may cause involuntary muscle contractions. A partial seizure can spread to other parts of the brain, sometimes resulting in a generalised seizure. Such a seizure would be termed a secondary generalised seizure, to indicate that it was initially a partial seizure.

5.3.1 Epileptic seizure anticipation

Previous research has been directed towards the anticipation of epileptic seizures by nonlinear analysis. As mentioned, epileptic spindles and seizures seem to arise spontaneously, and with no apparent reason. Some studies have concluded that a deeper characterisation of the underlying neuronal dynamics would allow new insights into the mechanisms of seizure generation on a longer time scale [57].

The use of nonlinear tools in the analysis of EEG has not been restricted to epilepsy studies, however. A paper by van den Broek discusses the use of nonlinear tools (such as correlation dimension and dimensional complexity – a modified correlation dimension) and concluded that “new parameters from the field of nonlinear dynamics can be an aiding tool in detecting effectual changes induced by anaesthetics” [11]. This illustrates the rapidly increasing importance which nonlinear analysis tools are playing in the field of neuroscience.

Klonowski et al [34] made 16-channel EEG recordings from human subjects, and from these calculated a range of standard quantifiers such as autocorrelation function, Lyapunov exponent and embedding dimension. They were interested in the effect of theory on these quantities, but found no consistent pattern of change. They speculated that quantities which were able to be calculated directly from the EEG, rather than from the embedded phase space description (sections 3.3.1 and 3.3.1.1), might be more efficacious and they concluded by suggesting that fractal dimension might be a useful quantifier.

Lehnertz et al [42] argue for the presence of a third state, between the inter-ictal and ictal states, heralding an imminent seizure. They introduce a method of measuring temporal changes in correlation dimension in localised areas of the brain, to attempt seizure anticipation in patients with mesial temporal lobe epilepsy. They call their measure neuronal complexity loss, and it allowed them to identify the location of the epileptogenic area in all ten of their examined patients. They also found that for patients for whom

surgery was not successful, there was a more diffuse localisation of estimated epileptogenic area, indicating that there were several foci for seizures.

5.3.2 Epilepsy diagnosis

A robust clinical diagnosis tool for primary generalised epilepsy (PGE) relies on witnessed generalised seizures or a diagnostic EEG seizure recording – these are difficult to obtain. A diagnosis that did not require this would be very useful, but does not yet exist in a usable form.

Work has proceeded in the identification and diagnosis of focal epilepsies [72] however, classification of PGE remains an intractable problem. My colleagues, Willoughby et al [84] were thought to have found evidence of non-ictal abnormalities in subjects with PGE in an exploratory study, however further investigations have yet to yield a useful diagnostic tool.

5.3.3 Cell Swelling

Rapid depolarisation and repolarisation in neurons results in an osmotic gradient across the cell membrane. This causes water to diffuse into the cell, to equilibrate the internal and external ionic concentrations. As the neurons and astrocytes swell with incoming water, they occupy a proportionally greater amount of the volume. This causes an increase in electrical resistance of the extracellular fluid (ECF), and a decrease in the resistance within the cell – note that the resistivity of the ECF does not change. Changes in resistance are due to changes in volume. These are not the only changes that occur though, since cell-swelling will also alter ionic concentrations, changing the resistivity of the various components. This is quite a complicated system, and much work has sought to understand the mechanisms involved. There have been several studies that examine the relationship between cell swelling and epilepsy, and some of these are mentioned in [54]. Impedance has also been used to examine other phenomena, such as hypoxia [43], although they are not

examining cell-swelling, but rather a (hypoxia-induced) cerebral edema.

5.3.4 Relations between cell-swelling and impedance

The brain can be viewed as neurons resting in a bath of extracellular fluid (ECF). The neurons' cell wall is a lipid bilayer, which can be modelled electrically as a capacitor. The surrounding liquid is predominantly resistive, as is the inside of the neuron. As an electric model, this is a resistor (the intra-cellular fluid) in series with a capacitor (the cell wall) all in parallel with another resistor (the ECF) (figure 5.3). Hence, if we pass a lower-frequency electrical signal through the system and measure the impedance, we will be measuring the impedance of the ECF. If we pass a high-frequency signal however, we will be measuring the impedance of the whole system. This idea is illustrated in figures 6.2, 6.3 and 6.4.

As the cells swell, they occupy a greater proportion of the tissue – with a corresponding reduction in the volume of ECF. Thus, any current passing only through the ECF (a low frequency current) must take a more convoluted path between the cells, travelling further. The current will also be condensed into the (now smaller) spaces between the cells. Both of these result in an increase in impedance. Hence, there is a relationship between cell swelling and impedance.

Schwan and Kay [69] performed some early work examining the electrical resistance of various body tissues, in situ. In particular, they looked at the effect that blood had on measurements since its conductivity is, in general, about ten times greater than the conductivity of tissue. For our experiment [54], we have assumed that blood is 40% cells, and the brain is 80% cells by volume – which allows us to estimate that brain tissue has a resistivity approximately three times that of blood. We estimated that a cell volume increase of 10%, caused by an influx of electrolyte into the cells, would cause a 66% impedance increase, which would be measured as a 30% impedance

increase due to current also passing through the blood. Our apparatus lacked the ability to identify non-homogeneous cell expansion.

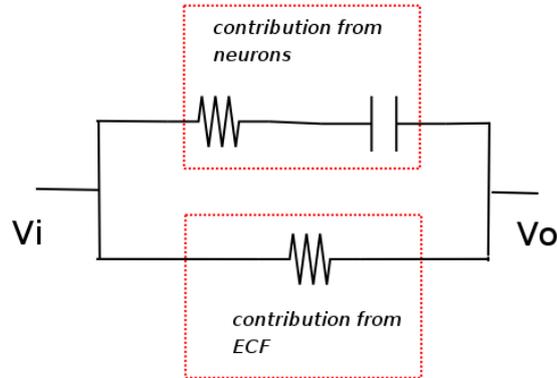


Figure 5.3: An electrical model for brain impedance

This figure is a theoretical electrical model of the brain. It considers the brain to be comprised of neurons as homogeneous black boxes, resting in an ionic bath of isotropic concentration. The upper branch of the model represents the neurons. They consist of a lipid-bilayer cell membrane, which acts as a capacitor, and an internal ionic solution which acts resistively. The lower branch represents the ECF, which is an ionic solution and hence behaves as a resistor. At relatively low frequencies the impedance of the capacitor is very high, so the measured impedance will be the impedance of the extracellular fluid. At high frequencies, the capacitor has a low impedance, and the circuit is effectively two resistors in parallel (the resistance of the extra- and intra-cellular fluid).

More detail regarding this subject, as well as information regarding the hardware that were designed to estimate the impedance can be found in section 6.1.4.

5.4 Nonlinear brain research

There has been much use of nonlinear analyses in other areas of brain research. Frank et al [18] have made investigations involving recordings of EEG and of the subjects tapping their finger. In particular, the relationship between behaviour and neo-cortical activity is examined. In one experiment, a subject was asked to tap their finger between audible stimuli (off-beat). Every 10 beats, the beat-to-beat interval was decreased. At a particular frequency (usually ranging from 1 to 3.2Hz) the subject was unable to continue, and switched

to on-beat finger tapping. They speculate that the neural control system undergoes a phase transition at that point. While the finger-tapping was occurring, MEG (magnetoencephalogram) was recorded, and the data gained from this was analysed using a variety of linear and nonlinear tools. They found statistical phase-locking, and it was interpreted as part of the self-organising nature of the neuronal oscillators.

It is also important that people develop algorithms for the purpose of modelling cerebral function, eg. Scheler [64]. She makes the observation that projecting a high-dimensional space into a single dimension (such as time) can only lead to a huge information loss – an idea that has quite serious implications for Takens’ embedding theorem, see section 3.3.1.1. This seems obvious, however it is easy to take for granted the recording of electrical signals within the brain as “how it is done”. Scheler’s idea is that by understanding the interactions between neurons at a low level (i.e. between few neurons) then the detail recorded from the brain (by whatever method) will not lose its multidimensional information.

5.4.1 Mental task classification

Watanabe et al [80] use a Lempel-Ziv compression² to assess the complexity of the data, in an attempt to distinguish between different mental tasks. They found a change in binary complexity of the signal affected by whether the subject had their eyes open, closed, or they were performing mental arithmetic. The complexity of the signal can be interpreted as the possible reduction in message size due to compression, or the size of the dictionary required to encode the data – both of which are related.

Automated mental task classification is related to the development of a brain-computer interface (BCI), an application on which much effort is focused, and one that is related to epilepsy research (both involve analysis of EEG

²Lempel-Ziv is a dictionary-based lossless compression mechanism (the algorithm upon which gzip is based).

data, and can also involve classification and prediction algorithms), however a discussion of BCIs is outside the scope of this thesis.

5.4.2 Modelling of the brain as a system

Breakspear et al [7] have developed a theoretical dynamical model to describe the behaviour of the brain during PGE seizure. In seeking a formal description of what occurs during a seizure, they developed a model that exhibits bifurcations that replicate real brain behaviour. They suggest that healthy brain operation reveals only “occasional and weak nonlinearities” in an otherwise linear stochastic system, and that it is pathological brain states that cause it to manifest a strong nonlinear component. In this way, the transition from linear to nonlinear can be seen as a bifurcation (refer to figure 3.1 for an illustration), and different seizure types are represented by different nonlinear dynamics. Once they had developed their model, they recorded EEG data, and then used nonlinear analyses to compare the model with the empirical EEG data. Breakspear et al used a nonlinear predictor to test for nonlinearities, where small errors suggest a good fit and hence nonlinearities in the data. They also applied a surrogate analysis (using phase randomisation, section 3.4.2) to this algorithm to assess the null-hypothesis that any difference was due to stochastic linear behaviour.

Their model was *relatively* simple (it can be described by 8 first-order differential equations) and consisted of cortical neurons interacting with neurons in the thalamic reticular nucleus and specific thalamo-cortical relay neurons. There were, however, many parameters controlling the behaviour of the various neuronal structures. By controlling the excitability of the the cortical pyramidal cells in the model, they were able to affect system stability, and found a bifurcation at a supercritical instability. Beyond this, they observed increasing amplitude oscillations – after a few cycles, slow wave oscillations began to appear. As the parameter representing pyramidal excitability is de-

creased, the amplitude of these waves decreases, and the system returns to its original pre-ictal state. They found that their model accurately represented an absence seizure, however a tonic-clonic seizure model did not perform as well. They suggest that their model explains the asymmetry between pre- and post-tonic-clonic seizure EEG, attributing it to the large parameter space that must be traversed before “a stable linear regime reemerges.”

5.4.3 Spreading depression

The spreading depression (SD) was identified by Leão in 1944 [40], and it is a slowly moving “wave” of electrical inactivity that passes across the neocortex, reducing EEG activity. It can be caused in several ways, including electrical stimulation [40], physical stimulation (our own experiments) and the application of various neurotoxins (such as potassium chloride (KCl) [32]). It can also be induced by large injections of saline solution (presumably affecting ionic concentrations within the brain). SD leads to a heightened level of Fos³ [23, 8] and a heightened impedance (refer to section 5.3.3), although both of these findings are counter-intuitive, since the EEG power is suppressed. It is thought that SD is linked to migraine [8] and epileptiform activity [54].

The spreading depression is initiated at one location, but will then propagate slowly across the cortex, at approximately 2 - 4 mm/min [8], and the mechanism of propagation is unknown. Larossa et al [38] have also shown that there is a measurable neuronal synchronisation that precedes a spreading depression, that there are oscillations up to 1 mm ahead of the spreading depression onset. They found that SD could be retarded by blocking glutamate receptors or by preventing ATP production in astrocytes using fluorocitrate (a

³Nerve cell activation initiates several events, one of which is a possible depolarisation. Another is the expression of certain genes, which can be expressed within minutes of stimulation. c-Fos proto-oncogene is such a gene, and has a rapid and short-lived expression following nerve stimulation [23]. The presence of the c-Fos protein in a cell, and its relative abundance, can be ascertained using immunohistochemistry, and gives an indication of nerve activity.

metabolic poison that selectively targets astrocytes, blocking the TCA cycle⁴).

This means that although SD causes a decrease in EEG power, it causes an increase in Fos (indicating increased activity), an increase in cell swelling (also indicating increased activity), and is partially caused by glutamate⁵ release from cells in the brain. The meaning of this is not yet clear, but it is possible that a spreading depression represents increased neuronal activity, but decreased synchronicity (and hence a lower-power EEG), which would also provide a link between SD and epilepsy.

⁴The tricarboxylic acid cycle is a series of enzyme-catalysed reactions that allow cells to metabolise energy. Blocking this process causes starvation of the cells.

⁵Glutamate is released during excitatory neuronal activity.

Chapter 6

Data collection

This chapter describes the three source experiments yielding the data used in this thesis. The first to be described is the rat experiment – these can be viewed as a model of human epilepsy, and are discussed in section 6.1. The second data set is sourced from the human experiment that we performed, which sought to find human parallels to the findings from previous rat experiments, and is described in section 6.2. Finally, the third source of data was the human paralysis experiment. This ground-breaking and exciting work was performed to further analyse and clarify some of the conclusions from the original human experiment. It is introduced in section 6.3.

As mentioned in section 1.1, my work represents an extension of the experiments' original methodology.

6.1 Pharmacologically-induced rat epileptogenesis

This experiment was a collaboration between Flinders University, Flinders Medical Centre, and Sahlgrenska in Gothenberg, Sweden. It involved the collection and analysis of electroencephalogram (EEG) data from rats in which epileptiform activity (epilepsy-like seizures) have been pharmacologically in-

duced – a state referred to as *acute epileptogenesis*¹.

6.1.1 Preparation

Inbred adult male Sprague Dawley rats, known not to exhibit spontaneous spike-wave discharges, were used. They were surgically prepared by Dr Marita Broberg. While anaesthetised with pentobarbitone, a right-atrial catheter (for drug administration) was inserted through the jugular vein, and holes were drilled into the skull for implantation of the screw electrodes (EEG and impedance recording), which were sized so as to rest against the dura. A reference electrode was attached anteriorly over the frontal sinus and an earth electrode to the occipital bone, each using a screw into the skull. The six electrodes were coupled directly to an IC socket, which was fixed to the skull using dental cement. The cement also covered the electrodes, wires and catheter. The end of the catheter was temporarily sealed to prevent bleeding.

Fluorocitrate and citrate experiments required a *well* for direct access to the surface of the dura. This allowed the drug to be injected directly into the rat's brain. Some experiments also involved the measurement of extracellular ion concentration (measured by an ion-sensitive electrode - ISE), which also required a well. I did not analyse ISE data. The well was never seated closer than 4mm to the nearest electrode, and was sealed with wax post-operatively. The animals were allowed to recover for 1 - 3 days, after which an experiment was performed.

6.1.2 Experimental procedure

The rat was paralysed by an injection of succinylcholine (scoline). In isolation this will often cause cardiac arrest, so a protective injection of scopolamine was given first to block the cardio-inhibitory effects. The rat stopped breathing, and was attached to a ventilator. If necessary, the protective wax in the

¹Epileptogenesis refers to the causes of epileptic activity.

rat's head-well was removed. The rat was then placed in a Faraday cage, and the pre-amplification IC was attached to the exposed IC socket. If the experiment was measuring ion concentration, the ion-sensitive electrode was inserted through the well so that it punctured the dura, and penetrated brain tissue to a depth of 1 mm – a process controlled using a micro-manipulator.

Once the signals were being recorded reliably, the epileptogenic agent was injected into the rat, via the implanted catheter. The rat's EEG was recorded while the drug affected the rat and until the rat's EEG returned to an apparently normal state. This typically required between 20 and 150 minutes, depending on the choice of epileptogenic agent (the maximum experimental duration was 150 minutes). During this time the rat was given repeated injections of scopolamine and scoline to maintain the paralysis.

The well in the rat's head exposes the brain to the air, and with time the surface of the brain would dry somewhat. Because of this, we periodically added drops of saline to the well to wet the surface of the brain. At the conclusion of the experiment, the rat was killed by deep anaesthesia, and trans-cardial perfusion. After death, the positioning of the electrodes was verified by Dr Broberg.

6.1.3 Pharmacological agents

The drugs in the various studies have been chosen based on their ability to target various aspects of theories of epileptogenesis.

There are two main effects which a neuron in the brain can convey to other neurons: excitatory and inhibitory. It is supposed that epileptic seizures are essentially due to one or both of over-excitation and under-inhibition, which cause increased brain activity. This appears to make the brain vulnerable to some form of resonance.

6.1.3.1 Picrotoxin

It is thought an increase in activity within the brain can lead to epileptic seizures. Picrotoxin produces seizures by blocking the chloride channel of the GABAA receptor [45] thereby preventing chloride-induced hyper-polarisation and thus blocking inhibition, leading to increased brain activity. The resulting seizure has no apparent location responsible for the initiation. Although this is thought to mimic the neuronal activity during a primary generalised seizure, it isn't known exactly how this increase in brain activity results in seizures. As mentioned, epilepsy has an element of heritability, and a gene identified as "epilepsy causing" has been found to cause chloride channel dysfunction [87], which lends credence to the picrotoxin model. It is known that GABAA channels are to be found throughout the brain [56], explaining picrotoxin's generalised effect.

6.1.3.2 Kainic Acid

Kainic acid is an excitatory amino acid agonist to the AMPA and kainate receptor [85], and is well known to produce repetitive seizures (thought to be a result of excess excitation). Experiments have shown that kainic acid initially affects limbic structures or the cerebral cortex, and for this reason has been used as a model of temporal lobe and generalised seizures [49]. The administration of 10 mg/kg kainic acid has been found to produce convulsive seizures without toxic effects [49, 44].

6.1.3.3 Fluorocitrate

A small dose of fluorocitrate is selectively taken up by astrocytes in preference to neurons, is reversible (there is no tissue damage), and is known to cause seizures [86]. Fluorocitrate blocks aconitase – an enzyme which converts citrate to isocitrate in the tricarboxylic acid cycle. This greatly reduces a cell's ability to metabolise energy. A significant dose of fluorocitrate will destroy tissue since

the cells will die from lack of ATP, however small doses reversibly interfere with cell function. The fact that it is selectively taken up by astrocytes means that it can be used to examine the role of astrocytes in brain function. Fluorocitrate is injected locally into the cortex, and its influence remains local. Therefore, this can be used to examine the effect of local astrocytic disability and explore the relationship to epilepsy and seizures.

6.1.3.4 Citrate

Uptake by astrocytes is not the only effect of fluorocitrate – it can bind with free calcium, potentially leading to other physiological effects. Citrate is an analogue of fluorocitrate, but is considered neurologically inactive, although citrate also binds to calcium. For this reason, citrate was used in an attempt to show that any observed effects of the fluorocitrate injection were not due to its binding with calcium.

6.1.4 Impedance recording

6.1.4.1 Cell-swelling theory

The impedance recording is a means by which we can estimate cell swelling, which can occur during abnormal brain operation [54]. The cell wall is composed of a lipid bilayer, which is highly electrically resistive. However, because it is thin, it can cause capacitative effects. The extracellular fluid (ECF), surrounding the cells, is approximately resistive. Because the brain is in a confined space, any cell swelling that occurs will result in a decreased volume of ECF, which will change the electrical properties of the tissue (figures 6.4, 6.2 and 6.3). Impedance is measured by applying an alternating current of constant amplitude (a *constant current*) to the brain (figure 6.1), and measuring the voltage. Depending on the degree of cell swelling, the tissue will show changes in impedance that are dependent on the frequency of the stimulus current.

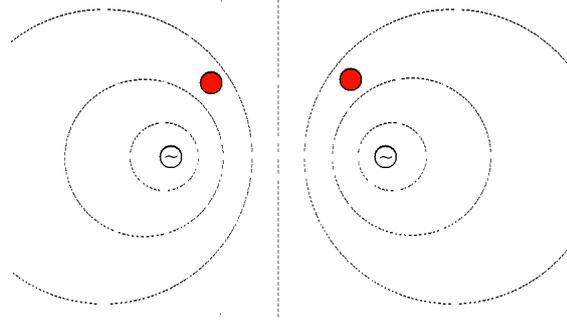


Figure 6.1: EEG electrodes and impedance drivers

The electrodes marked “~” deliver the constant current to the brain, causing an electric field to form (approximately as shown). The density of this electric field is dependent on the impedance (Z) of the tissue, because the driver voltage will increase if Z increases to maintain constant current levels. The red electrodes are EEG electrodes. The low frequency component signal (0 - 1000 Hz) recorded at the EEG electrodes is the EEG, and the high frequency signal component (above 5 kHz) is the voltage due to the constant current driver from which we can, with calibration, calculate the impedance of the tissue.

6.1.4.2 Impedance hardware

The impedance system consists of two driver electrodes which apply a constant AC current of 0.1 mA peak-to-peak at 50 kHz to the brain². This signal is passed through a reference resistor and also through the stimulating electrodes, injecting the current into the brain. The impedance hardware measures the voltage across a reference resistor, as well as the voltage across the EEG electrodes (refer to figure 6.1), and from this measures the impedance – an analysis performed in real time.

²For the purpose of deciding the optimum frequency at which to stimulate the brain, a range of frequencies between 2 kHz and 700 kHz were examined, with 50 kHz yielding the best estimation of brain impedance [54].

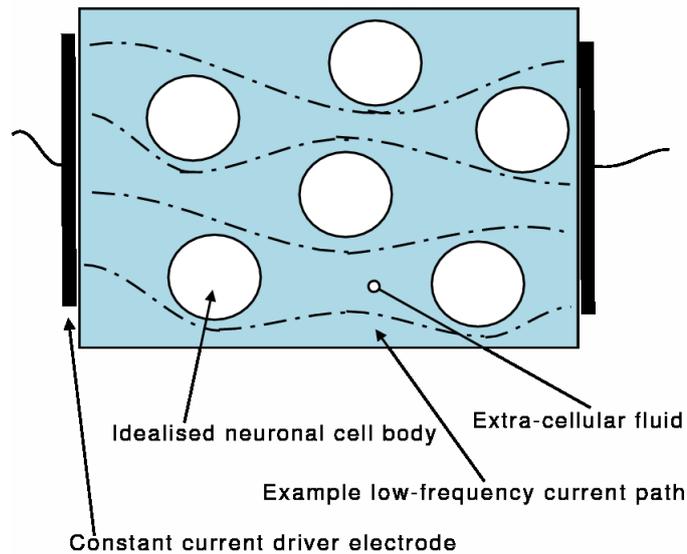


Figure 6.2: The path of low-frequency current through neural tissue
 This figure illustrates the path that a low-frequency current (eg. 50 kHz) takes through tissue. The blue area represents the ECF, the circles represent the neurons' cell bodies, and the dotted lines represent the path that the current takes through the tissue. The current cannot pass through the cells because the lipid-bilayer membrane has a capacitative effect, making the through-cell path a high impedance path.
 The elements in this diagram (and figures 6.2 and 6.3) are not to scale. Typically, the extracellular space in the brain is approximately 20% of the tissue volume.

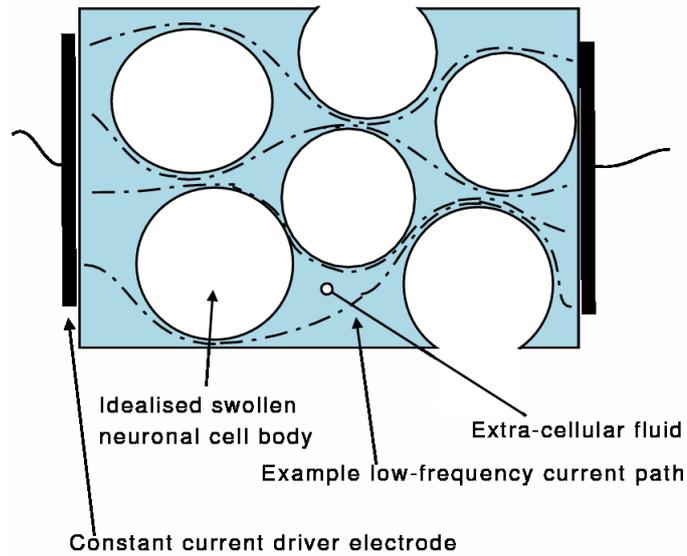


Figure 6.3: Swollen cells, and their effect on low-frequency current passage
 Because the swollen cells occupy a greater proportion of the tissue (with a corresponding reduction in the amount of ECF) the current must take a more convoluted path through the cells, travelling further. The current is also condensed into the (now smaller) spaces between the cells. Both of these result in an increase in impedance.

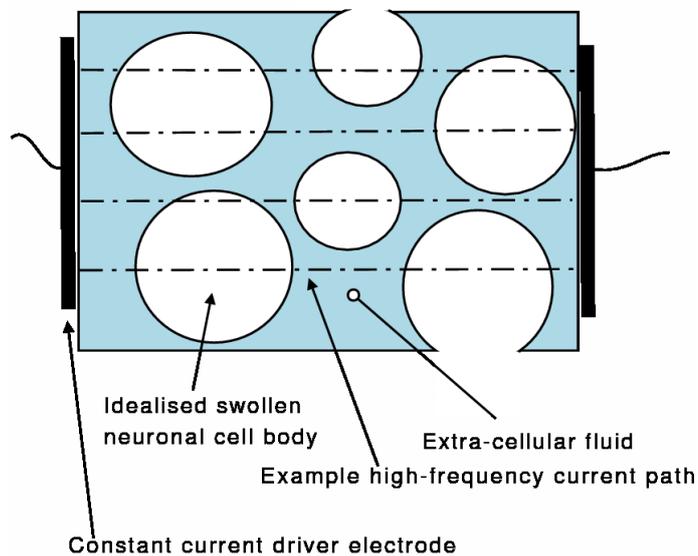


Figure 6.4: Path of high-frequency current
 Notice that a high-frequency current (eg. 500 kHz) can pass through the cells unhindered. This shows that the passage of the high-frequency current component is unaffected by cell swelling. Notice, in particular, that the route taken by the high-frequency current is the most direct.

The hardware that makes the impedance recording was developed primarily by Professor Torsten Olsson in Gothenburg, Sweden. He came to our lab in March 2004, and I worked with him to develop software to calculate the impedance from the data in real time. The output from the device consists of two voltage measurements – one within the device (due to the constant current applied across a known impedance), and one measured between the EEG electrodes. Because the frequency range of the EEG and the impedance stimulus are so different, we are able to measure both EEG and Z at the same electrodes. The software consisted of a C library to examine the amplitude and phase variations of the impedance signal, and calculate the impedance.

The experimental setup was quite complicated, because the Swedish impedance hardware was not designed to integrate with our experimental hardware. Figure 6.5 illustrates the arrangement of the hardware. The Z stimulus control computer manages the stimulus hardware, and regulates the constant current stimulation. The stimulus hardware also records the 50 kHz signal at the EEG channels, and calculates the impedance. The hardware uses a 12-bit ADC, but in order to obtain a better quantisation resolution a shifting scale was implemented. This meant that the 12-bit ADC was used as though it were the lower 12-bits of a 16-bit ADC, yielding better resolution but causing the data to wrap (figure 6.7) when traversing high-bit boundaries. An unwrapped, low-resolution recording of the impedance was stored on the stimulus control computer. The high resolution, wrapped impedance data are converted to an analogue signal and recorded in the recording hardware alongside the EEG, to provide a high resolution synchronised recording of EEG and impedance.

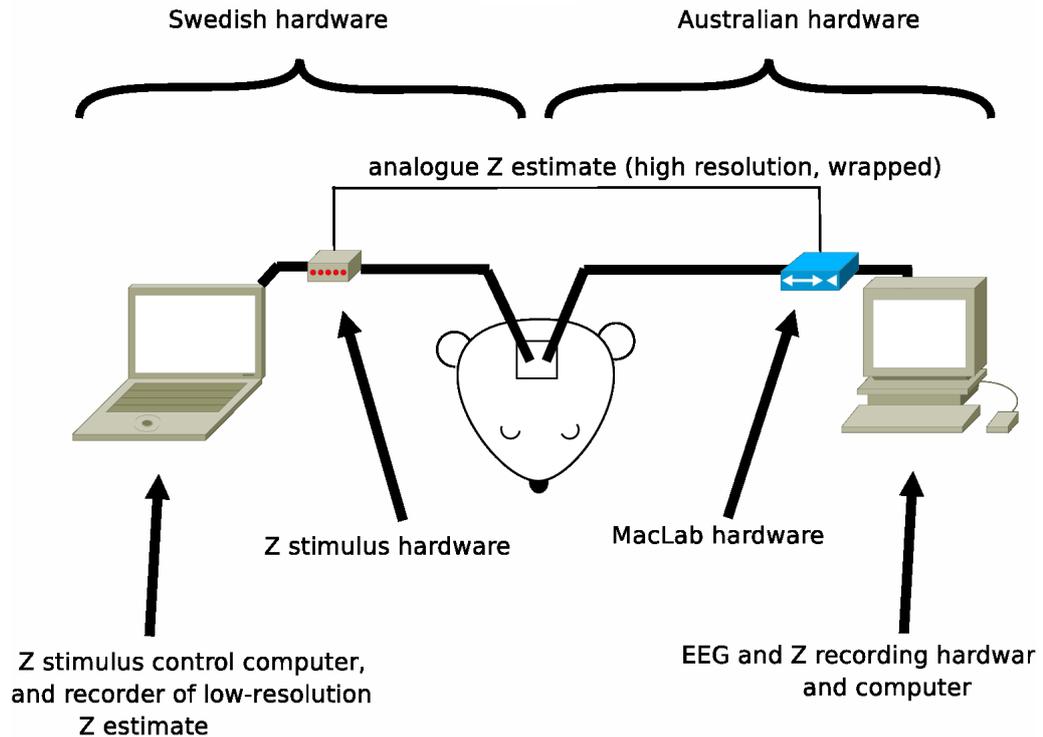


Figure 6.5: Impedance recording experimental setup

The Z stimulus computer controls the behaviour of the Z stimulus hardware – regulating the amount of applied current, and the applied frequencies. It also calculates the impedance from the feedback received from the rat by the stimulus hardware. This is stored internally in a low-resolution format, using a Microsoft Windows™ memory-mapped file. The stimulus takes the calculated impedance, and converts it to an analogue signal, for synchronised recording alongside the EEG in the EEG recording hardware. This high-resolution analogue signal contains wraps, due to limitations in the Z stimulus hardware.

6.1.5 Recorded signals

There are three signals that are measured from the rat:

1. EEG: two channels of EEG are generally recorded.
2. Brain impedance, Z (section 5.3.3)
3. Ionic concentration of Potassium in the extracellular fluid.

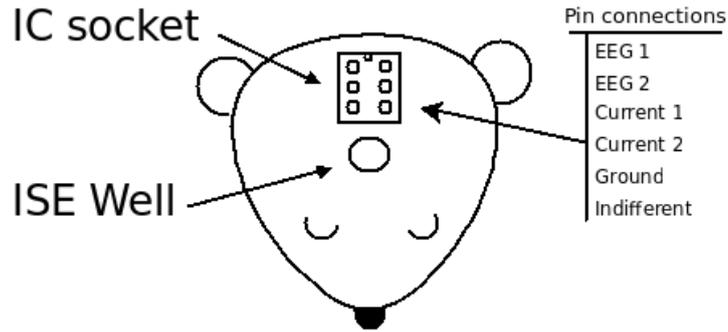


Figure 6.6: The rat head, experimental setup

This image shows a typical setup of the equipment, and the means by which the signals are controlled and recorded. There are six screw electrodes that record or inject electric potentials, and these are accessed via the IC socket. Two of these channels are connected to a constant current oscillator, and apply this stimulus to the dura. The other two record the EEG and, because the signals occupy such different frequency ranges, can also be used for the measurement of the impedance. The reference electrode is attached at the frontal sinus, and the ground is attached to the bone posteriorly – over the cerebellum.

There are six screw electrodes that are connected to the IC socket (figure 6.6). Two of these are used for injecting the constant-current that allows the measurement of the impedance. Two are responsible for recording two EEG channels and an impedance channel (the impedance data are in the high-frequency component, and the EEG data are in the low-frequency component). There is also a ground, attached anteriorly and a reference electrode attached to the skull at the frontal sinus. The constant current oscillator produces signals at 50 kHz, and by recording the voltage that results at these frequencies we are able to measure the impedance of the brain. The EEG is low-pass filtered at 2 kHz, so that none of the constant current driver signal is present in the recorded EEG.

These are excellent data, recorded from a standard animal model, in a very low noise environment. Because the animals are paralysed, there is no EMG artifact. The recording of impedance as well as EEG (at a very high sampling rate) provides a perspective that is granted to very few researchers.

6.1.6 Post processing of impedance data

There were several limitations to the recording of the impedance data, largely because the Swedish hardware was not designed with consideration for the Australian hardware (figure 6.5). Our experiment used custom hardware developed at the Flinders Medical Centre, and this was combined with custom hardware developed in Sweden.

The Swedish impedance hardware digitised the signal at 50 Hz, but was not synchronised with any other channels, for example the EEG. It was stored in a Windows memory mapped file, which made analysis on our Linux systems difficult. The Swedish hardware had been modified to allow an analogue output of impedance in real time, which our hardware recorded alongside the EEG (sampled at 4 kHz) to ensure synchronisation, however the impedance was wrapped. Thus, we had two representations of the impedance signal. One was high-resolution and synchronised with the EEG, but had wraps. The other was low-resolution, and un-synchronised, but had no wraps. We needed to generate high-resolution, synchronised, unwrapped data.

The wrapping in the impedance data was a problem because it made it impossible to visualise long term trends in the impedance. This problem was exacerbated because the response time of the ADC is such that approximately 17 samples are measured while the wrap occurs. If we closely examine a wrap, we see that it is not a true discontinuity (figure 6.7). The data recorded on the Australian equipment were split across multiple files to allow portability and analysis on older hardware.

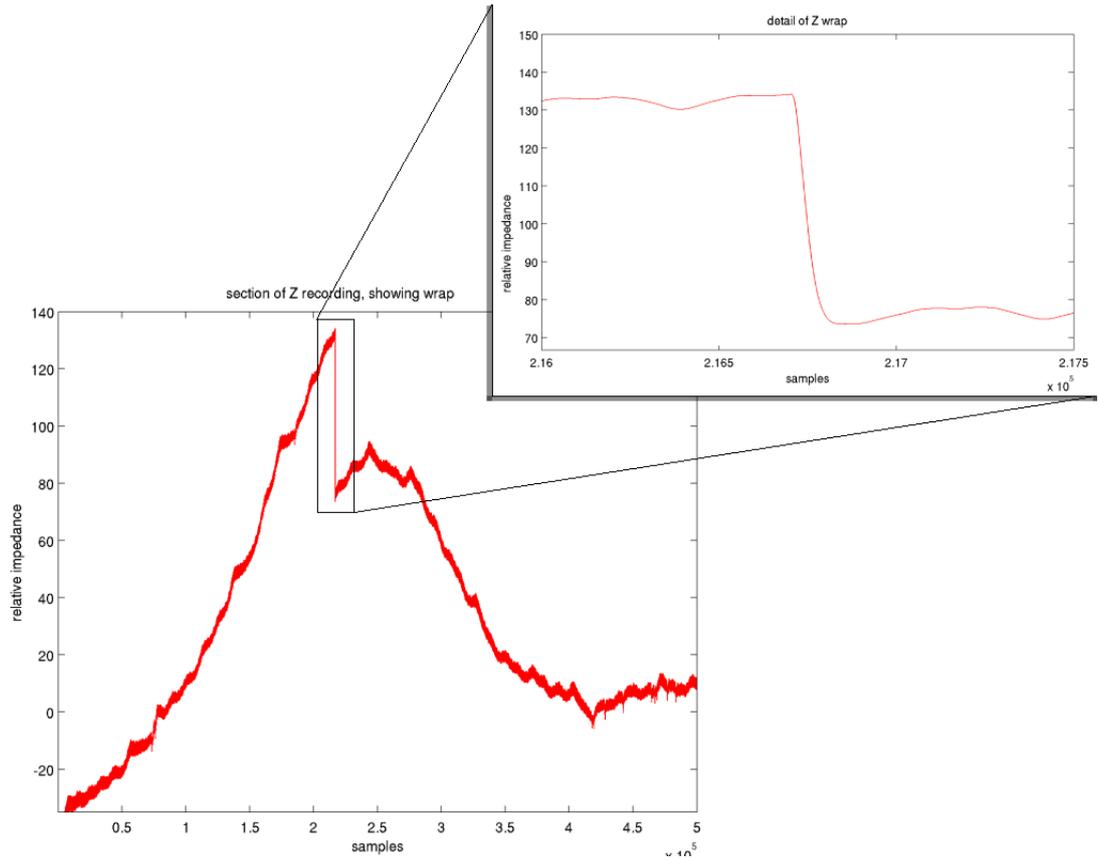
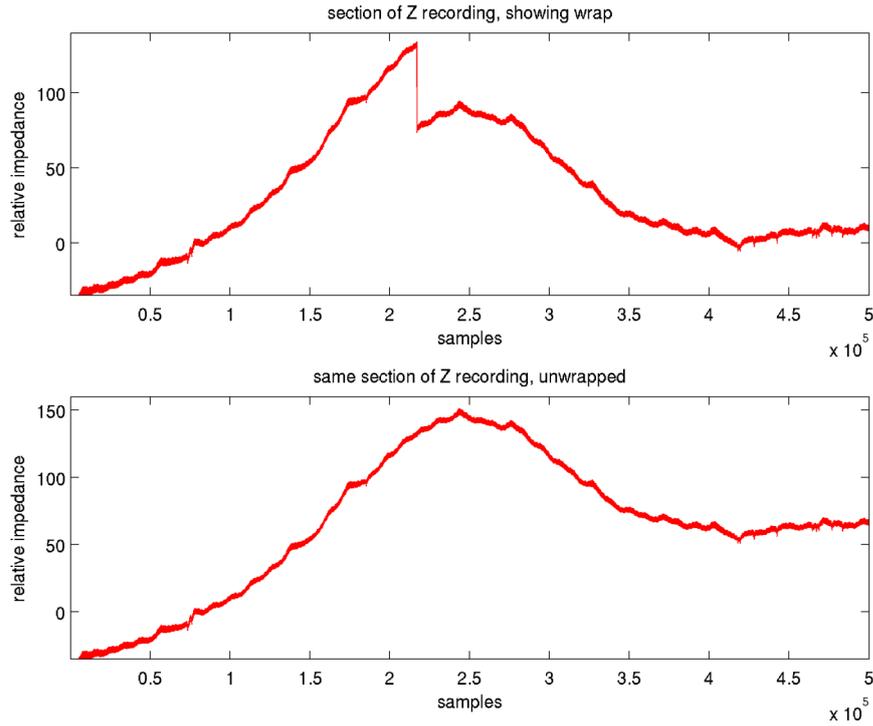


Figure 6.7: An example of wrapped data – sampled at 2kHz
An example of wrapped data. The main plot has discontinuities. The subset plot details a close view of a discontinuity, illustrating that it is not a true discontinuity, but rather occurs over several samples.

To facilitate analysis, I wrote software to read the memory mapped file (C++ and Matlab). It read in the multiple text files containing the high-resolution, wrapped impedance and EEG data. It unwrapped the impedance data, as shown (figure 6.8).

These data were then aligned with the impedance data obtained from the memory mapped file. This needed to be done because there were discontinuities in the recording that were actual, and ought not to be removed (the addition of saline to the well caused an abrupt impedance change, due to the additional conduction pathway). Figure 6.9 shows two data sets. The upper set is the wrapped data recorded at 2 kHz. The lower set are the data recorded at 50 Hz and stored in the memory mapped file. Notice that in the lower set there are still discontinuities at about the 500th and 4000th samples – these are



*Figure 6.8: Wrapped data and unwrapped data – sampled at 2kHz
This figure illustrates the input data that contain wrappings and the output data that have been repaired.*

due to the application of saline to the well in the rat’s skull. In this situation, the program would determine that the discontinuity should not be removed (because it was also present in the memory mapped file), and would “undo” its removal. Once this process was completed, we had contiguous, synchronised files ready for analysis.

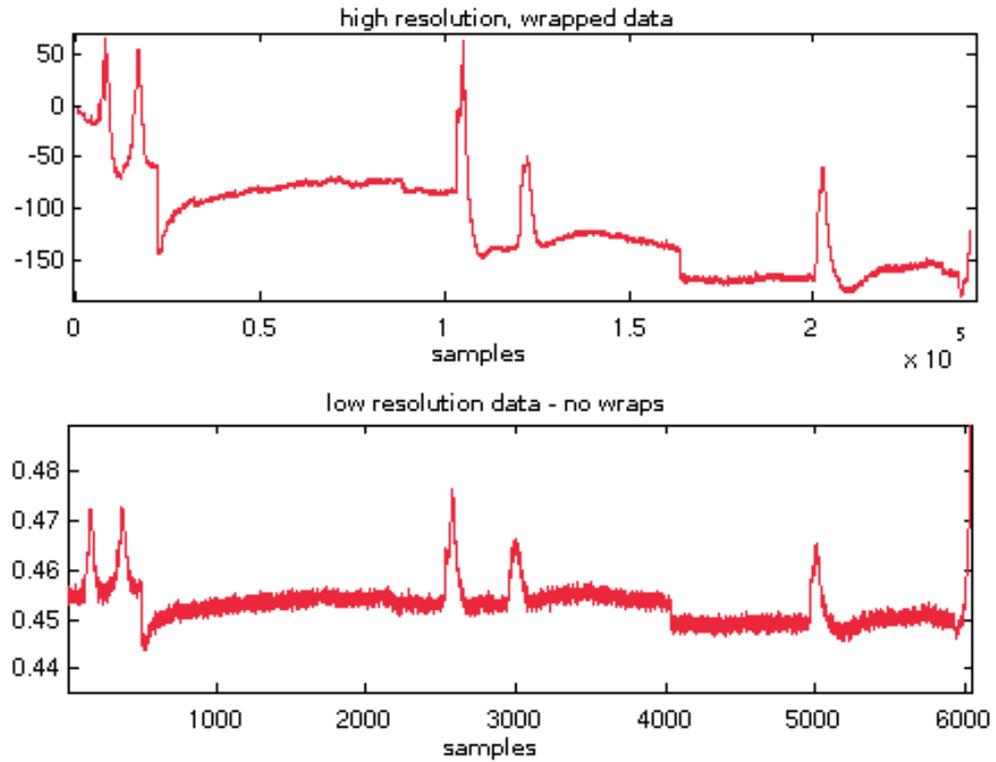


Figure 6.9: Real discontinuities

There are some discontinuities that are not a measurement artifact caused by the dynamic shifting of the ADC's range, but rather are real and are caused by (for example) the application of saline to the exposed dura. Such discontinuities should not be removed because they are real.

6.2 Human data acquisition

In 2003, we ran a pilot study that showed elevated gamma in inter-ictal EEG in subjects with PGE relative to controls [84]. The larger experiment described here represents a follow-up to those very exciting and promising results.

6.2.1 Preparation

Volunteers were recruited from within the Flinders Medical Centre (Adelaide, South Australia), after being recommended by their referring specialist, controls were also recruited by advertisements. These referred volunteers exhibited various neurological disorders, and were matched to control subjects (by age,

handedness, etc) – for the experiments in this thesis, we were only interested in primary generalised epilepsy and control groups. Each subject was booked in to the lab to conduct an experiment.

The subject’s hair was washed, if necessary, without conditioner, so that hair oils were reduced – these can affect the behaviour of the electrode-scalp interface. The appropriately-sized cap (EASYCAP GmbH, Germany) was fitted to the subject by our EEG technician, and conductive gel was injected into the electrode-scalp cavity. The subject was then moved into the Faraday cage, where they remained for the duration of the experiment. They were seated and connected to the Neuroscan EEG amplifiers which amplified and digitised the signal, allowing it to be recorded to computer at 16-bits, 2000 samples per second.

As with the rat experiments, the data are of excellent quality. The EEG cap contains a dense array of 128 electrodes, and signals are digitised at a high sampling rate. Because the experiment is conducted in a Faraday cage, noise levels are very low.

The impedance of the interface at each electrode was measured and verified. The electrodes and gel were adjusted as necessary to reduce the impedance at each electrode to less than $5k\Omega$ – this sometimes involved the light scraping of dead skin cells from the surface of the scalp. For the duration of the experiment, the subject was seated in front of a computer monitor and custom keyboard, built for the project by the Biomedical Engineering Department of the Flinders Medical Centre.

The experiment was automated, using custom experiment design implemented in the Presentation (Neurobehavioral Systems, California) software package. It was written so that every subject was presented with an identical (or, as close as possible) experience. A pre-recorded automated voice presented the information as well as any required feedback. If needed, the experiment manager could communicate with the subject via a microphone outside the

cage. This was sometimes needed if the subject did not understand the instructions, or had difficulty with the task.

6.2.2 Experimental procedure

Whilst in the cage, we recorded the subjects' EEG. During this process, they performed a series of mental tasks. These tasks were designed to selectively exercise different brain regions (corresponding to different forms of mental processing). We did this because we are interested in gamma, and expected that there was a correlation with mental activity. The tasks were

- Eyes closed: This is a basal state of operation of the brain, and can be viewed as background EEG.
- Eyes open: The subject opens their eyes, and stares at the computer screen, while doing nothing else. This can also be considered as a basal state, but including the contribution from the visual cortex.
- Finger tapping left: The subject taps their left index finger as rapidly as possible.
- Finger tapping right: The subject taps their right index finger as rapidly as possible.
- Mental rotation of computer generated shapes: The subject is presented with two geometric shapes and needs to decide if they are rotated versions of each other, or if one is flipped (a mirror image) relative to the other.
- Visual Discrimination: They are presented with an original shape, and four other shapes - one of which matches the original. The subject must decide which image is the same.
- Auditory Discrimination: Identification of subtle word differences. A series of similar sounding pairs of words are presented to the subject, who needs to decide whether pairs of words are identical.

- Reading: The subject silently reads a passage of text and, once finished, answers some questions.
- AVLT (auditory verbal learning test): This task tests a subject's ability to recall words from a previously presented list.
- Maze: the subject must solve a maze presented to them. An un-revealed maze is presented on the screen, and the subject must navigate by trial and error. Correct moves result in progress through the maze, while incorrect moves result in no progress, and audio reinforcement. To complete the task, the subject must traverse the maze with no incorrect moves twice in succession (the task ends after 8 minutes).
- Subtraction (serial sevens): The subject repeatedly subtracts seven from a large number, quietly in their head. In this task, it is particularly emphasised that the subject should not move or vocalise as they perform the mental operations.

The EEG was recorded for the duration of the experiment, and was labelled to indicate the start and finish of the various parts of the experiment. For the purpose of analysis, tasks that had a correct or incorrect response from the subject (such as the visual discrimination task, where they may identify the wrong shape) had the EEG divided into two sets – a correct set and an incorrect set.

6.2.3 Data processing

The data were saved, catalogued, and the details of the subject were entered into a mySQL database, to allow simple indexing of groups of subjects. The recording was examined, and areas of EEG in which excessive EMG (muscle artifact) was present were marked as unuseful.

6.2.4 Experiment design

Ideally, it would be possible to record subjects' EEG at the time of seizure, and then examine EEG immediately prior to the seizure, as well as post-seizure EEG. This might, for example, allow the detection of a pre- or post-ictal change of state. To achieve this, we would need to admit a subject to the laboratory and ask them to remain there until a seizure occurred – however long that may be. This would be very inconvenient for both the subjects and the researchers, and would likely result in vastly decreased subject availability, and therefore a much smaller data set.

Conversely, the experiment as formulated requires the subject to remain at the laboratory for a predetermined amount of time, and perform a predefined set of tasks. The fact that subjects are in the lab for a similar duration, and are performing similar tasks, means that the experiment is more consistent than if they were waiting an indefinite (and variable) period for a seizure to occur. Also, because this approach will yield a generic sample of inter-ictal EEG, any results found therefrom will be more generally applicable.

6.3 Paralysed human data collection

This experiment was very similar to that described in section 6.2, and that was the base for its design. Because of the ethical issues surrounding the unnecessary paralysis of a healthy human, all subjects for this experiment needed to be knowledgeable in the field, and have an understanding of the safety implications of the procedure. Each individual experiment needed a separate ethical approval process, and the committee considered the knowledge of the subject and their relationship to the experimenters (among other things) each time.

6.3.1 Preparation

To date, only three subjects have participated in this experiment – all were knowledgeable in medical fields.

The subjects were prepared for EEG recording as in section 6.2.1. They were seated in a reclined, fully supported position and connected to the Neuroscan EEG amplifiers which digitised the signal and allowed it to be recorded to computer at 16-bits, 5000 samples per second.

These data are of excellent quality. The EEG cap contains a dense array of electrodes, and signals are digitised at a high sampling rate. Paralysed subjects means that no EMG is present in the EEG. Because the experiment is conducted in a Faraday cage, noise levels are very low.

For the duration of the experiment, ECG, EMG, expired CO_2 , blood pressure and pulse oximetry were monitored to ensure subject comfort and safety.

While paralysed, the subject required ventilation. They had, prior to the day of the experiment, practised with the laryngo-pharyngeal airway to increase familiarity with the procedure. All ventilation and monitoring equipment was outside the cage, with the inter-connectors routed through a small cage port. For the duration of the paralysed experiment, a neurologist and an anaesthetist remained inside the cage with the subject, and a second anaesthetist remained outside the cage, using the monitoring and ventilation equipment to monitor the subject.

Once connected and comfortable, the subject completed the mental tasks as described in section 6.3.2 in an unparalysed state. Once the tasks had been completed in an unparalysed state, the subject was prepared for paralysis.

The subject was injected with glycopyrrolate (0.4 mg i.v.) to reduce oral secretions, as it had been found that the insertion of the laryngeal mask airway caused a powerful salivation response. Pharyngeal lignocaine was applied as a local anaesthetic prior to insertion of a laryngo-pharyngeal airway. A manual sphygmomanometer cuff was applied to their non-dominant arm and inflated

to a pressure 1.5 times that of blood pressure, preventing the paralytant from affecting the arm and therefore allowing communication, while paralysed, using the custom keyboard. The subject was injected with 20mg of cisatracurium to cause muscular paralysis. To assess muscular paralysis, extensor digitorum brevis was electrically stimulated, at a level established prior to paralysis, at the right common peroneal nerve. Complete paralysis occurred approximately 5 minutes after the injection of the paralytant. For the duration of the experiment, the subject's seated position had the complete support necessary during paralysis, so their position did not change when they were paralysed.

It is important to note that, because of the presence of the sphygmo-manometer, the experiment had to be completed quickly. Anoxic paralysis (due to lack of blood supply) of the communicating arm began at about 18 minutes, at which time the cuff was removed. This allowed a temporary recovery of movement in the arm, before the paralytant acted, and the arm became paralysed again.

6.3.2 Experimental procedure

The tasks in this experiment were similar to those in section 6.2.2, though with several differences, because of the limitations imposed by the paralysis. As mentioned, these tasks were first performed in an unparalysed state, and then repeated with the subject paralysed.

- Eyes closed: This is the basal state of operation of the brain, and can be viewed as background EEG.
- Left eye open: the contribution from the visual cortex to the EEG. The eye was held open. This is another basal state of the brain.
- Photic stimulation: a strobe at 16Hz was used so stimulate the eye. This was performed with both eyes closed, and then with the left eye held open.

- Finger tapping left and right: This task is, of course, affected by the paralysis
- AVLT (auditory verbal learning test): This task tests a subject's ability to recall words from a previously presented list.
- Odd-ball: 180 discreet tones are presented to the subject, sounding for 50ms and spaced at 1 second. There are two tones, 500Hz and 1kHz, and they are presented at random, with the high tones comprising 25% of the tones. This is known as an "odd-ball" experiment, because the subject has to seek that which does not belong - the "odd one out".
- Subtraction (serial sevens): The subject mentally subtracts multiples of seven from a large number.

6.3.3 Data processing

The data were saved, and catalogued, to allow simple indexing of groups of subjects. The recording was examined, and areas of pre-paralysis EEG in which excessive EMG (muscle artifact) was present were removed.

Chapter 7

Developed algorithms

This chapter provides detail regarding the custom algorithms that were developed or modified for this project.

7.1 Entropy

Entropy is estimated via

$$H = - \sum_{i=1}^n Pr(x_i) \log_2 Pr(x_i) \quad (7.1)$$

where x is the binned EEG channel.

This function was based upon a function written by Dr Kenneth Pope and uses a ranking and bin algorithm to estimate the amount of information in a signal. For this project, I rewrote the code in vectorised form to increase performance. A window of 1000 samples was typically used.

Prior to entropy estimation, data were scaled to zero mean and unit variance. This removes the variance-sensitive portion of the entropy estimation, so that the measure describes the shape of the probability mass function (pmf). We also included a correction factor to reduce the bias arising from the analysis of finite sampled data.

7.2 Mutual Information

This analysis was based upon a function written by Dr David Mewett as part of his doctoral studies [51].

There is evidence that generalised epilepsy is associated with an increase in synchronicity of EEG activity across the brain [53]. Mutual information (MI) measures similarity between signals, and is able to reflect increasing sameness between EEG channels. It is calculated thus

$$MI = \sum_i \sum_j Pr(x_i, y_j) \log_2 \left(\frac{Pr(x_i)Pr(y_j)}{Pr(x_i, y_j)} \right) \quad (7.2)$$

where x and y are the binned EEG channels. All histogramming used 100 bins along each dimension.

After the two data sets are ranked (to ensure consistency between different data sets), their two-dimensional histogram is evaluated, based on the ranked data. These data represent the joint probability mass function (JPMF). The JPMF is used in conjunction with the two individual PMFs to find the mutual information.

The computational expense of MI precluded analysis of all pairwise combinations of electrodes, and a subset were chosen. To examine local synchronicity, 46 pairs of adjacent electrodes distributed across the head were selected. To examine non-local synchronicity, 448 pairs of electrodes from different brain regions (figure 7.1) were selected. These pairs of electrodes were chosen prior to any analysis, and remained consistent for the analysis of all subjects.

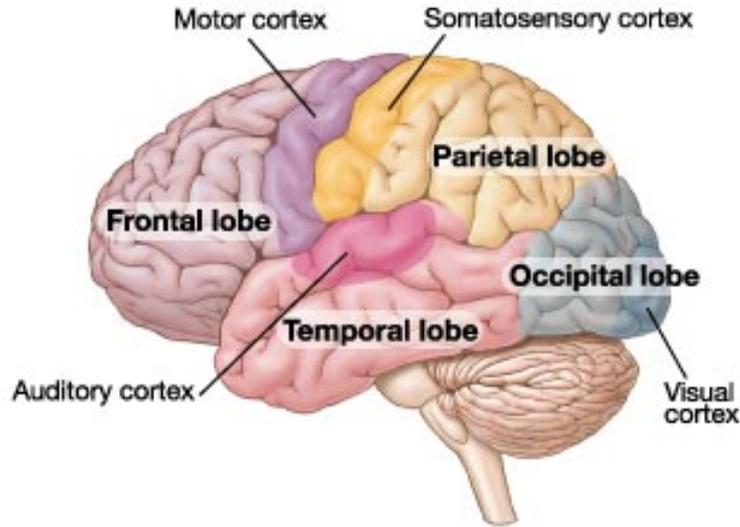


Figure 7.1: Schematic illustration of the brain, indicating the brain regions

http://wps.prenhall.com/ca_ph_wade_psych_1/0,7394,604018-,00.html

Brain regions represent the broad areas where different functional processing occurs. There are 4 main brain regions that are present in both hemispheres: frontal, temporal, parietal and occipital. These regions have been defined by grouping similar cognitive functions.

7.3 Correlation Dimension

Correlation dimension is an estimate of the number of degrees of freedom of a signal [34]. This is one of the simpler dimension estimates, and this function was written by me.

Correlation dimension is estimated from the correlation integral, which measures the number of pairs of points separated by less than a distance ϵ ,

$$C_x(\epsilon) = \Pr(D_x < \epsilon),$$

where D_x is the distance between randomly chosen points of the embedded EEG signal. Correlation dimension is estimated from the slope of the plot of correlation integral,

$$CD = \lim_{\epsilon \rightarrow 0} \frac{\log(C(\epsilon))}{\log(\epsilon)} \quad (7.3)$$

The function embedded the data at various dimensions to find the optimum

embedding dimension. Vectors¹ were subsampled so that approximately the same distribution of vectors was retained, while reducing the memory requirements. A distance matrix² was calculated, sorted and scaled to the interval (0 1], and the distances outside a range of interest ($D_x > max_{dist}$) were discarded. The remainder were binned into a histogram with a logarithmic bin spacing. The cumulative sum of this histogram is the correlation function.

Estimating the correlation dimension then involved calculating the limit in equation 7.3.

7.4 Linear Discriminant

The theory of discriminant analysis was introduced in section 4.1. The algorithm written here was for the simple case of two classes, to find the degree of match between the test data and example data representative of each class.

The pairwise covariance matrix of each representative data set was found. The test data were demeaned, and the pseudo-inverse of the covariance matrix calculated. The degree of match between the test data and each group was then calculated,

$$match = testdata' * pinv(covariance\ matrix) * testdata.$$

Experimentation demonstrated that LDA was markedly inferior to neural networks (results not shown), so LDA was not investigated further.

7.5 Surrogate analysis

The concept of surrogate analysis and the reasons for its use are discussed in section 3.4.2. We wish to create surrogate data that accurately mimic the

¹Each vector is a point in phase-space, corresponding to the state of the system's parameters at a given point in time.

²A distance matrix is an upper-triangular matrix containing the distance between each pair of points.

EEG for low order statistics, but lack any higher order statistical similarities.

7.5.1 Testing

The choice of surrogate algorithm is important because we must ensure that the surrogate data accurately mimic the EEG for low order statistics. Because of this, we compared different surrogates with EEG using linear statistics. All of the tested surrogate algorithms destroy higher-order statistical structure, so the choice of surrogate was based on its mimicry of the low-order statistics. Hence, the surrogate data were examined with Fourier analysis, auto-correlation and amplitude distribution functions. The results from the surrogates were then compared against those from the EEG.

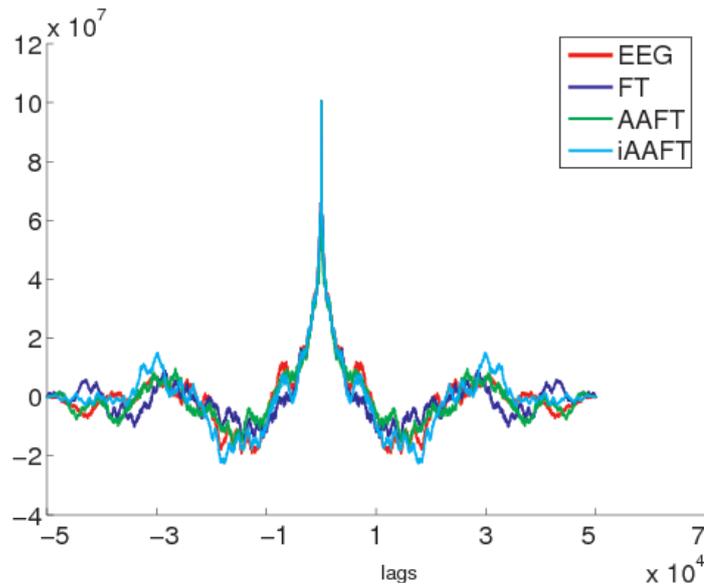


Figure 7.2: Autocorrelation function of EEG and three different surrogates

This figure shows a comparison between the autocorrelation function of three surrogates and that of EEG. Ideally, the autocorrelation of the surrogates would mimic that of the EEG. It is clear that, while imperfect, the closest match to the EEG is the iAAFT surrogate.

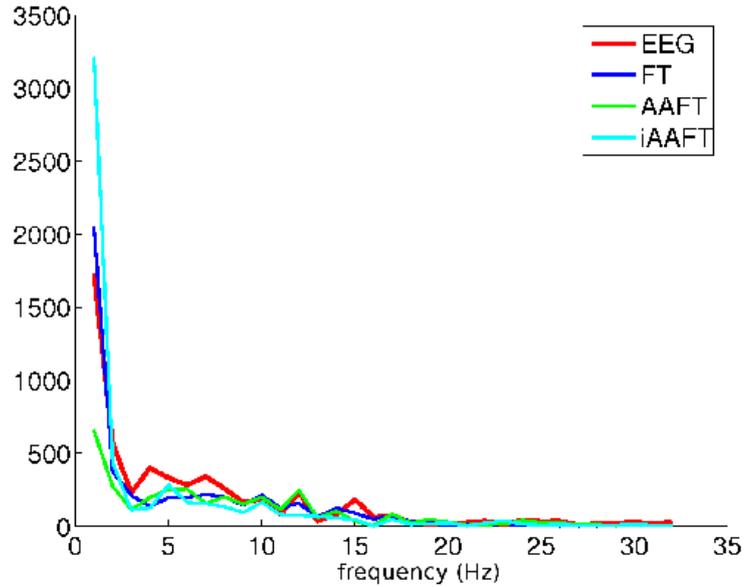


Figure 7.3: FFT of three different surrogates, and EEG

This figure shows a comparison between the FFT of three surrogates and that of EEG. Ideally, the FFT of the surrogates would mimic that of the EEG.

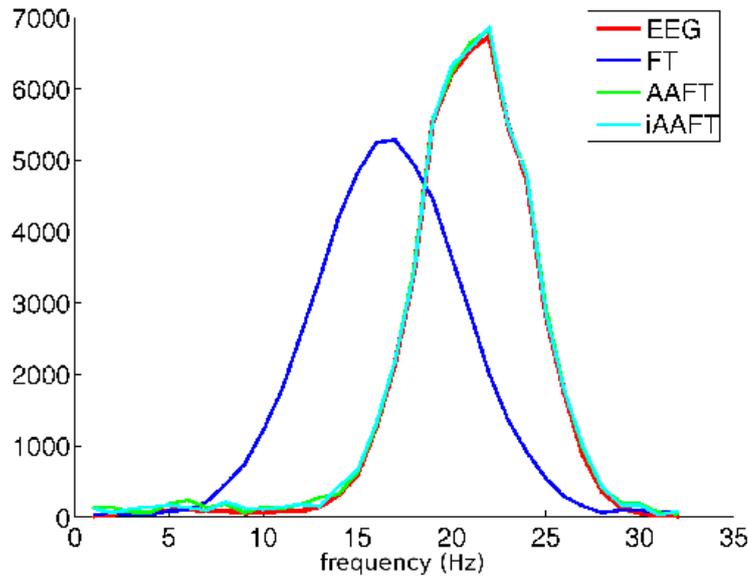


Figure 7.4: Amplitude histogram of surrogates vs EEG

This figure shows the amplitude histogram of the surrogates and EEG, highlighting the inadequacy of the FT surrogate.

As shown, each surrogate has its strengths and weaknesses. The FT surrogate is a close match of the autocorrelation, but poor for the FFT and the amplitude distribution. AAFT matches the amplitude spectrum closely, but is a relatively poor match for the FFT and the autocorrelation. The amplitude

spectrum resulting from the iAAFT matches that of the data closely, and its FFT is also a better match than that of the AAFT algorithm. These results are unsurprising when one considers the method of construction of the surrogates. Based upon these experiments, we chose to test the algorithms using the iAAFT method of surrogate generation.

7.5.2 Analysis

Rat EEG data (4 kHz, 12-bit, superdural screw electrodes) were used to synthesise 20 iAAFT surrogates (10 per channel). The data were a whole picrotoxin experiment, because it was considered representative of the type of nonlinear changes that we could expect to see in other rat experiments and hopefully human experiments also. The EEG and surrogates were analysed by the nonlinear quantifiers using a sliding window. The results were averaged, and compared with the results of analysis of EEG. This provided an idea of the sensitivity of the analyses to nonlinearities in EEG, and lack-of-sensitivity to linear statistical changes.

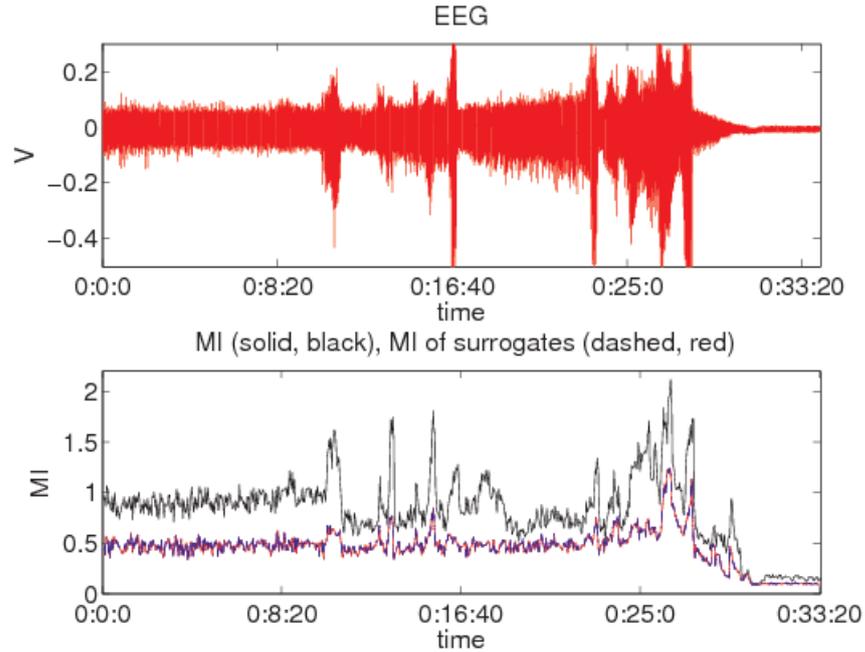


Figure 7.5: MI of surrogates (red; $\pm\sigma$ blue), vs MI of EEG (black).

This figure shows a comparison of a MI analysis of surrogate data vs analysis of EEG data. 10 pairs of surrogates were fabricated from the 2 channel EEG data using an IAAFT algorithm, the EEG and each pair of surrogates were separately analysed using the MI algorithm.

The upper axes shows EEG, and the lower axes show MI. The MI estimate between two EEG channels is shown in black. The mean MI between pairs of surrogates (one from each EEG channel) is shown in red, and $\pm\sigma$ shown in blue.

It is clear that, while there are obvious similarities between the MI of EEG and the MI of surrogate data, that the two are not the same. This is particularly striking because the MI of the surrogate data are very consistent.

The iAAFT surrogate applied was univariate, so we would not expect the cross-correlation to be preserved, contributing to the difference seen here.

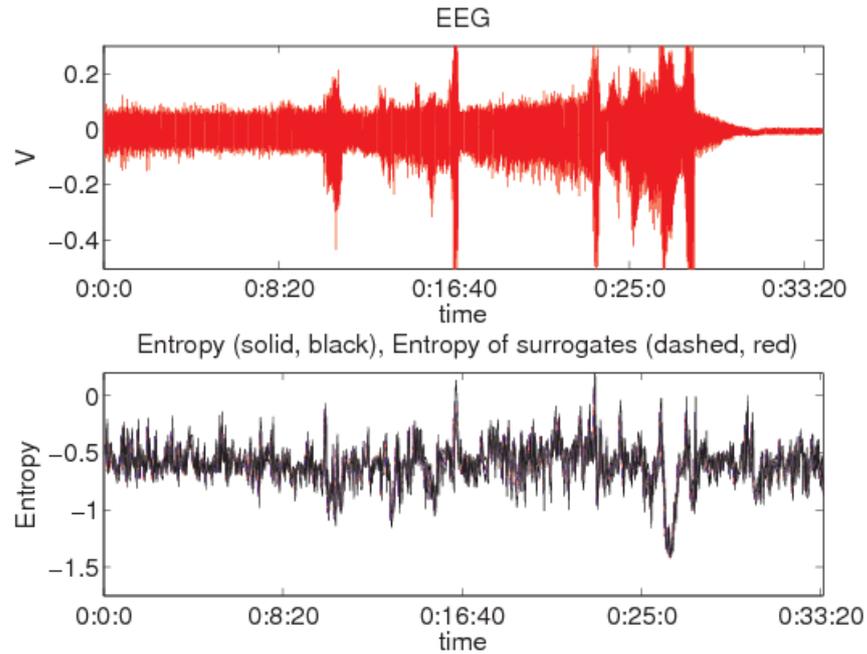


Figure 7.6: Entropy of EEG (red; $\pm\sigma$ blue) vs surrogate (black)

This figure is equivalent to figure 7.5, with the difference being that the tested algorithm is entropy. This figure shows entropy cannot distinguish between EEG and surrogate data. This could imply that the entropy analysis is insensitive to nonlinearities in the data, but it is also likely to be related to the choice of surrogate (because the *i*AFT surrogate preserves the amplitude distribution, which entropy measures).

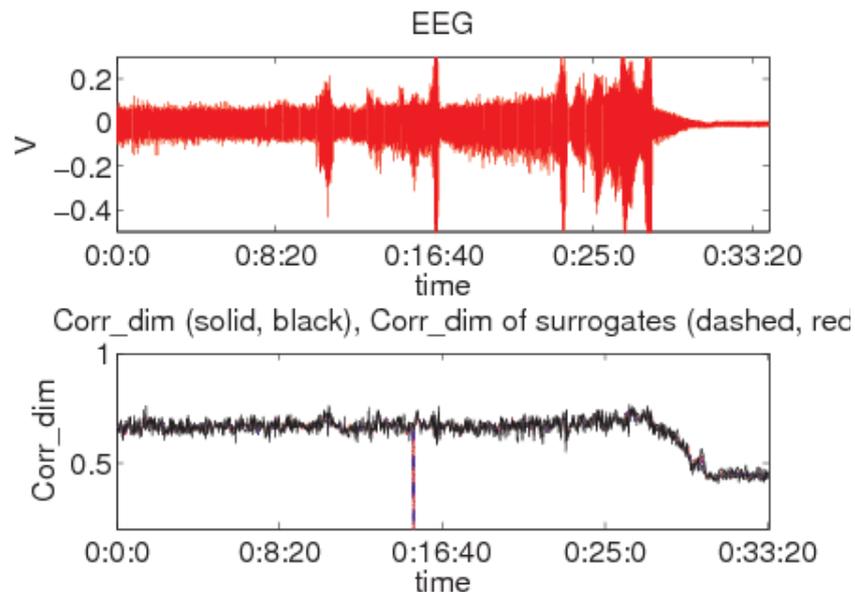


Figure 7.7: Correlation dimension of EEG (red; $\pm\sigma$ blue) vs surrogate (black)

This surprising result shows that the variations in the correlation dimension analysis are wholly described by low-order linear changes in the EEG, and hence that the correlation dimension algorithm is insensitive to higher-order statistical changes in the EEG.

Figures 7.5, 7.6 and 7.7 show the response of three analyses to surrogate data. The data are standard experimental data, recorded at 4 kHz. The surrogate data were produced using an iAAFT algorithm, and for each EEG channel 10 surrogates were produced – allowing MI to be tested on 10 pairs of surrogate data. Out of the three examined analyses, only MI show differences between the surrogate and EEG data.

We also examined graphs containing the analysis of EEG with the mean of the surrogate analysis subtracted. This is the *innovation* – the component in the EEG analysis results that were not accounted for by the analysis of surrogates. These are shown in figures 7.8, 7.9 and 7.10.

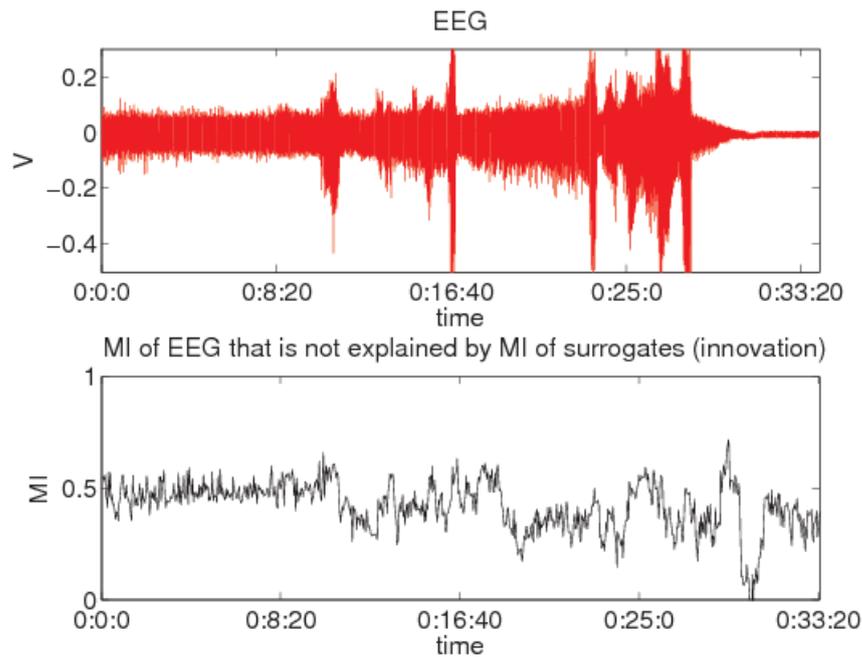


Figure 7.8: Innovation of MI of EEG with respect to MI of surrogates

This figure is produced from the same data as figure 7.5, except that the MI plot shows the estimated MI of the EEG with the mean MI of the surrogates subtracted. This can be called innovation: the portion of the MI of the EEG that is not accounted for by the MI of the surrogates – “the nonlinear bit”. There is a rapid decrease in MI-innovation after the first spindle and a slight increase in MI prior to seizure onset – there is no obvious EEG change associated with this change in MI.

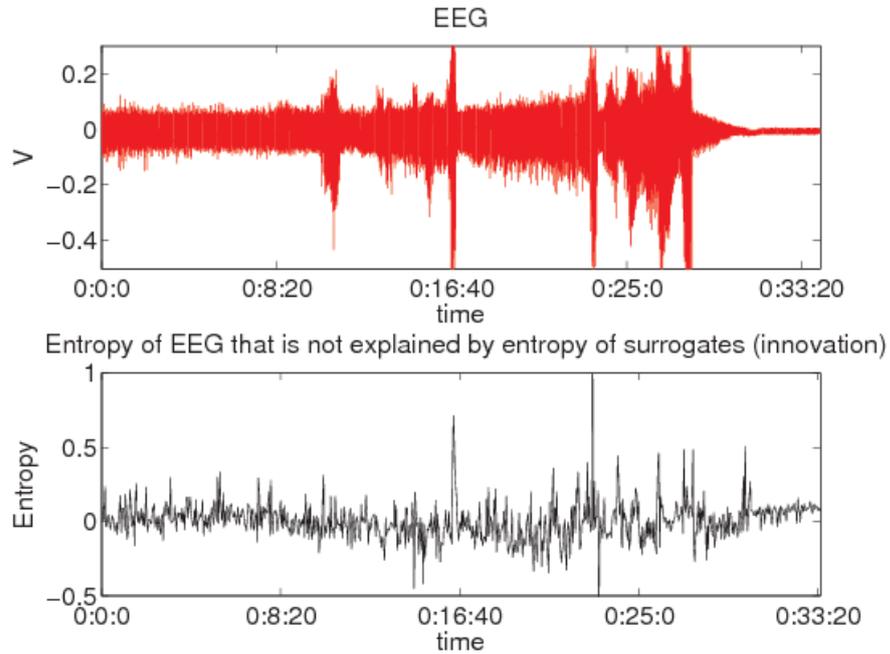


Figure 7.9: Innovation of entropy

Although the innovation component of the entropy analysis is much smaller than that of the MI (figure 7.8), there is a decrease (starting at 8:20) and a pre-ictal increase (starting at 18:00).

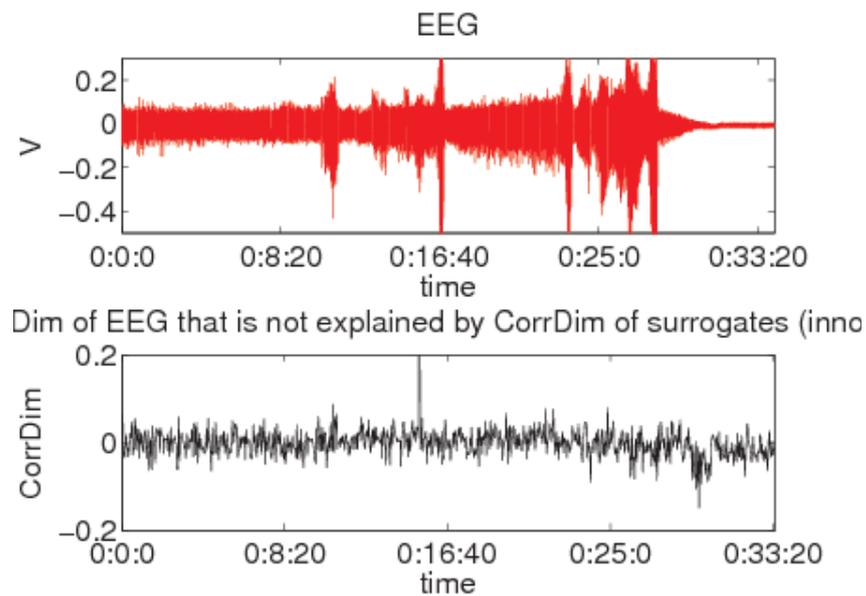


Figure 7.10: Innovation of correlation dimension

The correlation dimension analysis shows essentially no innovation in the EEG data analysis. This implies that the correlation dimension algorithm is not usefully detecting nonlinearities in the EEG data – nonlinearities that we know to exist based upon the MI and entropy analyses.

7.5.3 Discussion

MI innovation

MI exhibits innovation in the analysis of EEG relative to surrogate. While it appears that much of this can be related to events in the EEG (such as spindles and seizure) there are some patterns in the MI that are not. Perhaps the most prominent of these occurs near 18:00 – a small increase in the MI innovation. Note also that there is a decrease in MI innovation after the first spindle, and a further decrease after the seizure.

Entropy innovation

The entropy innovation (figure 7.9) is much smaller than that in the MI analysis. Its very small amplitude and (relatively) high noise level make it very difficult to draw many conclusions. Entropy calculation is less sensitive than MI to nonlinear changes in EEG. It is possible that there is an increase in the entropy innovation immediately prior to the seizure, however the noise in this signal precludes a firm conclusion – this would need to be examined over many experiments.

Correlation dimension innovation

The correlation dimension analysis is utterly insensitive to higher-order nonlinear statistical changes in the EEG data. Correlation dimension is a calculation that requires many stationary data to enable a robust estimate, and this has proven difficult to satisfy with the sampling rates that we have used in these experiments. It may be that the lack of nonlinear sensitivity and specificity is attributable to a lack of data.

Chapter 8

Experiment: Nonlinear analysis of rat epileptiform EEG

We recorded extra-dural EEG and impedance data from respiration, paralysed rats. Animals were injected with epileptogenic¹ drugs, via an implanted venous catheter, in doses designed to induce epileptiform activity.

The EEG was analysed using entropy, mutual information, correlation dimension and auto-correlation algorithms (the *analyses*). The experimental data and analyses were examined in an attempt to elucidate some of the underlying mechanisms occurring during the progression from normal EEG through to seizure.

Impedance generally increases as the experiment progresses. Ictal EEG is related to transient increases in impedance, increases in MI and decreases in entropy. The analyses reflected changes coincident with (and sometimes in advance of) EEG events. However, there was no evidence of a consistent predictive element, and little evidence of a pre-ictal intermediate state. There appears to be a relationship between the EEG (and analyses thereof) and the impedance data, but this is of undetermined consistency.

¹A drug that induces epilepsy-like symptoms.

8.1 Objective

This exploratory experiment was conducted in parallel with the development, implementation and refinement of the analyses. Although there was no hypothesis, per se, its purpose was to examine the extent to which the analyses consistently reflected changes in EEG, whether they could detect changes during inter-ictal EEG or spreading depression, and whether they had a consistent relationship with impedance data.

8.2 Data analysis

Data were gathered as described in section 6.1. The method of impedance recordings and their relationship to cell-swelling is described in section 5.3.3.

Data were processed by software written in Matlab (Mathworks). Discontinuities in the impedance data were repaired (section 6.1.6), the EEG was analysed using several nonlinear tools (chapter 7) and the results were stored, ready for visualisation (figure 8.1). Surrogate analysis 7.5, demonstrated that the implemented entropy and correlation dimension algorithms show little that was unexplained by low order statistics. That they are not exclusively sensitive to nonlinear dynamics does not make them useless, but it should be considered in the interpretation of the results. Entropy is a measure of the information content of the data, and its sensitivity to linear dynamics was expected (section 7.1). It is useful because it provides a means of assessing fluctuations in mutual information with information content, hence acting as a control. Similarly, correlation dimension provides a measure of degrees of freedom of the signal, which is useful to consider when interpreting MI variations. For these reasons, they were included in the analysis.

The experiments typically recorded 2-channel EEG, so this was used as a standard. The analysis took the form of a moving window, yielding changing analysis results as the experiment progressed. The window size was 1000

samples for all analyses except correlation dimension², for which it was 2000 samples. Windows were spaced³ at 1/4 their width (250 or 500 samples). Analysis results were stored coincident with the most recent EEG used in their calculation⁴.

²As mentioned (section 7.5), correlation dimension requires many data to produce a stable output.

³This refers to the separation between adjacent windows. That the distance between adjacent windows is only 1/4 their length implies that adjacent windows overlap.

⁴This ensures that events in the analysis reflect only contemporary and historical EEG, *not* future values.

8.2.1 Display

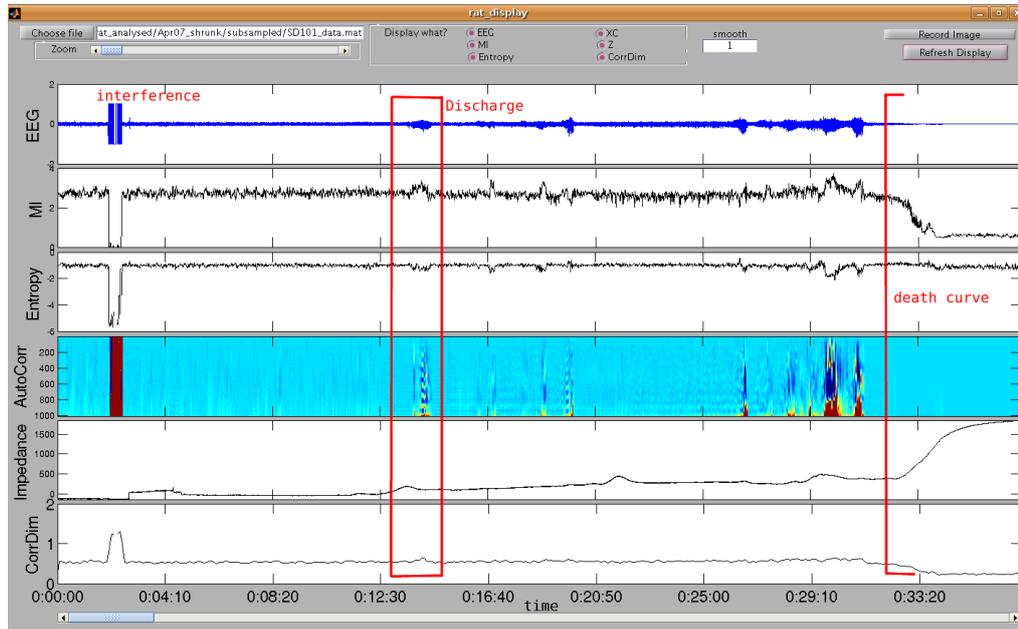


Figure 8.1: Rat display GUI, kainic acid experiment

This GUI was used to examine the results of the analyses, and shows typical behaviour of a rat experiment. The EEG channel (uppermost) shows the development of the seizures. The impedance channel (second from bottom) shows a steady increase until the “death curve” (starting at approximately 33 minutes). There are three small impedance transients at 00:13:00, 00:21:30 and 00:29:00. These are likely to be incidents of seizure and spreading depression (section 5.4.3). The death curve, at the end of the experiment, occurred immediately following the administration of a lethal dose of pentobarbitone, which also caused the decrease in EEG power (due to sedative effects). The strong interference at approximately 00:02:00 is the result of a connection adjustment to the rat’s head-piece (manipulating the rat for the connection of the ion sensitive electrode).

Notice the “smooth” tool at the top of the GUI. The ability to smooth the data was used only as a visualisation tool – none of the figures shown in this chapter have been smoothed.

A simple GUI was written to facilitate the display of data (figure 8.1). The impedance channel (second from bottom) shows characteristic behaviour, including the *death curve*⁵, illustrating a typical progression of an experiment in which an epileptogenic drug is administered.

Most of the channels are simple to interpret, however the auto-correlation channel is somewhat less so. We calculated the *auto-correlation function* for

⁵The death curve is the impedance change (cell swelling) resulting from the administration of a lethal dose of pentobarbitone.

each position of the sliding window, which was assembled into the image the same way as a spectrogram is assembled from a series of fast-Fourier-Transforms. Thus, each column represents the auto-correlation function of the EEG at a particular time, and the colours represent the strength of that function – blue (low), yellow (medium), red (high).

8.3 Results

First I will provide a brief overview of the experimental results. Figures are provided on the following pages, and are referred to in text.

Picrotoxin experiments (eg. figure 8.5) generally have fewer and shorter seizures but more dramatic impedance changes, compared to kainic acid experiments (eg. figures 8.2 and 8.3), which show continuous dramatic changes in EEG and repeated seizures. Kainic acid experiments often manifest pre-seizure EEG events that differ from spindles in that they are broad-spectrum, including high levels of gamma. Baseline impedance tends to show greater changes, but these are more gradual (partly due to the longer duration of the experiment) and are thought to be related to increases in gamma⁶ that occur throughout the experiment.

Fluorocitrate (eg. figures 8.6 and 8.7) does not produce seizures in the same manner as picrotoxin or kainic acid (by affecting neuronal excitability), but instead affects astrocytic metabolism, resulting in a very different effect. Few fluorocitrate experiments exhibit seizures, but instead show spreading depressions⁷. Citrate experiments show occasional impedance changes, but no spreading depressions attributable to the drug⁸.

⁶*Gamma* refers to EEG in the band 30 - 100 Hz, and is thought to be associated with higher level mental activity. Gamma was examined using spectrographic analysis, which is not shown here.

⁷Fluorocitrate is only injected into one hemisphere of the brain, so any resulting spreading depressions are constrained to that hemisphere. Seizures, however, can spread from one hemisphere to the other.

⁸Spreading depressions in citrate experiments appear to be caused by physical action on the brain, such as the insertion of needles, and not as a result of the drug itself.

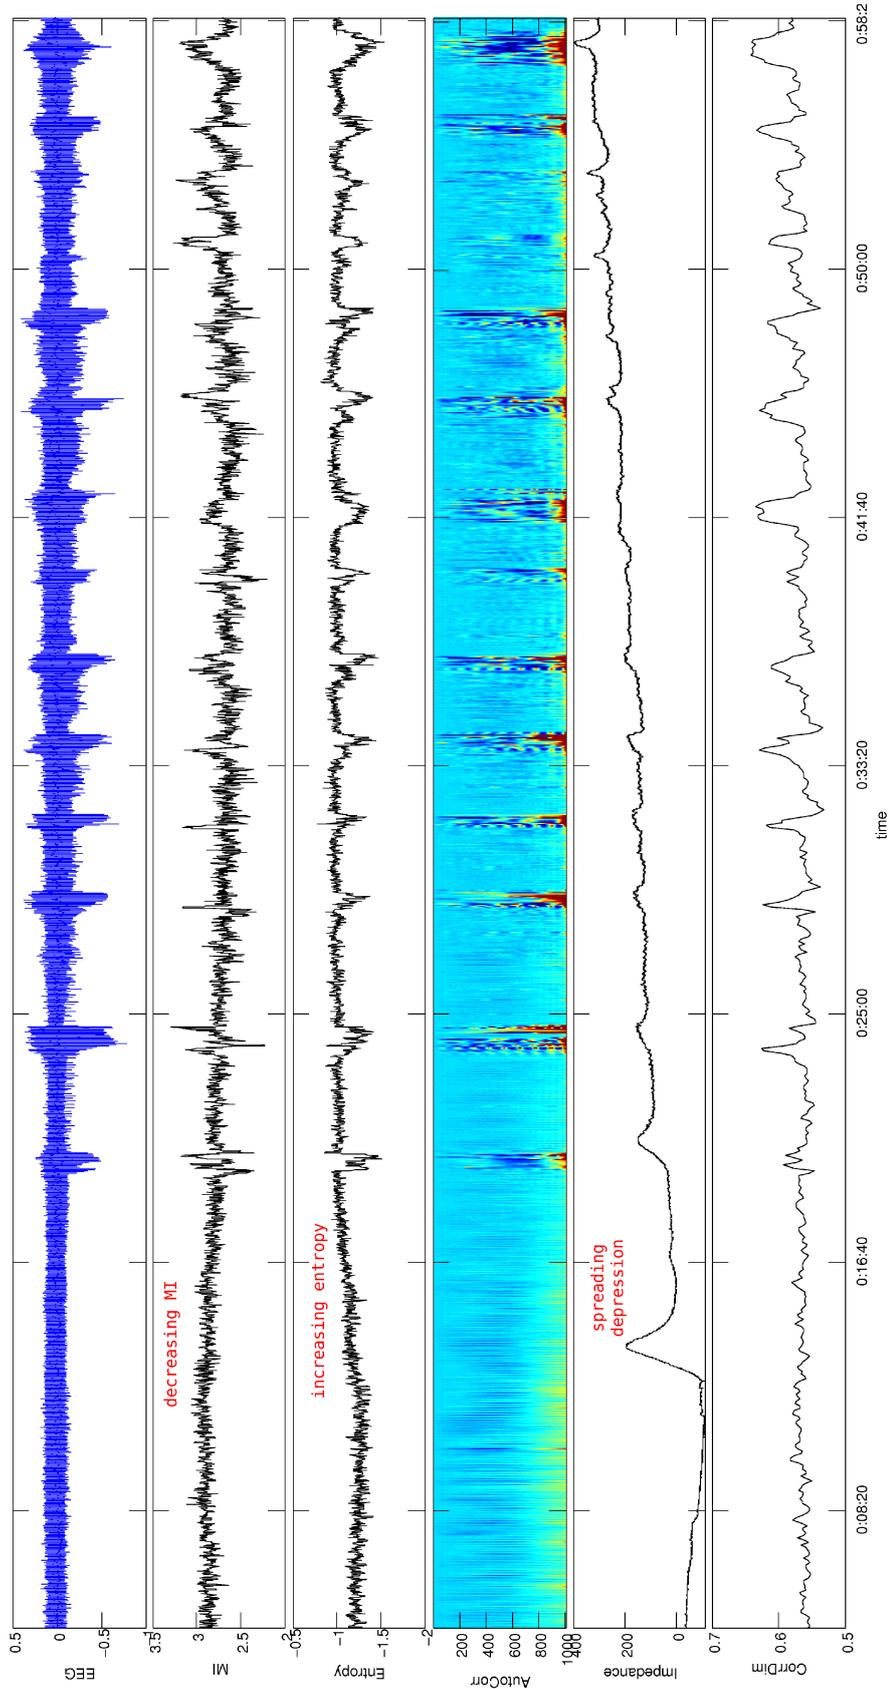


Figure 8.2: SD98, kainic acid

This experiment shows a series of seizures, and a coincident steady increase in brain impedance. There is a decrease in MI and an increase in entropy (at 00:12:30) preceding any ictal activity and possibly leading the beginning of the impedance increase. Gamma becomes prevalent in the EEG from approximately 00:16:40 (not shown), and steadily intensifies as the first seizure approaches, remaining elevated for the rest of the experiment. There are increases in correlation dimension during the initial stages of impedance transients which rapidly decreases as they progress.

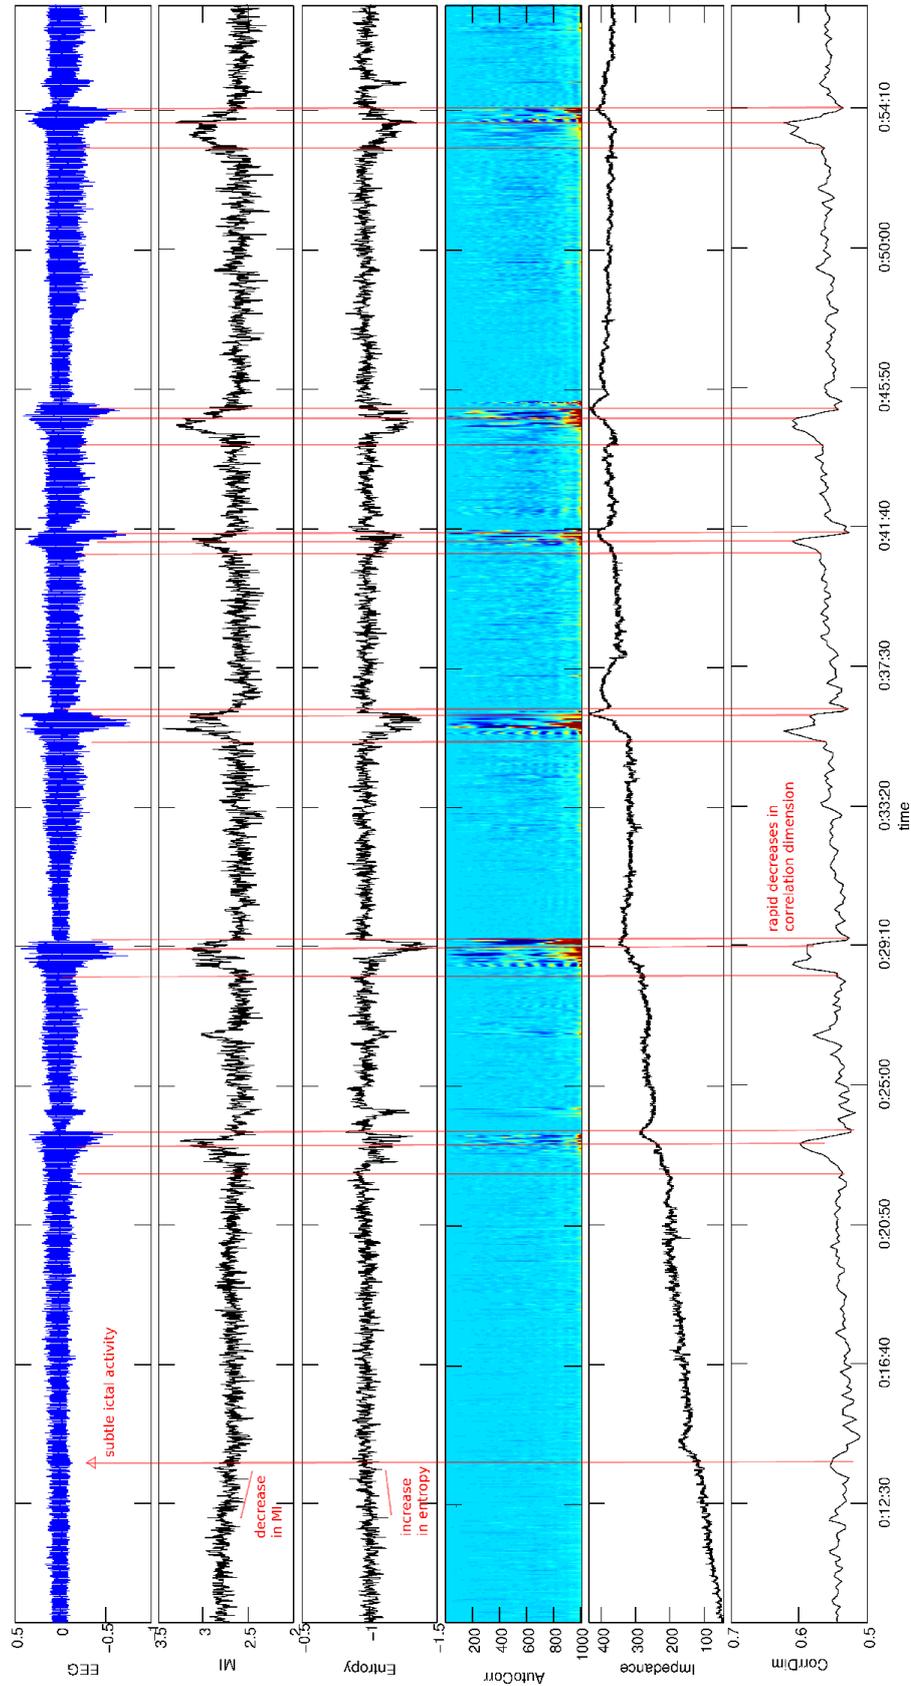


Figure 8.3: SD64 kainic

In this kainic acid experiment, the impedance shows an increase prior to ictal EEG activity, or changes in the analyses. The impedance transient at 00:13:30 is slightly later than a subtle burst of ictal activity (red arrow). There is a slight decrease in MI and a slight increase in entropy at 00:12:30 (prior to the aforementioned ictal activity). A detail of this early activity appears in figure 8.4. Annotations are in red.

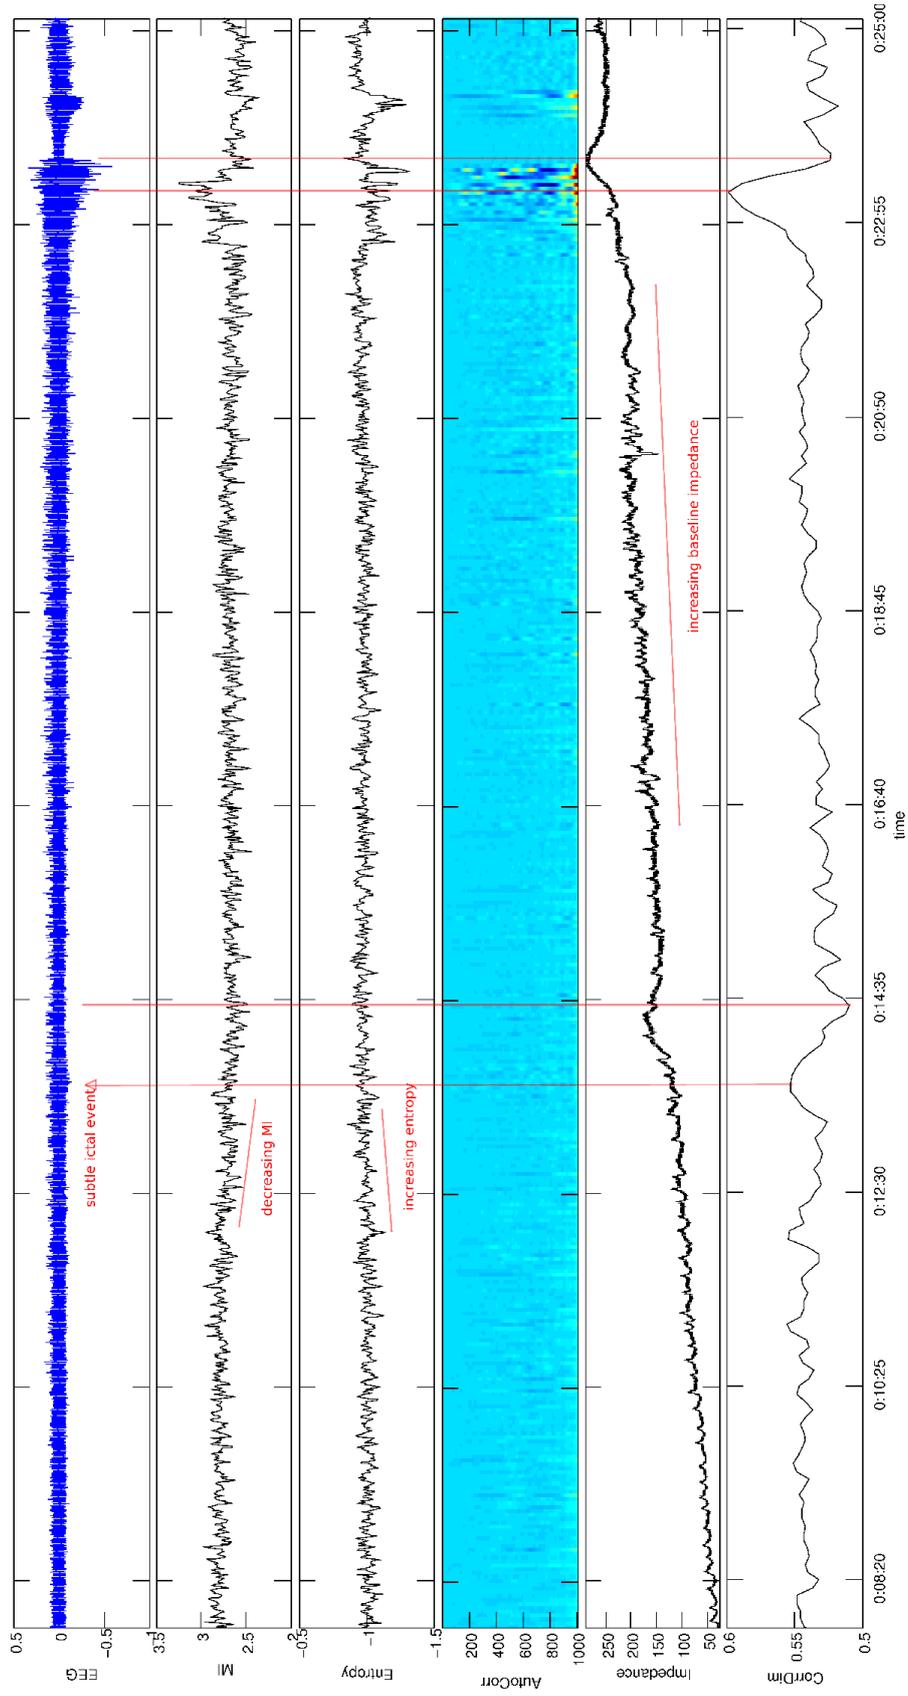


Figure 8.4: SD 64 zoom kaimic

A detail of the same data as figure 8.3. Note the decrease in MI and increase in entropy visible at 00:12:30. The ictal activity is apparent from 00:13:30 to 00:14:00 and the impedance transient occurs between 00:14:10 and 00:14:40, approximately coincident with a change in correlation dimension.

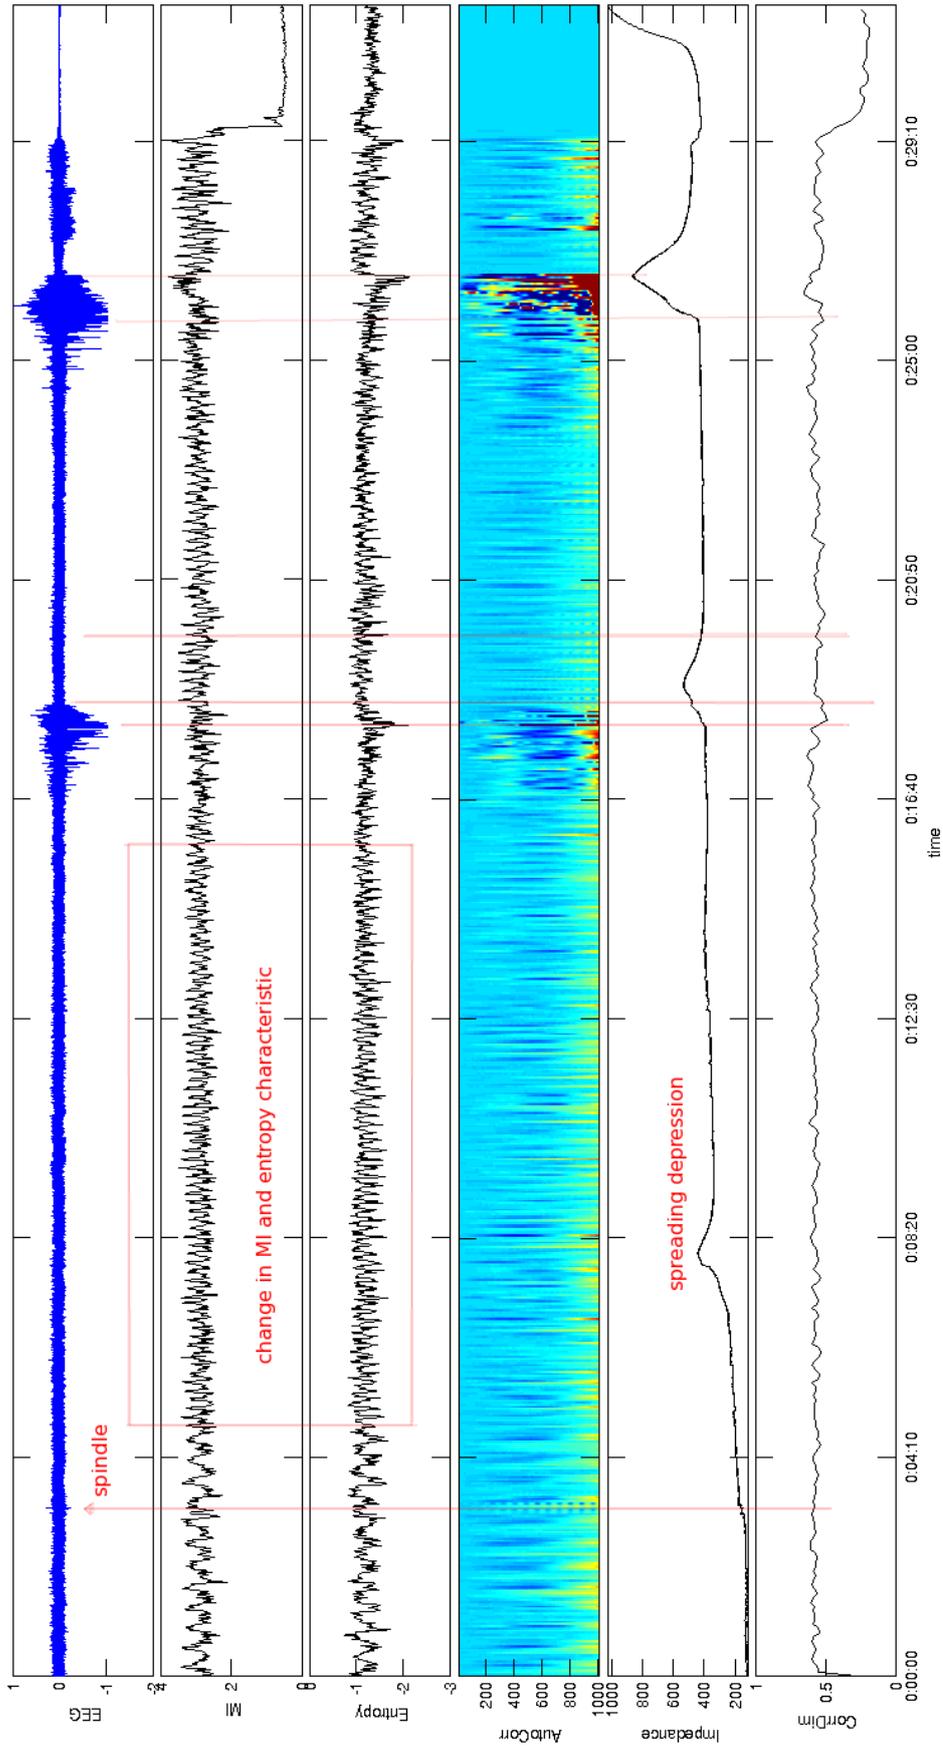


Figure 8.5: SD94 - picrotoxin

As shown, the impedance steadily increases coincident with the first spindle (apparent in the EEG and autocorrelation at approximately 00:03:00). There are 3 impedance transients, and the latter two appear to correspond to the two seizures. The second seizure shows stronger auto-correlation. The first impedance transient (at 0:08:20) has no associated novel EEG activity, although MI and entropy show a change in behaviour in this region. The two impedance transients are likely spreading depressions. The rat was anaesthetised at 00:29:10 and the death curve starts after 0:30:30.

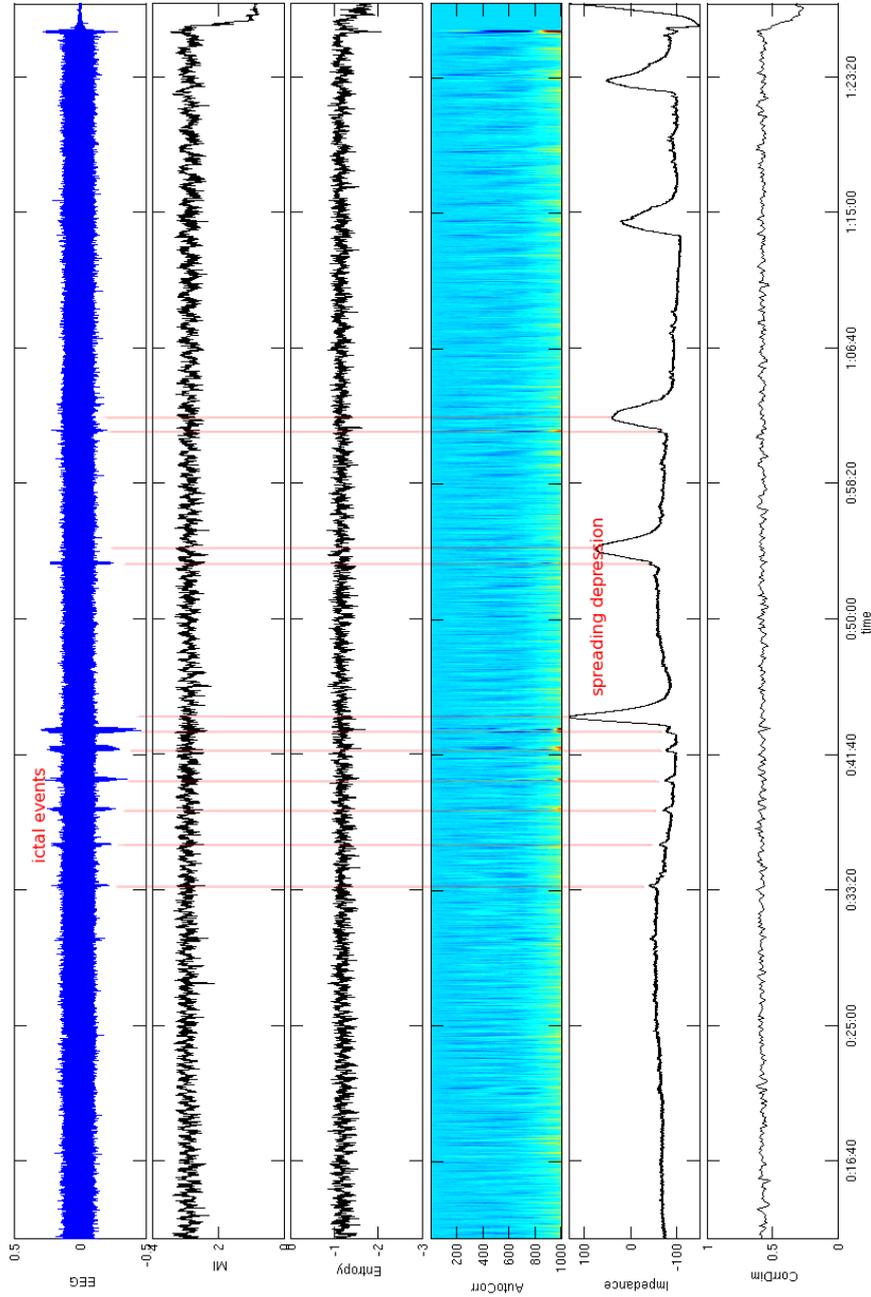


Figure 8.6: SD79 - fluorocitrate

This figure shows some EEG activity that could be ictal between 00:30:00 and 00:43:00 – this activity is mirrored in small impedance changes, and is indicated by a detail of these events in figure 8.7. Following this, there are five spreading depressions, and the beginning of the first three follow ictal EEG activity. The MI and entropy analyses show much smaller changes for the seizures in this experiment, suggesting that they are different from the ictal events occurring due to the administration of picrotoxin and kainic acid.

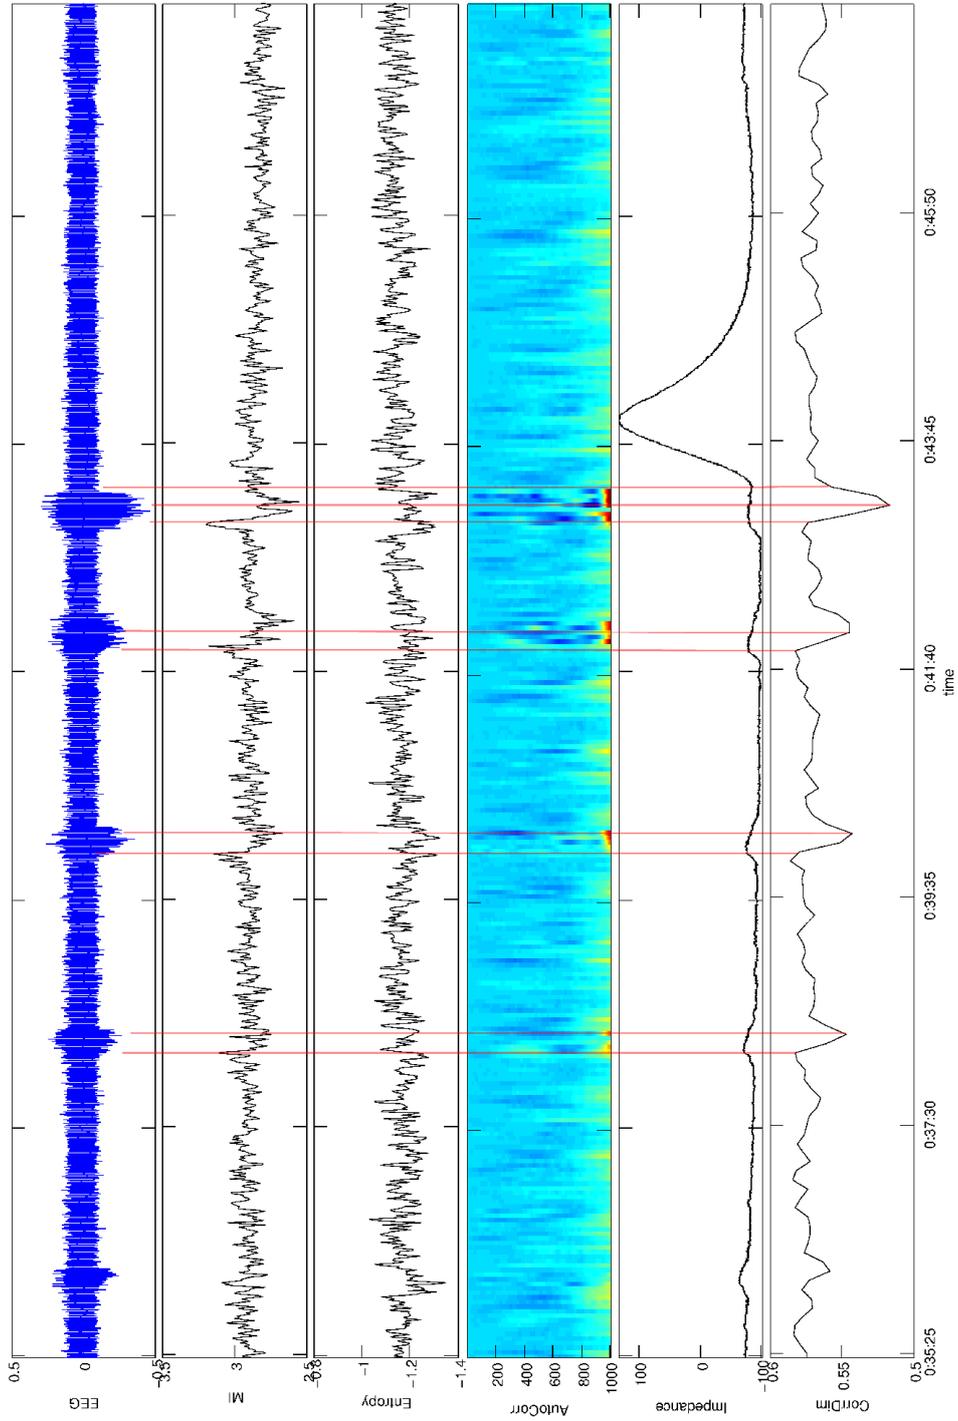


Figure 8.7: SD79 - fluorocitrate, detail
 This figure shows a detail of the data in figure 8.6. Specifically, it shows that the ictal EEG events manifest in MI, auto-correlation and correlation dimension, with more subtle effects also visible in the entropy.

In general, impedance data complement EEG data: epileptiform activity is associated with changes in brain impedance. However, there are cases where there are changes in impedance unaccompanied by obviously related changes in EEG and visa versa. Thus, it appears that both impedance and the analyses are affected by epileptiform activity, however, they respond differently and there is not a consistent relationship between them.

8.3.1 Analysis of EEG

Changes in the analyses may correlate with with impedance with no clear change in contemporary EEG (figure 8.3), although spectrographic analysis sometimes shows changes in gamma power. Spreading depression (section 5.4.3) tends to reduce EEG activity and increase brain impedance (figure 8.5).

8.3.2 EEG seizure activity

During ictal activity (figures 8.2, 8.3, etc), there was an increase in MI suggesting that there is an increase in brain synchronicity across the hemispheres⁹. Entropy generally decreased, indicating a loss of complexity in the EEG¹⁰. Ictal activity usually results in an increase in auto-correlation, which appeared as narrow, repetitive bands of *sameness*.

Correlation dimension usually increases during the initial stages of ictal activity, before decreasing sharply. The increase may begin slightly in advance of obvious ictal EEG activity, but is usually in advance of the associated transient impedance increase. Such transients are often synchronous with the sharp decrease in correlation dimension (figures 8.2 and 8.3).

Impedance transiently increases during, or immediately following, ictal

⁹The increase in MI during ictal activities cannot be attributed to the increase in EEG amplitude, because the MI calculation operates on rated data – a process that ignores the absolute amplitude of the data.

¹⁰Recall that prior to entropy analysis the data are demeaned and scaled to unit variance, reducing sensitivity to signal amplitude changes. This implies that the observed entropy changes are not merely the result of a change in EEG amplitude.

events (figure 8.2).

8.3.3 Inter-ictal periods and seizure anticipation

A decrease in baseline¹¹ MI can precede any apparent ictal EEG activity (figures 8.2 and 8.3), although this is sometimes possibly attributable increases in gamma EEG energy (eg. SD98, figure 8.2).

There are experiments where a slight increase in baseline entropy may precede seizure activity (figures 8.2, 8.3 and 8.4) but possibly not changes in baseline impedance.

Autocorrelation periodically shows bands in inter-ictal periods that are similar, but less pronounced, to the ictal bands. These are more common in kainic acid (eg. figures 8.2 and 8.3) experiments, but are not absent in picrotoxin animals (figure 8.5 shows a spindle at 00:03:00, apparent in the auto-correlation, that is coincident with a rise in impedance).

8.3.4 Cell swelling

Impedance has a complex relationship with EEG. Cell swelling can occur as a baseline change, apparently as seizure approaches. We also observe transient impedance changes that appear to be associated with ictal activity or spreading depression. Impedance recordings show several behaviours: incidences of spreading depression and cell swelling that *precedes* ictal EEG activity (figure 8.3), cell swelling that *lags* ictal EEG activity (figure 8.5), and cell swelling that is associated with increases in gamma EEG power (figure 8.2). There are impedance changes that are not associated with ictal activity (figure 8.5 – the first impedance peak, at 00:08:30, is not coincident with any obvious EEG activity, whereas the second shows a seizure, marked decrease in correlation dimension and changes in MI and entropy).

¹¹I use the word *baseline* to denote a change in the general trend of the analysis – the trend without the transient changes such as those associated with ictal activity.

8.3.5 Spreading depression

It appears that seizures might induce spreading depressions (figures 8.4 and 8.5), though they can also arise without obvious associated EEG changes (figure 8.5, earlier). Such spreading depression can be associated with a change in baseline MI, chiefly in kainic acid experiments (figure 8.2).

Spreading depressions in fluorocitrate experiments cause changes that MI appears unable to detect (figures 8.6 and 8.7).

8.4 Discussion

8.4.1 Pre-ictal EEG

Although some experiments showed changes in the impedance and analyses that were in advance of obvious ictal EEG, we did not see any indication of a consistent pre-ictal state heralding the immanence of a seizure.

Figure 8.5 shows an interesting pre-ictal phenomenon. There are changes in MI and entropy, commencing 10 minutes in advance of the first ictal activity. This was not observed in other experiments.

8.4.2 Ictal EEG

Ictal EEG activity is associated with increased impedance which, as a direct correlate for cell-swelling, implies that ictal brain activity might *cause* cell swelling¹². Ictal EEG activity has increased synchronicity and contains less information, compared with non-ictal EEG. The correlation dimension of ictal EEG varies in a complex manner. From the observed results, it appears that a seizure detection algorithm could be implemented using the EEG, impedance and analyses as input variables.

¹²Repeated neural firings cause the release of glutamate into the extra-cellular space, where it is removed by astrocytes. Sufficient quantities will cause them to swell because of osmotic imbalances (section 5.3.3). Osmotic imbalances can also lead to neuronal swelling, because of the movement of ions required for action potentials to occur. This is well described in scientific literature, eg. [54].

There are sometimes changes that occur in both the analyses and the impedance data prior to seizure onset. When they occur, these changes are typically 1 - 2 seconds in advance of apparent changes in EEG. It is worth attempting seizure prediction by exploiting these changes which, while having no immediate practical application, would demonstrate the existence of a short, detectable, pre-seizure state.

8.4.3 Seeking a predictable relationship between EEG and impedance

While there are changes in the analyses that reflect changes in EEG, none of the analyses demonstrate a clear and consistent relationship with changes in the impedance¹³. While it seemed possible that changes in impedance reflected EEG changes (section 5.3.3) and might demonstrate the existence of a pre-seizure change in EEG occurring in parallel with changes in impedance, we have not been able to unequivocally observe this.

To view this another way, it is unclear whether there is a consistent relationship between EEG data, the analyses, and impedance data. There appear to be changes in the analyses that are indicative of impedance changes, but this is difficult to quantify¹⁴ by visual examination. In order to better examine the data for such a relationship, a good approach would be to attempt to model the impedance data – making an estimate of the impedance from the EEG and EEG-derived analyses. This is examined in chapter 9.

8.4.4 Spreading depression

Interestingly, while brain changes occurring during spreading depression are detectable in kainic acid and picrotoxin experiments, they are not detectable

¹³This is unsurprising, since all the analyses are derived from the recorded EEG only and not the impedance

¹⁴For example, it is possible that each analysis, in isolation, does not have a consistent relationship with impedance but that the analyses together do.

in fluorocitrate experiments. It may be possible to detect spreading depression from EEG and analyses, however it appears that such a model would be drug-specific (due to differences between the epileptogenic drugs).

8.4.5 Surrogate data

Surrogate data analysis (section 7.5) allows verification that the changes observed in nonlinear analyses are due to changes in the nonlinear components of the EEG. However, it is very computationally expensive¹⁵ and for this reason was not able to be implemented here (the analyses themselves were examined in section 7.5, however similar simulations were not run for every rat experiment). Improvements in computer power will, in the future, make this a non-issue and we recommend making use of surrogate analysis as an intrinsic component of EEG signal analysis, not merely part of an algorithm validation procedure.

We must also remember that, assuming the EEG data are equivalent to those in section 7.5, that the correlation dimension is examining only low-order, linear statistics and that the entropy is not exclusively sensitive to higher-order statistics.

8.5 Conclusion

There are changes in impedance that occur as the experiment progresses. Ictal EEG is related to increases in impedance, increases in MI and decreases in entropy, in accordance with our understanding of seizure processes. While the various implemented analyses reflect changes coincident with (and sometimes in advance of) EEG events, there is no evidence of a consistent predictive element. Nor did we detect evidence of a pre-ictal intermediate state.

¹⁵The generation of iAAFT surrogates is computationally expensive and, because many surrogates must be generated and analysed (to obtain $p < 0.05$), the analysis time is greatly increased.

There appears to be some relationship between the EEG (and analyses thereof) and the impedance data. Whether this is consistent is something to be investigated.

Chapter 9

Experiment: Rat impedance modelling during epileptiform events

We recorded extra-dural EEG and impedance data from respiration, paralysed rats. Animals were injected with kainic acid, via an implanted venous catheter, in doses designed to induce epileptiform activity.

The EEG was analysed using entropy, mutual information, correlation dimension and auto-correlation algorithms (the *analyses*). The analysis results and the raw EEG were used to train a feed-forward neural network to model the recorded impedance from the EEG and analyses only. This attempt was not successful, therefore we conclude that impedance data are (at least partly) independent of EEG. We recommend the continued recording of impedance data along with EEG data.

9.1 Hypothesis

The use of nonlinear analysis of EEG, coupled with contemporary and historical EEG values, can be used to estimate impedance.

9.2 Introduction

The recording of impedance adds complexity to the experimental apparatus, so it would be useful to demonstrate that the impedance data convey new information not present in the EEG data. Thus, we attempt to model the impedance as a function of EEG and the analyses, to demonstrate the worth of impedance recordings.

Successfully estimating impedance from EEG might seem difficult due to the holistic and stable nature of impedance and the local and unstable nature of EEG. However, it seemed possible that locally-derived EEG and analyses might reflect some aspects of generalised brain electrical activity. For this reason, we decided to attempt such estimation.

9.3 Method

9.3.1 Data Analysis

For this experiment, the rat EEG data were gathered as described in section 6.1, and analysed as described in section 8.2.

Each channel¹ was mapped to the range $[-1\ 1]$ across all animal experiments – this maintained differences between animals, but ensured that the amplitude of each channel (across all animals) was constrained.

9.3.2 Data assembly

Some rat experiments contained discontinuities in the impedance and EEG channels (eg. due to technical artifacts). These experiments were excluded from the modelling.

Because of computer limitations, we sub-sampled the impedance data – these were *goal data*. Each of these data were associated with a vector of *train*

¹The channels were: EEG1, EEG2, MI, entropy, correlation dimension, auto-correlation and impedance.

data: periodic windows of historical EEG and analyses data (700 historical records, spaced at 8 sample² intervals). This was repeated for each of the seven rats, and the results concatenated to yield a vector of impedance values (*goal data*), associated with a matrix of contributory historical EEG and analyses values (*train data*). This is illustrated in figure 9.1.

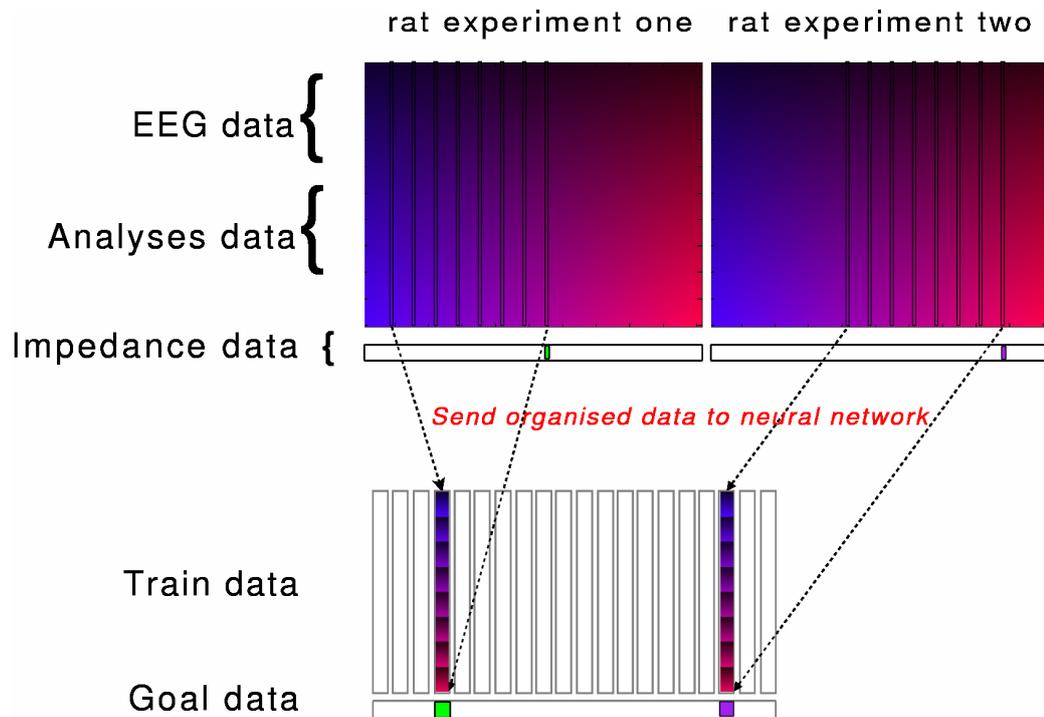


Figure 9.1: Neural net training data organisation

This figure illustrates the method by which the data from each rat were assembled into a matrix ready for training the neural network. The two matrices (top) represent the EEG and analyses from two experiments (channels by time). Shown in each impedance data set are two chosen values of impedance, each associated with a series of vectors in the analyses/EEG data. This series of vectors is linearised and inserted into the training data matrix as shown, while the goal data are sub-sampled directly from the relevant impedance data. The train data and goal data are ready for presentation to the neural network.

In the actual experiment, there were 700 historical vectors (instead of 8, as shown here), 7 rat experiments (instead of 2, as shown here) and there were many impedance values and associated analyses vectors (two are shown in this figure).

²Sample refers to a sample of the analyses, which are at a lower sampling rate than the EEG. Each sample is of equivalent duration to 250 EEG samples (the specifics of the EEG analysis are described in section 6.1).

9.3.3 Neural Network modelling

All modelling was performed using a feed-forward neural network with 50 neurons in the hidden layer. After modelling, the network's training history was verified to confirm appropriate training had occurred (eg. examining the gradient history). The task of the neural network was to estimate the impedance measured in the experiment, based on the EEG and analyses thereof. We trained a simple feed-forward network, but because train data contains *historical* data pertaining to each value of impedance goal data, the feed forward network behaves in a manner similar to that of a feed-forward network with tapped delay. This is because the network has knowledge of historical values. However, this approach has the benefit of allowing us to present the network with all the data (including historical data) for all rats, simultaneously. This would be difficult if we were presenting EEG and analyses that were contemporary with the impedance value we wished to model, expecting the network to use delay to retain knowledge of historical values.

As an initial step, every experiment was included in the training and those same experiments were then simulated to estimate the impedance. To emphasise, *we simulated the same experiments that had been used to train the network*. Such a process could not test the network as a predictive tool, but only as a transfer function (EEG and analyses \rightarrow impedance).

We then performed a leave-one-out system of training and testing, whereby a network was trained on all applicable experiments *except one*, and then simulated on that experiment to estimate the impedance – a process repeated for each experiment. Leave-one-out analysis requires a new network to be created and trained for every modelling attempt. It is used when there are insufficient data to split them into *train* and *test* groups. In this experiment, each rat's data were successively isolated from the other data. A new network was then trained on the remaining data and simulated on the isolated rat's data. This approach examines the generalisability of the trained network, and

(by repeated application to each rat’s data) gives an indication of the expected performance of the hypothetical single network that we would implement if there were more data.

9.4 Results

I first tested a network by training it with all available data, and then testing it on individual rats (a subset of that data). As can be seen in figure 9.2, the neural network is able to approximate the impedance signal, although noisily. This means that, although not demonstrating a generalised predictive ability, the network is able to form a transfer function that can represent impedance from the analysis of the EEG data. This demonstrates that the setup and configuration of both the training and goal data is appropriate, and the network is correctly configured and trained.

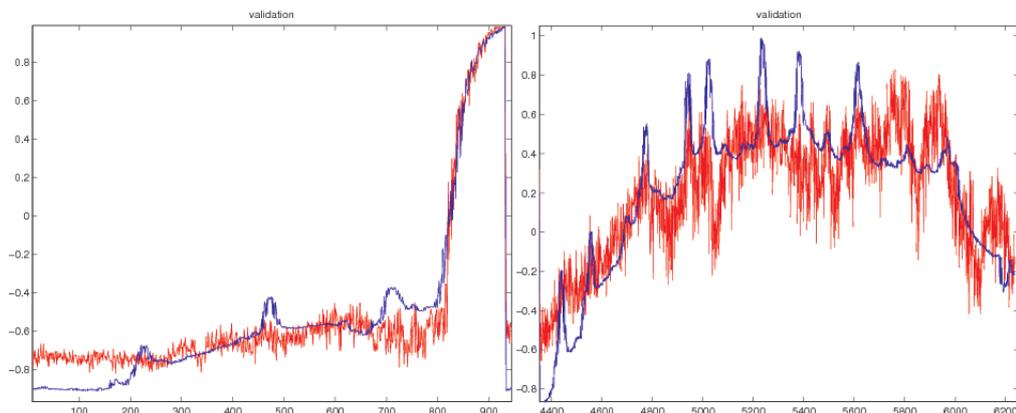


Figure 9.2: Estimate of impedance – familiar data

Each graph shows two data sets. The blue data are the actual impedance data that the network was attempting to model. The red data are the network’s estimate of the impedance. Over all the experiments, the RMS difference between the goal and estimate, when performed on familiar data, was 0.17 ± 0.07

For this test, the data to be tested were included in the training data, so this model does not test predictive ability. However, it demonstrates that the setup and training procedures for the neural network were performed correctly – a validation of the experimental setup.

Next, we implemented a leave-one-out analysis, so that the rat data to be tested were *not* included in the training. If this were successful, it would mean

that a generalised transfer function could be trained on some rats and then simulated on others. This would allow an estimate of impedance, from EEG alone, that would render the recording of impedance unnecessary. As shown in figure 9.3, however, the network was unable to estimate the impedance and hence cannot be used as a general tool for predicting impedance.

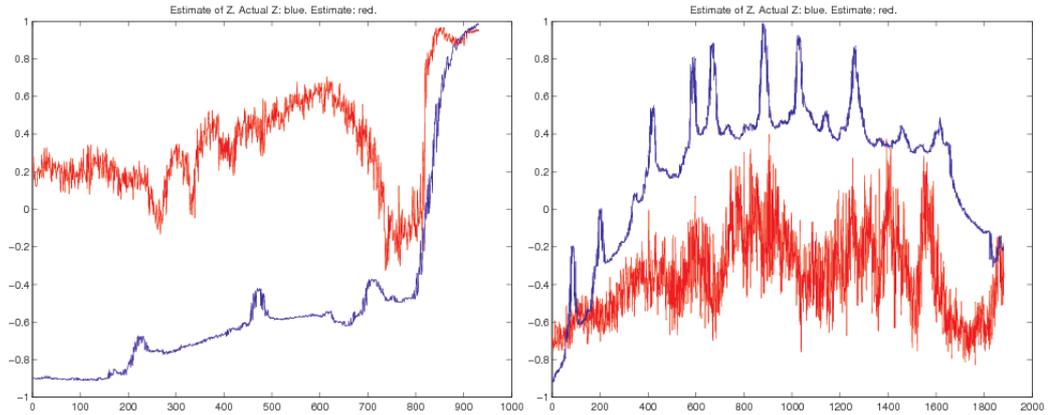


Figure 9.3: Predictive test of impedance estimate

Each graph was produced in a similar fashion to figure 9.2, except that the test data were not included in the training. The blue data are the actual impedance data that the network was attempting to model. The red data are the network's estimate of the impedance. This figure thus illustrates the effectiveness of the model as a generalised predictive tool (*i.e.* the ability to learn about the relationship between impedance and EEG in all rats **except one** and use the trained model to test that relationship in the unseen rat). As you can see, the performance of this model is substantially worse than that in figure 9.2. These figures are representative of the other results: The RMS difference between goal and estimate, on unfamiliar data, was 0.79 ± 0.19 .

9.5 Discussion

Figure 9.2 demonstrates that the configuration and training methodology are appropriate because it shows a reasonable estimation of the impedance recording based only on EEG data. The data shown in this figure are the raw output of the neural network, and could be improved by additional network training or post-simulation filtering. However, networks that better model the training data tend to become over-fit to the train data – and lose generalisability³ as a

³*Generalisability* is the ability of a network to learn a rule on one set of data and then apply that rule to novel data.

result. For this reason, we refrained from training a network to exactly model the training data.

An attempt to form a generalised model of the relationship between EEG and impedance is made in figure 9.3. As can be seen, there is little relationship between the impedance estimate and the measured impedance. This figure is representative of the all impedance estimate attempts. This failure does not mean that it is impossible to estimate the impedance from the EEG, but it suggests that a more sophisticated methodology is required, if it is to work.

As mentioned, not all data were presented to the neural network. This was due to CPU and RAM limitations. If all the data were presented, the training matrix would be approximately 17000 x 17000 instead of 2100 x 8500. Although it would be impractical to manipulate on our current computers, it is possible that such an analysis would yield more useful results.

9.5.1 EEG vs Impedance

These results suggest that the formation of an estimate of impedance is impossible from current EEG recordings, and that a much greater number of electrodes⁴ are required. Because impedance, unlike EEG, is a holistic brain property, it seems likely that it provides data that are (at least partially) independent of EEG. Such a conclusion is supported by this work, and we recommend the continued recording of impedance data in parallel with EEG.

9.5.2 Further work

The nonlinear methods and modelling software could be adapted to attempt seizure prediction. Much effort has been devoted to seizure prediction but, to our knowledge, none of the attempts to predict the onset of seizures have used impedance recordings as a signal source. Because we have shown that

⁴Given sufficient electrode density for EEG recordings, impedance estimation is likely to be possible, however with only two electrodes successful modelling remains elusive.

the data in the impedance channel might be complementary to the data that were derived from the EEG, impedance data could potentially assist attempts to predict the onset of seizure.

9.6 Conclusions

In this experiment we have attempted to create a model to allow an estimation of rat brain impedance changes from a knowledge of EEG and nonlinear analysis thereof. We found that the impedance estimate does not adequately track the measured impedance. Based on this analysis we conclude that the impedance cannot be predicted from the EEG at this stage, hence measuring impedance experimentally will yield information about the brain that is not present in EEG and EEG-derived analyses.

Chapter 10

Experiment: Human data analysis

Subjects were chosen from control (15 people) and primary generalised epilepsy (9 people) groups. Each subject's EEG was recorded while they performed a series of mental tasks. All EEG was pre-ictal, ie. no epileptiform EEG activity was recorded. The data were stored and were analysed offline using entropy, correlation dimension and mutual information algorithms (the *analyses*).

The results of these calculations were compared between subject groups, between mental tasks and between brain regions to evaluate the effectiveness of the analyses at detecting changes in EEG.

We found that there are detectable differences in the EEG, dependent both on the task being performed, and on the subject group.

10.1 Hypothesis

Nonlinear analysis of non-ictal scalp EEG can detect differences dependent on the mental task being performed, and the presence or absence of primary generalised epilepsy (PGE).

10.2 Introduction

This experiment follows from our earlier exploratory ($n = 3$) findings that there were detectable differences in the gamma frequencies of EEG in subjects with primary generalised epilepsy [84] – results that were not confirmed by our larger follow-up experiment. Based upon this work, it yet seems possible that there are differences, not apparent in a spectral analysis of EEG, but detectable in phase space using nonlinear analysis.

10.3 Methods

EEG data were collected and stored from human volunteers, in control or active¹ PGE groups, as described in section 6.2.

10.3.1 Data Analysis

Software was written in Matlab (Mathworks) to analyse the data using entropy, correlation dimension and mutual information algorithms (Chapter 7).

10.3.2 Data display methods

There were two methods by which the data were visualised.

Local MI variations (adjacent electrode pairs) To examine regional variations of local MI, we calculated the mean and *standard error of the mean* (SEM) of the local MI results within each brain region. We displayed the results from the control and PGE groups (figure 10.1).

Regional MI variations To examine variations in regional MI, we calculated the mean and SEM of the regional MI results between each combination of brain regions. We visualised the results from the control and PGE groups.

¹*Active PGE* refers to subjects who were either untreated, or were ineffectively treated.

Figure 10.2 shows the left-frontal region vs all other brain regions; the other results are shown in the appendix.

10.4 Results

We found that the various analyses were useful to varying degrees in making these assessments. In particular, the most useful analyses for distinguishing between the groups and tasks were the two mutual information analyses, as shown. The entropy and correlation dimension analysis (shown in appendix A) did not yield results useful for discriminating in this way, and are not shown in this chapter. This was unsurprising, considering the results of surrogate analysis (section 7.5).

10.4.1 Local MI variations (adjacent electrode pairs)

Figure 10.1 illustrates variations in local MI, by task and by brain region. Depending on the task being performed, there are regional changes in localised MI in control subjects. There are clear differences between the PGE and control groups in several of the brain regions. Also the variation across tasks is similar across brain regions, in both PGE and control groups.

10.4.2 Mutual information analysis by brain region

Figure 10.2 shows the mean MI between the left frontal cortex and each other region². Each axis represents a brain region (axis title) and shows MI (relative to eyes-closed task) between that brain region and the left-frontal region. MI by brain region shows a difference between the PGE and control groups is qualitatively less marked, than that in the local MI analysis. There is a consistent variation across tasks, and this occurs no matter which brain region

²Figures detailing the other regions are shown in the Appendix (section A.1) – these include the remaining seven regional MI plots (with different regions of interest), as well as figures showing entropy and correlation dimension changes.

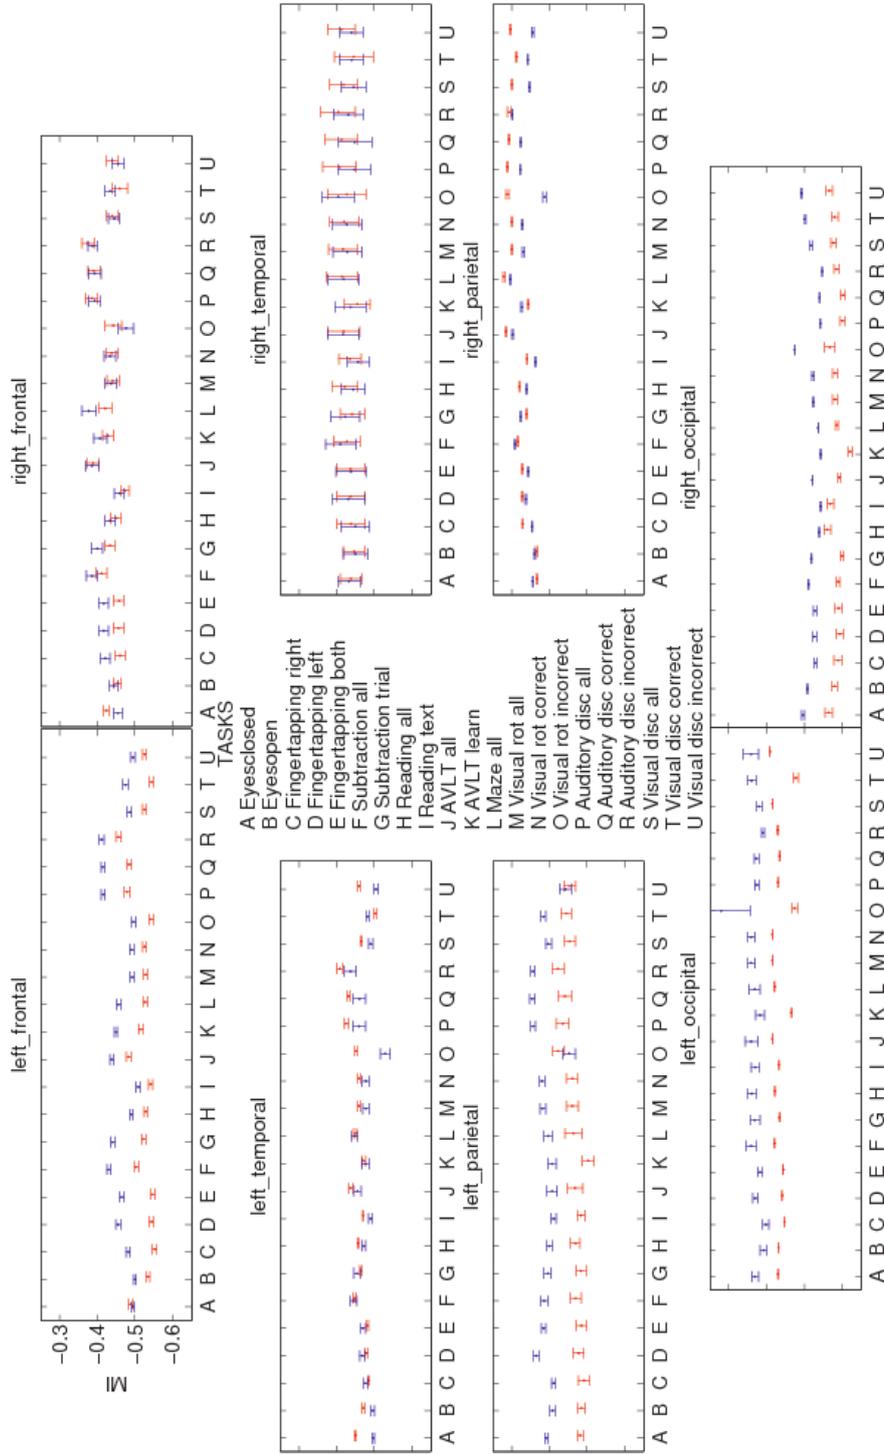


Figure 10.1: Local MI variations during tasks, control (red) and PGE (blue) subjects. Each axis in this figure corresponds to a brain region. The error-bars represent the mean and SEM of the MI of each task, within each brain region and task. The MI was calculated between pairs of adjacent electrodes, and averaged within brain regions. Thus, the data shown represent the manner in which local MI varies depending on mental task.

is the region of interest.

It appears that the subtraction task produces results that are different to the other tasks, and it also shows separation of the PGE and control groups, which no other task consistently shows.

10.5 Discussion

10.5.1 Local MI variation

When examining local MI, PGE vs control (figure 10.1), there were many tasks and regions that exhibited altered MI. Left frontal, left temporal, left occipital and right occipital all showed elevated MI in PGE vs control subjects. Individually, these differences were not statistically significant however.

There were many task/region pairs³ that showed significant ($p < 0.05$) differences in local MI. Because there are so many comparisons, we would expect to see 5% appear significant even if there was no difference between groups. However, the fact that approximately 18% were reported as significant suggests that there are differences between the tasks.

10.5.2 Regional MI variation

Subtraction is elevated in PGE vs control, in all regions (frontal region, figure 10.2, other figures are in the appendix). Eyesclosed is suppressed in PGE vs control in most brain regions, as are the visual discrimination and reading tasks.

There were many tasks/region pairs that showed significant differences between control and PGE groups. Every task had at least one region that showed significant ($p < 0.05$) differences between groups. Of all the comparisons made, approximately 20% were reported as significant. This suggests that there are

³A task/region pair refers to a comparison of MI between adjacent electrodes in a task in a region, with MI between adjacent electrodes in another task in another region.

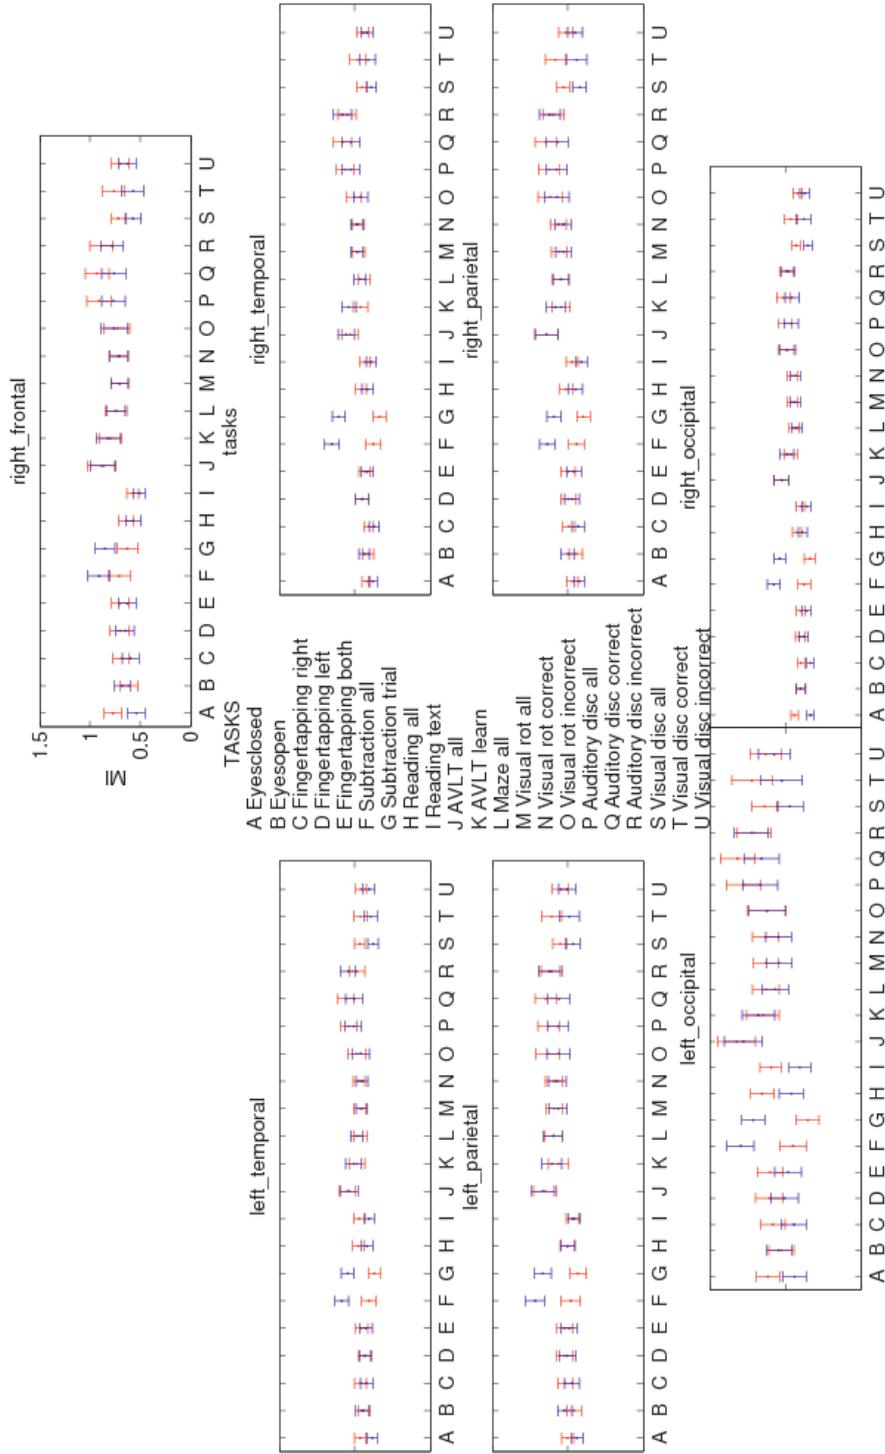


Figure 10.2: Regional MI variations during tasks, left frontal vs other brain regions, control (red) and PGE (blue) subjects. This figure shows the mean MI between the left frontal cortex and each other region, for control subjects. Each axis represents a brain region (axis title) and shows MI (relative to eyes-closed task) between that brain region and the left-frontal region, for each task. Error bars correspond to standard error of the mean (SEM).

differences between the two groups.

10.5.3 Group Classification (diagnosis of PGE)

These results suggest that the subtle group differences might be exploited to diagnose PGE in an unknown subject. This is more difficult than identifying statistical differences between two populations because there is a large intra-group variance and a relatively small inter-group variance. There are many data to use – small differences over many data may allow accurate classification. Success will depend on there being sufficient *statistical independence*⁴ between the variables. This is the subject of the following chapter.

10.5.4 Task-classification

These results also suggest that it may be possible to identify the task being performed by the subject using the analysis of EEG. This would have implications for brain-computer-interface (BCI) research. Robust task identification will require several variables with sufficient statistical independence because of small inter-group and large intra-group variances. This is the subject of chapter 12.

10.6 Conclusion

There are differences in the inter-ictal EEG between control and PGE subjects that are detectable by local and regional MI analysis. However, the intra-group variance is quite large, and the inter-group variance is quite small. These results demonstrate that it may be possible to classify an unknown subject into their group (control or PGE), if there is sufficient statistical independence between tasks. Similarly, there are some detectable differences between tasks,

⁴If two variables, β and α , are statistically independent, then one will learn nothing about variable β by examining variable α .

that may allow classification between certain tasks. These two classification problems are the subject of the following two chapters.

Chapter 11

Experiment: Diagnosis of PGE in human subjects

Subjects were chosen from control (15 people) and primary generalised epilepsy (8 people) groups. Each subject's EEG was recorded while they performed a series of mental tasks. All EEG was inter-ictal, ie. no epileptiform EEG activity was recorded. The data were stored, and were analysed offline using entropy, correlation dimension and mutual information algorithms (the *analyses*).

These data were used to train, via a leave-one-out method, a memoryless feed-forward neural network to classify the subject as either control or PGE. The networks had 15 logsig nodes in the hidden layer, and 2 logsig nodes in the output layer.

The best results showed an accuracy of 87%, an excellent result, apparently allowing an accurate diagnosis of PGE from analysis of non-ictal EEG. These exciting results should be pursued and duplicated¹ and an improvement in sensitivity and specificity sought.

¹It is important to duplicate this experiment, because of its post-hoc nature. New experiments need to be performed, and the developed algorithms applied, to address the potential problem of "over-mining" the data (despite great care being taken, over-mining is an inherent problem with this kind of analysis and needs to be discounted).

11.1 Hypothesis

It is possible to identify subjects with primary generalised epilepsy by analysis of experimentally gathered inter-ictal (apparently normal) EEG data.

11.2 Introduction

Chapter 10 sought variations between control and PGE classes. These variations suggested that classification of an unknown subject might be difficult, mainly because of the small inter-group variation and large intra-group variation. That there were multiple variables exhibiting these differences would be helpful, provided the variables were independent. Given sufficient statistical independence between several variables, a multi-variate classification process can take advantage of many small differences to draw a reliable conclusion.

11.3 Method

A brief summary of the method follows:

We have 24 subjects to test. We iterate through the subjects, testing each subject in turn. Each test subject is quarantined from the remaining subjects. The un-quarantined subjects are then used to train a feed forward neural network. Once trained, the network is used to classify the quarantined subject. We then proceed to the next subject, repeating this procedure.

See figure 11.1 for a graphical representation of this process.

11.3.1 Data collection

EEG data were recorded from control (15) and active primary generalised epilepsy² (PGE, 8) subjects, as described in section 6.2. Software was written in Matlab (Mathworks) to analyse the data using entropy, correlation dimen-

²Active PGE refers to subjects who were either untreated, or were ineffectively treated.

sion and mutual information algorithms (the *analyses*, chapter 7). Both the Laplacian estimate and the raw EEG were used as a basis from which to estimate the various nonlinear quantifiers.

The EEG of each subject's tasks was divided into epochs³, of minimum length 3500 samples (1.75 s), and analysed separately. The results were then collated across all subjects.

11.3.2 Data organisation

Every *instance* of each subject's data were read. An *instance* is a set of analyses results, calculated from EEG of a subject performing a task. Because the EEG was epoched, there could be multiple instances per task. Each instance was a vector comprised of 623 values, as follows

- 16 upper-triangular 7x7 matrices for regional MI
- 46 values of adjacent electrode MI
- 128 values of entropy
- 1 value of correlation dimension

Each subject's *instances* were assembled into a large matrix, where a single column represented a single unique *instance* of a particular subject, and a row represented a particular variable across all instances of all subjects.

Missing data were referred to as a NaN (not a number). Rows with more than 10% NaNs were discarded, and other NaNs were replaced with the mean of that row (taken across *all* subjects), and thus contributed no innovation to the classification process.

This process resulted in a large matrix of 623 variables describing about 800 instances across 23 subjects. The matrix was stored in single, rather than double, precision to minimise memory requirements.

³A short segment of EEG.

11.3.3 Classification

Classification was performed using a memory-less feed-forward neural network. The algorithm iterated through the subjects, quarantining their data successively⁴. For each iteration, 10% of the remaining subjects' instances were randomly selected for use as validation data⁵, and the remaining instances were used as training data. Networks had a hidden logsig layer of 15 neurons, and were trained for a maximum of 3000 epochs on the other 90% of the un-quarantined data. The trained neural network was then used to classify each instance of the quarantined subject. Networks were trained with a gradient descent with momentum algorithm⁶, and the validation data were used to control early stopping, which was employed to help prevent network-overtraining⁷.

The process of quarantining, initialising, training and testing a network was repeated for each subject. In this way, we proceeded as though we had a data-base of 23 known subjects, and tested a 24th unknown subject, and then repeated this for every subject – a process known as leave-one-out training and testing.

A summary of the method is shown in figure 11.1.

⁴For each iteration, all instances of a subject's data were quarantined simultaneously.

⁵The use of separate validation and training data means that while the network is being trained, it examines its performance against the validation data. As the network becomes more trained, it can tend to become *over-trained* (too specific to the train data). When this occurs, the performance of the network using validation data will begin to decrease, and training will cease.

⁶The addition of momentum to the gradient descent algorithm helps prevent the network remaining in a local minimum.

⁷Network over-training occurs when the network becomes specific to the training data and lacks generalisability. The use of validation data allow an assessment of network generalisability (during training) that monitors the network's ability to generalise to novel data.

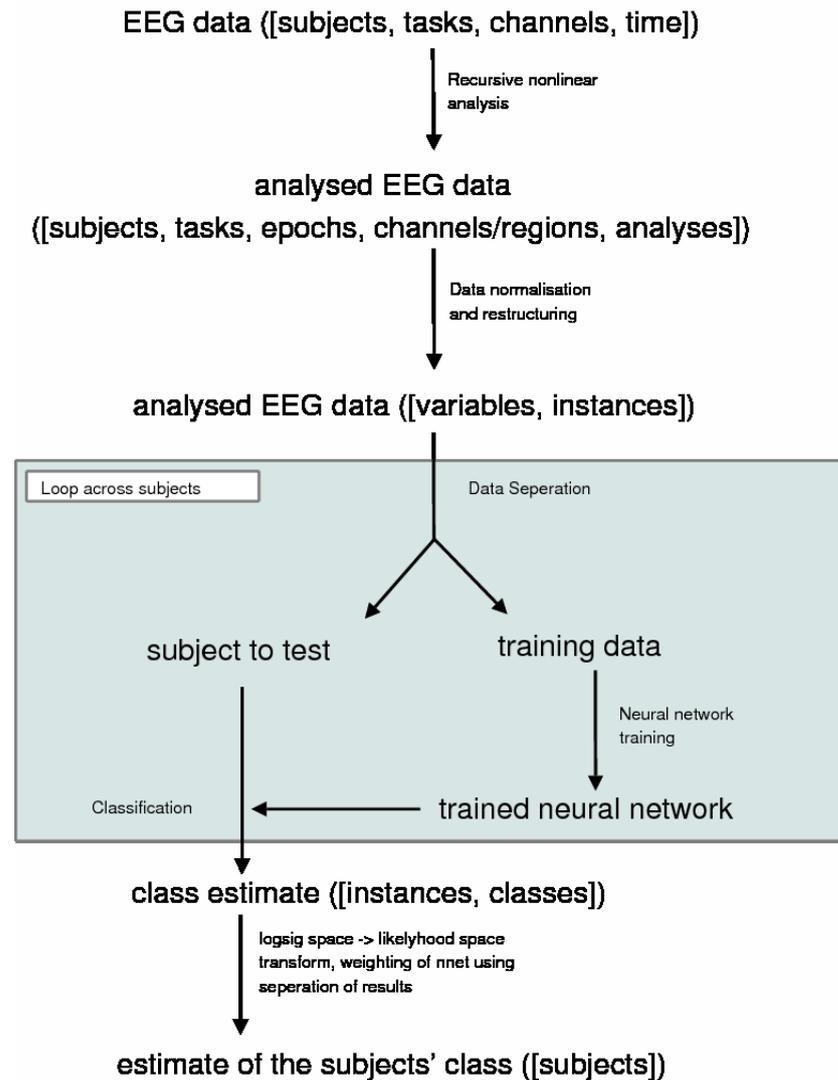


Figure 11.1: Summary of analysis of EEG leading to an estimate of the subjects' class

This figure shows the processes performed on the EEG to classify the subjects. In summary, the EEG data are divided into tasks, and then further divided into epochs – each epoch is analysed separately by each analysis. The resultant data are normalised across subjects and within each analysis, and organised into a matrix of dimensions [variables, instances], where each instance contains a analysis results from an epoch for a given task from a given subject.

We now iterate through the subjects. For each subject, we quarantine their data, train a neural network on the remaining data, and then use that trained neural network to classify the quarantined subject's data. The results of the classification are then mapped to likelihood space to allow averaging within each subject (across instances) of the results weighted by the estimated reliability (which is simply the average difference between the two classes – a large difference indicates that the neural network was more confident about its results).

11.4 Results

Results data were formed by averaging across the chosen instances⁸ and are shown in an XY plot (showing the network’s assessment of each subject). The same data were reduced to one dimension by subtracting the control likelihood estimate from the PGE likelihood estimate, and were displayed as a receiver operating characteristic (ROC⁹), showing classification performance as the decision threshold is changed.

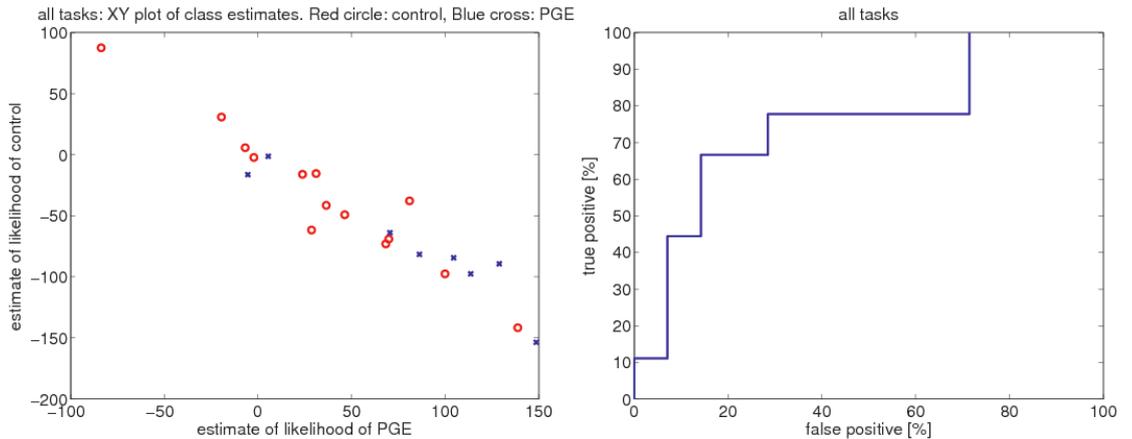


Figure 11.2: Subject classification based on all tasks. XY plot (left) and ROC plot (right)

The XY plot shows the network’s estimation of each subject’s likelihood of being in each class. Blue circles represent control subjects, and red crosses represent PGE. The ROC plot shows how true- and false-positives vary as the decision threshold is changed. The performance of the system is measured by the area under the ROC curve – perfect performance would go straight up the y-axis, and across the top of the graph.

Figure 11.2 shows the classification results for each subject when averaged across all task-instances. There were several variations to the method discussed above, all of which produced slightly different results.

⁸Recall that instances were associated with tasks. We were screening to only use certain tasks to form a decision (as described in the figure title).

⁹An ROC plot shows the relationship between true-positives and false-positives as the threshold for identification is changed. This is an atypical method of showing neural network results. Standard machine-learning protocol involves the generation of a confusion matrix (true- and false-positive, true- and false-negatives) and several statistics that describe classification performance. However, these require a classification cut-off value – a choice in the trade-off between true- and false- positives. The ROC plot shows multiple cut-off choices, and is therefore more informative.

11.4.1 Task-based breakdown of classification

A task-based classification breakdown was performed to examine the relative performance of a subset of the tasks.

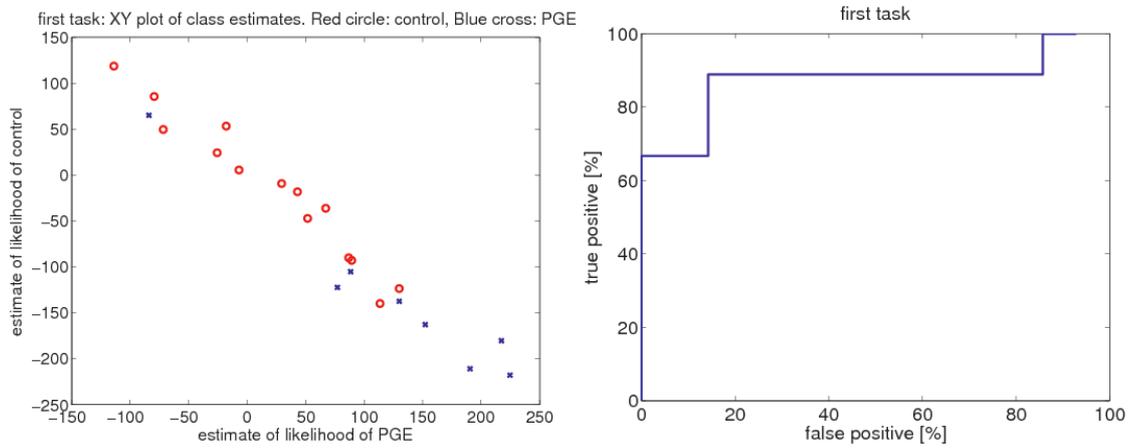


Figure 11.3: Subject classification based on eyesclosed (first) task only, XY plot (left) and ROC plot (right)

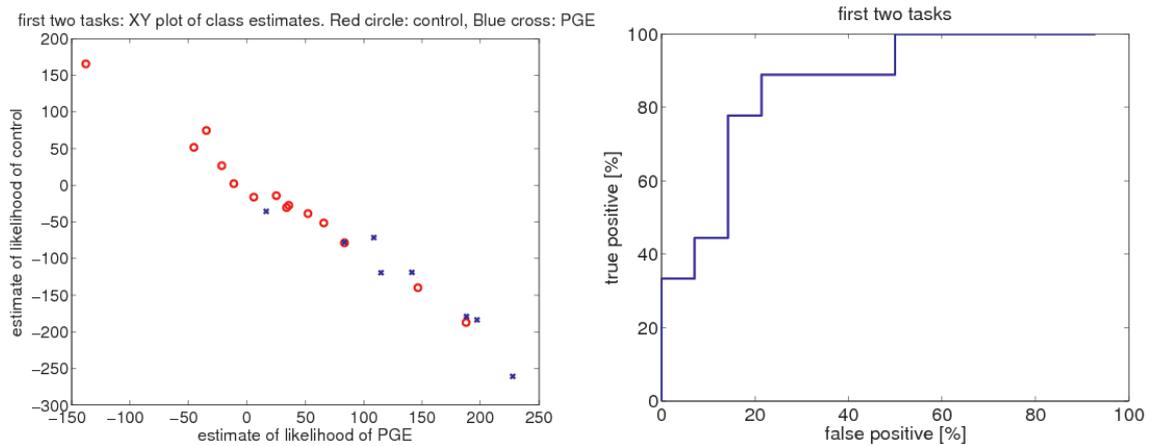


Figure 11.4: Subject classification based on eyesclosed and eyesopen tasks only, XY plot (left) and ROC plot (right)

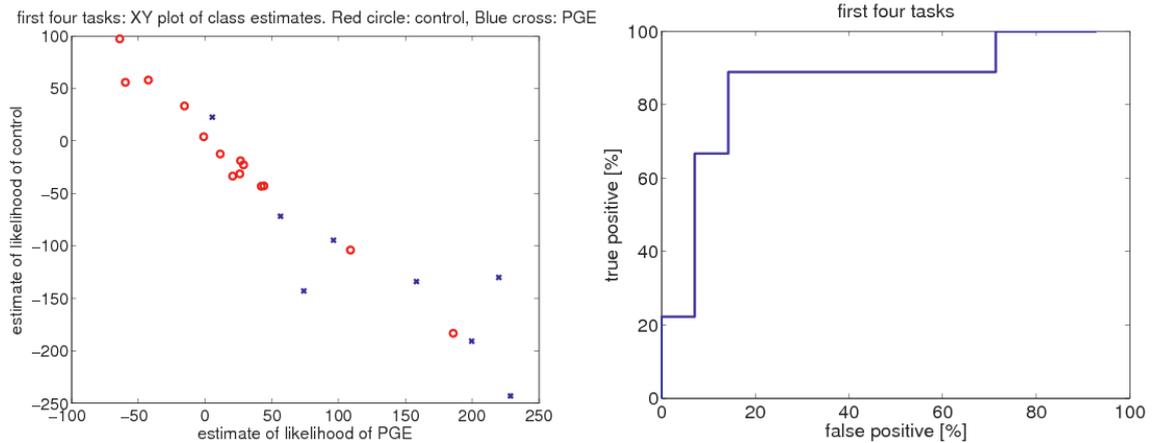


Figure 11.5: Subject classification based on first four tasks, XY plot (left) and ROC plot (right)

Figures 11.3, 11.4 and 11.5 show a good ability to classify subjects into their correct group. Use of the first task only confers the best results, and the results become worse as more tasks' data are used in the classification. Using these task subsets yields better classification than using all tasks.

11.4.2 Variation: Subject-based instances

Another network arrangement was attempted. Instead of organising data into task-instances, they were organised into subject-instances. This resulted in a matrix of size $[analyses * tasks, subjects * instances]$. Because some tasks contained more EEG epochs than others¹⁰, tasks with fewer epochs were replicated, within a subject, so that all tasks were equal in size. This produced a matrix of size $[17000, 2000]$ where each subject contained approximately 80 instances. Because there were so many data, a linear statistic was developed to reduce the number of variables (described in the Appendix, section A.3). The remaining data were used to train a feed-forward neural network in the same way as previously described. For clarity, I refer to this method (with subject-based instances) as the *subject-instance* classifier, and the earlier method as the *task-instance* classifier.

¹⁰Recall that successive epochs of EEG were used to form the instances.

Figure 11.6 shows the results that were yielded by this process when selecting the best 1000 variables. This number can be varied, and we tested several values, finding that while there was a variation in performance with the number of variables used, performance was reasonably consistent.

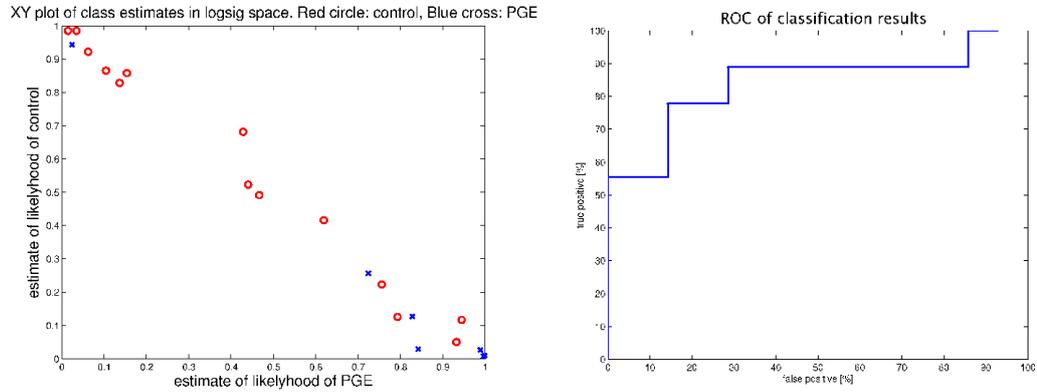


Figure 11.6: 200 best variables

The XY plot (left) shows two groups of points, corresponding to the estimated class of the subjects. The points are coloured according to their actual class. The X axis shows the estimated likelihood that the given subject is PGE, and the Y axis shows the estimated likelihood that the subject is control. These results were produced from a network that had 200 variables presented to it. The ROC plot (right) show the same data. It represents the performance of the classification system as the cut-off point for positive identification (of PGE) is shifted.

11.5 Discussion

These results show an ability to diagnose PGE in an unknown subject, and suggest an important application, which is the potential use of this technique as a diagnostic tool. Such a tool could provide a quantitative description of the subject's likelihood of manifesting PGE epilepsy symptoms (such as seizure) from a recording of non-ictal EEG. Such a tool does not currently exist, and would be very useful.

The results show sensitivity to the choice of data for training and testing. Before application in a clinical environment, this would need to be better understood.

11.5.1 Which variables are useful for classification?

We also conducted an examination of which variables were useful in classifying the subjects.

11.5.1.1 Laplacian analysis

The analysis of Laplacian EEG was never useful for discriminating between control and PGE subjects. The only analysis useful for classifying subjects was MI. This result implies that the levels of synchronicity between different areas of the brain are affected by PGE – *even* when the PGE subject is not exhibiting any symptoms.

11.5.1.2 Useful Tasks

We examined which tasks performed best in the individual task-based analysis. The useful tasks were eyesclosed, eyesopen, fingertapping and auditory discrimination.

Three of these tasks were the the first tasks performed by subjects, and were always performed in the same order. Later tasks were undertaken in random (across subjects) order, which would have made classification more difficult (because of differences in the fatigue level of different subjects when performing the same task). Based on this I expect that, were this classification algorithm to be applied to data that were consistent throughout the entire experimnt, the classification accuracy would be improved – simply because of the the greater number of useful data.

11.5.2 Implications for the application of this method in a clinical setting

In this analysis, there were approximately equal numbers of subjects examined in the control class and the PGE class. In a clinical setting, that would not be

the case – there would be many more non-PGE people tested. The identification of a small minority in a population is a difficult problem, in general the lower the incidence of the sought characteristic, the greater the rate of false-positives. It is likely that this procedure could suffer from problems similar to this, were it to be implemented in a clinical environment. Such considerations, however, are outside the scope of this thesis.

I have exercised great care in ensuring that the development of the algorithm has not been unduly influenced by the data being analysed. This is always a trap with this type of procedure – one creates an algorithm well-tuned to the specific data, but one lacking in generalisability. This is an inherent limitation of this approach to *post-hoc* development and classification. To ensure that the results seen here are actual, the algorithm (developed and tuned on these data) must then be applied to new data, to see if its ability to diagnose PGE remains, and to what degree.

11.5.3 Task performance variations

The performance of the *task-instance* classifier varies with which task-instances are presented to the network. In general, earlier tasks (eyesclosed, eyesopen, fingertapping) showed better classification accuracy than later tasks (AVLT, maze, visual rotation). The first few tasks were always presented in the same order, whereas the task presentation order of later tasks varied between subjects. Also, the useful tasks tended to be those that were unaffected by the subject's performance – tasks like maze, and visual rotation require interaction with the subject, which may increase variability.

11.5.4 Future directions

This procedure needs to be applied to more people, in a clinical setting, to validate the results.

A similar procedure could be used to investigate the relationship between

various neurological pathologies. We believe (unpublished) there to be a relationship between epilepsy and migraine. A three classification system between control, PGE and migraine, and control, PGE and schizophrenia might allow dendrogrammatic associations between brain pathologies to be established (e.g. PGE is similar to migraine, but different from schizophrenia, which is similar to Parkinson's disease, etc).

11.6 Conclusion

We have successfully demonstrated a method by which we can form a diagnosis for PGE in previously unknown subjects. The extent to which the observed good accuracy is replicated in a clinical setting should be examined by further testing.

Chapter 12

Experiment: Classification of mental task in paralysed human subject

EEG data were recorded from a healthy human subject, in unparalysed and paralysed states, while they performed a series of mental tasks. The data were stored, and were analysed offline using entropy, correlation dimension and mutual information algorithms (the *analyses*).

These data were used to train, via a leave-one-out method, a memoryless feed-forward neural network to classify the data from various tasks (mental task identification). The networks had 15 logsig nodes in the hidden layer, and 2 logsig nodes in the output layer. It was found that the network was unable to generalise using paralysed data, but that classification was improved when using unparalysed data – suggesting that the addition of muscle artifact to the EEG assisted task classification.

12.1 Hypothesis

It is possible to identify the task being performed by a paralysed human subject by analysis of EEG and subsequent automatic classification techniques.

12.2 Introduction

The ability to determine the “brain state” from an analysis of EEG would show that the task differences (chapter 10) can be used to classify the task being performed. Reliable task classification would demonstrate that classification for a brain-computer interface (BCI) is not necessarily reliant on muscle artifact in the EEG.

Aside from this, the ability to differentiate between tasks will provide a validation that the procedure (nonlinear analysis of EEG and subsequent classification therefrom) is able to discriminate between different brain states, and is not merely sensitive to more general patterns of brain operation (ie. differences between control and PGE).

12.3 Method

12.3.1 Data analysis

The data were collected as discussed in section 6.3. The EEG data in each task were divided into epochs of EEG for analysis, where the number of epochs depended on the number of data recorded in the subject’s task. For the purpose of analysis, each epoch was treated as a separate segment of EEG. This EEG was then analysed with the analyses discussed in chapter 7 (the *analyses*), and the results were stored.

12.3.2 Data collation

Because of the very limited number of experiments and limited time, only data from one subject was examined. This meant that any developed classifier would be subject-specific, reducing generalisability in a practical situation¹. This is a less ambitious target than training a generic task-classifier, but success

¹Having said this, many BCI systems are designed to be customised to a subject.

would show that task identification and classification is possible from EEG data.

Figure 12.1 shows the difference in power between paralysed and unparalysed subjects at rest. From this, we conclude that contamination of EEG by muscle is a much more pervasive and important problem than is commonly recognised – by 30 Hz, the difference is nearly an order of magnitude. For this reason, it seems plausible that much of the classification performed by many BCI systems is actually based on EMG (muscle artifact) rather than EEG.

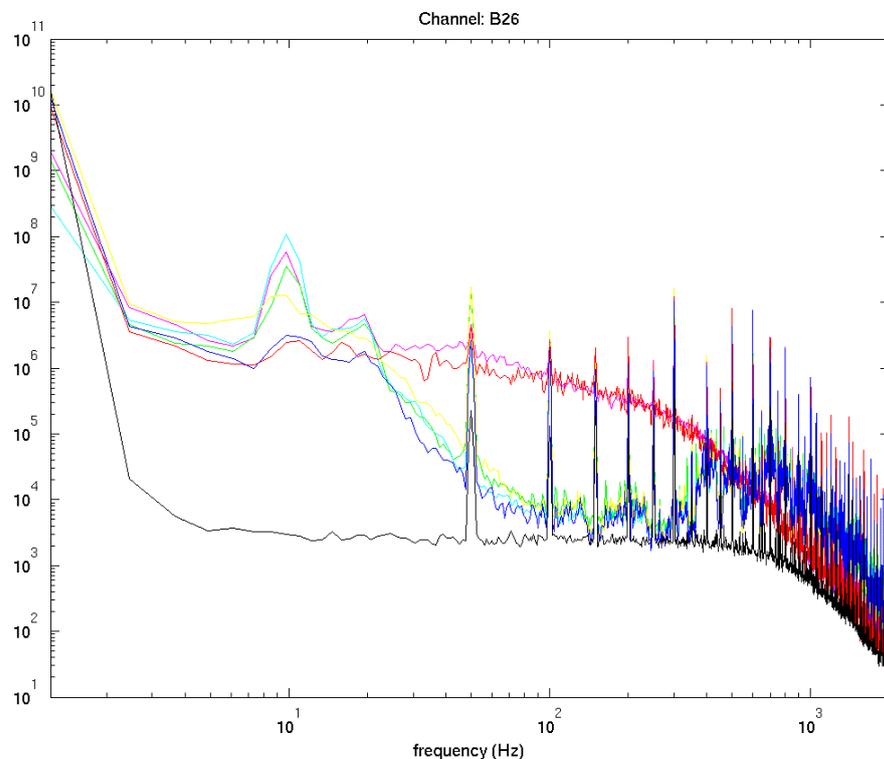


Figure 12.1: Graph of EEG spectral energy of paralysed and unparalysed subjects. This graph shows the large difference in gamma power between EEG recorded from paralysed (green, yellow and blue) and unparalysed (red and pink) subjects. The black trace was recorded from EEG electrodes in a bucket of water. From this, we conclude that muscle artifact is a much larger contributor to EEG than is commonly recognised [83]. The figure shows a clear 50 Hz (and harmonics) artifact, as well as a decrease in the noise floor above approximately 500 Hz due to the anti-aliasing filter (cutoff at 1 kHz)

12.3.3 Assembly of the training matrix

The saved analysed data were collated into a matrix, where each column represented the analyses from a single unique instance of the subject performing a particular task in a particular state (paralysis or pre-paralysis). Missing data were referred to as a NaN (not a number). Rows with more than 10% NaNs were discarded, and other NaNs were replaced with the mean of that row (taken across *all* subjects), and thus contributed no innovation to the classification process.

12.3.4 Classification

This process produced a matrix [variables, task-instances] from a recording of a single subject. This matrix was used to train a neural network to classify a particular *instance* as a *task*.

We chose 4 tasks to train and four separate but related tasks to classify (table 12.1). This ensured the separation of training and testing data. Also, a smaller number of classes tends to result in better classification (figure 4.2).

Table 12.1: Training tasks (paralysed data)

1. Baseline eyes-closed
2. Odd-ball high tones
3. Finger-tapping left
4. Auditory discrimination incorrect

We then attempted to use the trained network to identify the testing tasks (table 12.2)

Table 12.2: Testing tasks (paralysed data)

1. Final baseline eyes-closed
2. Odd-ball low tones
3. Finger-tapping right
4. Auditory discrimination correct

The tasks in the training and testing sets are not exactly the same – in particular the odd-ball response to high tones is demonstrably different to the low tone response (as demonstrated by evoked response analysis², not shown), however we considered that there should be sufficient similarity to allow classification.

We trained and tested the network in two ways. Firstly, we trained the network on the tasks in table 12.1, and then tested it on the same tasks. We then tested the same network on the tasks in table 12.2. This process was repeated 15 times. The results were transformed to likelihood space, and the network results averaged. The repetitions of the neural network simulation were selected-for based upon the confidence of the classification by multiplication by the variance of the class estimates³.

12.4 Results

The results for testing the network on the training data are shown in table 12.3. We achieved excellent classification when the test data were familiar to the network – demonstrating that the setup, training and testing stages of the network were correctly applied.

The results of testing the network on the previously unseen test data are shown in table 12.4, and show poor classification.

²Evoked response potential (ERP) is a mean of many EEG traces, synchronised by an event (in this case, the presentation of the tone). The mean shows the general EEG behaviour, without any added noise.

³A large variance implied that one class was favoured over the others.

Chosen tasks	1	2	3	4
Baseline eyes-closed	1	0.23	0	0.26
Odd-ball high tones	0.35	1	0.06	0
Finger-tapping left	0	0.27	1	0.17
Auditory discrimination incorrect	0.33	0.02	0	1

Table 12.3: Results of task classification attempt on familiar paralysed data

This table shows the neural network task estimates when tested on the (familiar) training data.

In this table (and tables 12.4 and 12.5), the numbers correspond to the network’s assessment of the likelihood that data from a particular task corresponds to each of the possibilities. For instance, the datum in cell $\{1,2\} = 0.23$, is the average network’s estimate of the likelihood that data from the baseline eyes-closed task actually came from the odd-ball. To increase readability, the results were scaled to the interval $[0\ 1]$.

Chosen tasks	1	2	3	4
Baseline eyes-closed	0.78	1	0.83	0
Odd-ball high tones	0.10	0.19	0	1
Finger-tapping left	1	0	0.56	0.06
Auditory discrimination incorrect	0.94	0.49	1	0

Table 12.4: Results of task classification attempt on **un**familiar paralysed data

This table shows the neural network task estimates when tested on the (previously unseen) test data.

Chosen tasks	1	2	3	4
Baseline eyes-closed	0	0.58	0.20	1
Odd-ball high tones	1	0.72	0.68	0
Finger-tapping left	1	0.38	0.87	0
Auditory discrimination incorrect	0.50	0	0.28	1

Table 12.5: Results of classifying attempt on **un**familiar **un**paralysed data

To examine whether the paralysis significantly altered the results of the classification, a short experiment was run to examine the effectiveness of the classification of unparalysed data.

12.5 Discussion

As can be seen, the analyses are poor at identifying the mental task from unfamiliar data, although performance was better when classifying novel analysed unparalysed EEG.

12.5.1 Limitations in the data

Because this experiment was performed on a single subject only, the developed neural network was tailored to a particular person. This lack of training data means that the feature space presented to the network was limited, which may have contributed to the poor performance (lack of training data diversity tends to result in a trained network that generalizes poorly).

12.5.2 Statistical differences between the tasks

In chapter 10, we showed that there are differences between the mean values of the tasks (figures 10.1 and 10.2) – note, however, that the differences we showed were in the SEM, which does not imply that a given sample is easily classifiable.

The work in this chapter has been unable to exploit any such differences to perform task classification on paralysed data. While not conclusively shown, it may be that while there are differences between the tasks' means, there are large variances in the populations so that an attempt to examine a single subject will result in problems because of large variation in the results. Repeating the process employed in this chapter, but with more experimental data, will allow us to draw firm conclusions and also attempt the construction of a generalised neural network task-classifier.

12.5.3 EMG artifact in EEG

The performance of the classifier is substantially better for the unparalysed data. This implies that the network is using the EMG present in the EEG to assist in the classification. Since (to our knowledge) *all* BCI research is being conducted on unparalysed data (human paralysis experiments are *very* rare), it is likely that they are using muscle artifact to assist in the classification process. This is not a problem in and of itself, however it should be recognised.

A major application of a BCI is as an augmentative device for a disabled person⁴. If this person has very limited movement (as could be the case for stroke resulting in paralysis) or highly uncontrolled abnormal movement (certain types of degenerative nervous disease) then the classifier may not be able to derive meaning from the muscle artifact. This means that it may be more difficult to build a functional BCI for the people who need it most.

12.6 Conclusion

The work in this chapter examined whether the analyses were useful in assessing the *type* of mental activity, and if such classification was affected by absence of EMG. We found that we could not reliably identify the task in a paralysed subject, leading to the conclusion that there is not a consistent distinct identifiable feature, that is detected by the analyses, when there is no muscle artifact in the EEG. Classification accuracy was improved in the unparalysed subject, suggesting that the analyses, were relatively insensitive to changes in brain state due to mental task and that EMG artifact can improve classification. Thus, they are unlikely to be useful for the development of a BCI, particularly for people who lack controlled muscle mobility.

⁴A person who lacks the muscle control to use a computer can then use a computer to assist them in their day to day life.

Chapter 13

Conclusions

This chapter summarises the conclusions and contributions of this thesis, and provides suggestions for future work.

13.1 Thesis summary

This thesis began with an examination of the literature – both in terms of signal processing theory and practice, as well as a description of epilepsy and its effect on the brain. I also discussed the *meaning* of signal processing as applied to the brain, as well as several methods of classification. Following this, I gave some examples of nonlinear signal processing in the field of neurological research, the validity of such research (some of the caveats and pitfalls) and some of the conclusions that were able to be drawn.

As a precursor to the chapters regarding my experiments, I described the sources of the data that I used, and the manner in which the data-collection experiments were conducted. This chapter also detailed my involvement in the data-collection experiments, and the work that needed to be performed on the data prior to analysis. I also described, in detail, the analyses that I applied for my EEG analysis, and the manner of their application. It is important to recognise the limitations of these analyses, and I showed how an analysis of surrogate data using said analyses aided in the interpretation of the results.

The remaining chapters all detailed the experimental work that I performed. The first, chapter 8, I refer to as an exploratory experiment in the analysis of rat EEG. I use the word *exploratory* because the experiment does not have a prescribed plan, but rather was conducted in parallel with the implementation and testing of the analyses. Despite this, I was able to observe that there were changes in EEG that appeared coincident with changes in measured brain impedance, suggesting that it may be possible to estimate impedance from EEG using a neural network.

Estimating impedance from EEG was attempted in chapter 9. The network was trained on a series of historical EEG and impedance values, from several rat experiments. A similar arrangement of historical values of EEG and the analyses (from a previously unseen experiment) was then presented to the network for simulation, to test its ability to estimate impedance. This attempt failed, but it seems likely that this was in part due to the limited number of data that were presented to the network (this limitation was mainly imposed by the number of the data needed from each experiment, and the computational load the training and simulation required).

The next experiment (chapter 10) examined human data. It sought statistical differences between mental tasks in the control class. We found that there were such differences, and these were apparent in the MI analysis – both the adjacent electrode and regional analysis. When we examined these task variations, in PGE subjects relative to control subjects, we found that there were class-dependent differences. When measured in units of standard deviation, the differences were small, but because there were many variables it seemed plausible that a PGE diagnosis in a single subject might be possible – if there was sufficient independence between these variables. This was attempted in chapter 11, using a neural network arrangement similar to that in chapter 9. We found that it was indeed possible to make a diagnosis of PGE in a previously unseen subject, using a neural network trained on a *training set* of

apparently normal EEG. Although there was some variation in performance, due to variations in the implementation, overall performance was reasonable.

The final experiment, in chapter 12, examined data from the human paralysis experiment. It used a neural network to attempt to identify the task being performed by a subject from the analysis of EEG data. This experiment was limited by the small number of data available at the time of writing (few such paralysis experiments have been performed). We found that, in a non-predictive test of the network (where the test data are included in the training data) performance was quite good. However, when a predictive test was conducted (keeping the test and train data separate – as was done in chapter 11 for the diagnosis of PGE) the ability to identify the task was very poor. This was somewhat improved by using un-paralysed data, suggesting that the network was using variations due to EMG artifact - an observation that is important for the development of a BCI for people lacking controlled motor function

13.2 Comments on the analyses

From the analyses performed in these experiments, some conclusions can be drawn about the reliability of the implemented algorithms. The algorithm that was most consistently useful, and conferred the most analytical and predictive power, was mutual information. When applied to human subjects, the method of estimating regional MI (by the calculation of MI between electrodes in the two regions and then averaging across the regions to yield an estimate of the average MI between two regions) was very powerful when used for the detection of PGE - it was the quantifiers based upon this that were most commonly useful.

The rat experiments showed that entropy is a useful analysis, despite the surrogate analysis demonstrating that it is sensitive to low-order statistical changes in the data.

13.3 Future work

13.3.1 Diagnosis of PGE from non-ictal EEG

This is perhaps the most exciting and potentially useful finding of this thesis. The development and implementation of this method could yield a diagnostic tool that could be extremely useful in a clinical environment. The most immediate requirement for developing this technique is to apply it to more people. I expect that doing this will make the diagnostic assessment more reliable, because the classifier will have more examples of the classes. Further investigation should allow us to target the analyses to the tasks and electrodes that are useful for classification. From the analysis performed so far, it appears that examining inter-region MI is the best method by which to determine the presence or absence of PGE – this knowledge means that our approach to classification can be more specific. By focusing only on certain tasks, people need not be subjected to such a drawn-out test. Thus an increase in specificity will allow us to make a much simpler, more rapid, and probably more reliable diagnosis.

An obvious generalisation of this method is whether it can be used to detect other neurological pathologies. Perhaps the first to be examined should be migraine, because of its known links with epilepsy. One could also use a technique to assess the similarities between these pathologies - for example, if it were difficult to distinguish between migraine and PGE, yet simple to distinguish between PGE and schizophrenia, then we could reasonably conclude that PGE and migraine were more closely related than PGE and schizophrenia.

13.3.2 Analysis of rat data

This work has shown the impedance recording provides data that are probably independent of the EEG data. This is concluded because we were unable to form a reliable model for the impedance data using the EEG data. This

doesn't mean that the analysis of EEG is not useful, merely that it does not yield data that are useful in estimating the impedance. It would be interesting to combine the impedance data, and the EEG-derived data, and attempt to make a prediction of the onset of seizure in the rat epilepsy models. This could be achieved by creating a new binary data set (called *seizure_flag*) that was mostly zeros, but was ones where there was a seizure. This could be used with a neural network configuration similar to that implemented in chapter 9. Instead of taking historical values of EEG and analyses as training data, and impedance as goal data, we would create a set of historical EEG, the analyses, *and* the impedance data as the training data, and use a future value of *seizure_flag* as the goal data. The amount of anticipation between the most recent train data, and the goal data would represent the duration of reliable anticipation of seizure (a time of zero would mean seizure recognition, whereas 10 seconds would imply that the system could anticipate a seizure by 10 seconds).

Seizure prediction is in some ways the holy grail of epilepsy signal processing work, but none of the attempts to predict the onset of seizures have used impedance recordings as a component signal to the neural network. Because we have shown the data in the impedance channel to be complementary to the data that were derived from the EEG, this could be a strong aid to attempts to predict the onset of seizure.

I suggest that this is very much worth attempting.

13.3.3 Classification of mental task in paralysed person

There is much work to be done to extend this. This work needs to be applied to several paralysed subjects, and an attempt made to generalise the task data across subjects – thus forming a generic classifier. Such a tool would be much more useful than the current subject-specific classifier that has been developed – it would likely also be much more robust.

13.4 Main contributions

- Implementation, vectorisation and integration of several nonlinear analyses with the laboratory recording arrangement (chapter 7)
- Implementation of an automatic un-wrapper for impedance data, and integration of impedance hardware into the laboratory setup (section 6.1.6)
- Creation of a GUI to facilitate easy visualisation of EEG, nonlinear analyses and impedance from experimental data (figure 8.1)
- Demonstration of the validity of recording EEG *and* impedance data from rats – as impedance cannot currently be estimated from EEG (figure 9.3)
- Development of a framework from which an attempt at prediction of seizure in a rat experiment can be made (section 9.5)
- Demonstration of detectable task and group (control vs PGE) differences in human EEG (figures 10.1 and 10.2)
- Development of a promising diagnostic tool for primary generalised epilepsy (PGE), able to operate on non-ictal data (section 11.5)
- Development of a framework from which to perform additional classification of paralysed human data from future experiments (chapter 12)
- Further demonstration that muscle artifact is prevalent in unparalysed human EEG. Also that such artifact aids task recognition – an observation that may be relevant to brain-computer-interface (BCI) developers (section 12.5.3)

13.5 Conclusion

This thesis has established that it is possible to identify *primary generalised epilepsy* from an analysis of non-ictal EEG only – a novel finding. The best accuracy of this diagnosis was 92%, although a real-world accuracy is likely to be somewhat lower. We think this has excellent potential. Further testing of this system on other data should allow the development of a robust diagnostic tool, usable in a clinical environment, for the diagnosis of PGE. Such a tool would be very useful to practising neurologists.

This thesis also suggests several directions in which further work is justified. In particular, an attempt at anticipating seizure in rats is worthwhile, as is a better examination of the task variation and detection in paralysed subjects. There is much work to be done!

Appendix A

Appendices

A.1 Appendix A: Additional Figures

A.1.1 Regional differences for entropy and correlation dimension analysis

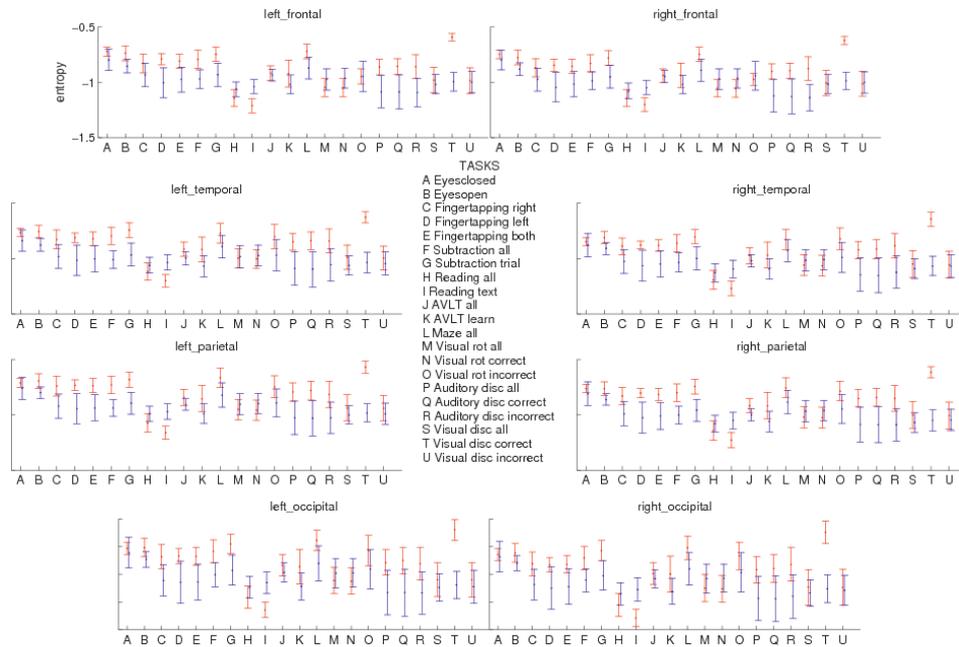


Figure A.1: Regional entropy analysis

This figure shows mean and SEM of entropy, when averaged across subjects (control: red, and PGE, blue) within tasks and brain regions. None of the task/region instances showed statistical significance between control and PGE groups.

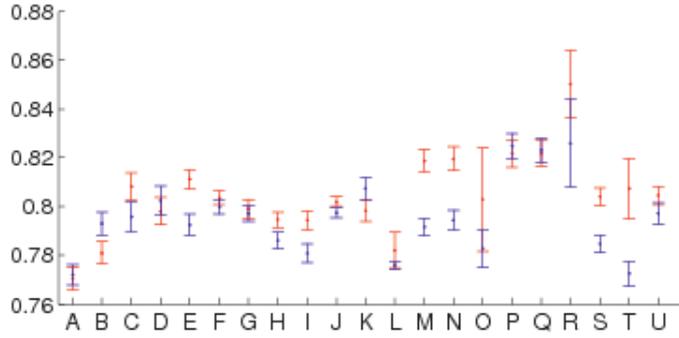


Figure A.2: Correlation dimension analysis

This figure shows the correlation dimension averaged across subjects and brain regions, none of which showed statistically significant separation between control and PGE groups.

A.1.2 Mutual information analysis by brain region, PGE vs control

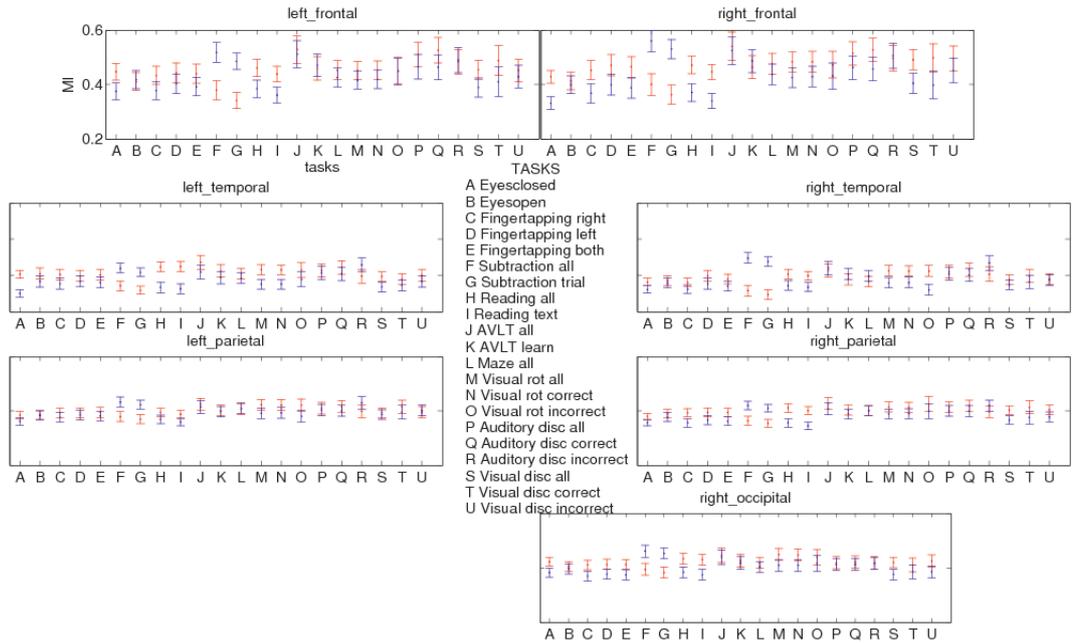


Figure A.3: Task-based mean and SEM of MI of Right Frontal vs other brain regions, PGE vs control

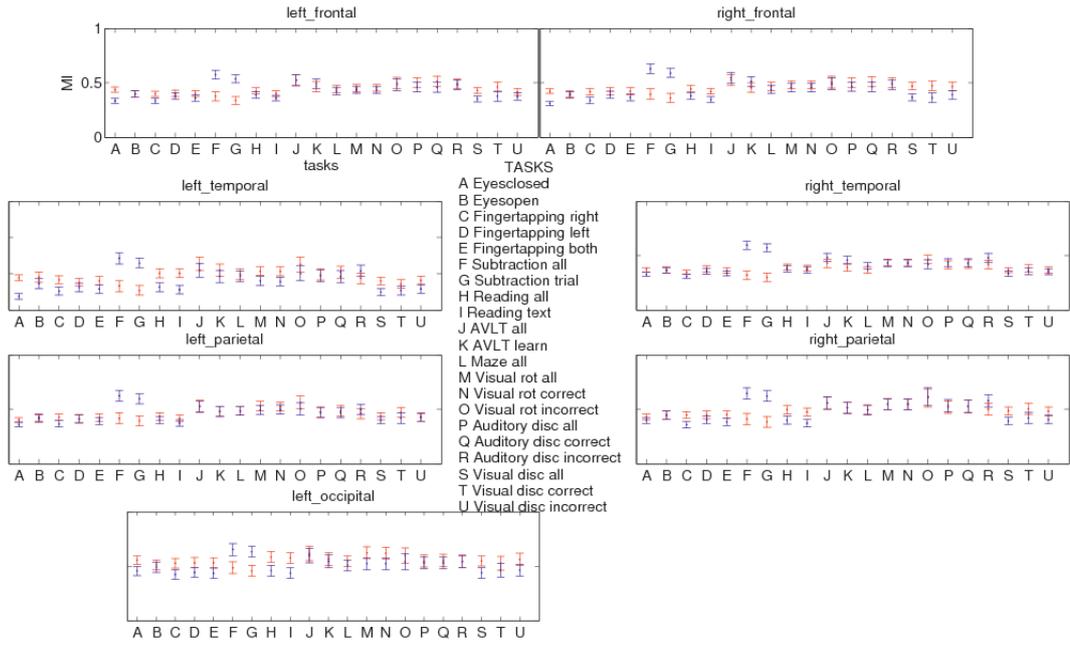


Figure A.4: Task-based mean and SEM of MI of Right Occipital vs other brain regions, PGE vs control

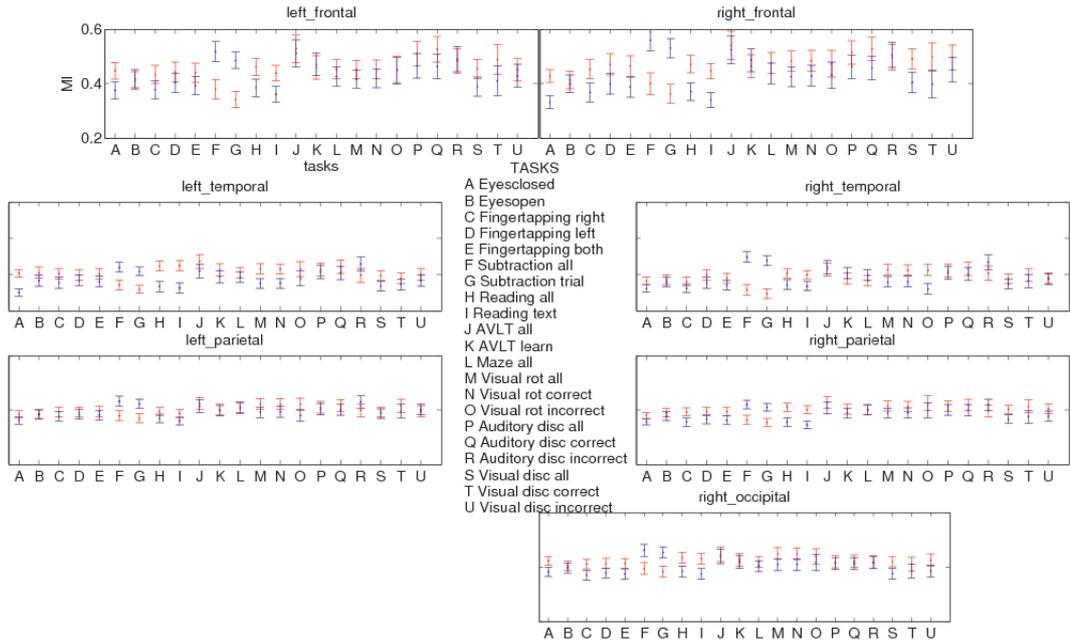


Figure A.5: Task-based mean and SEM of MI of Left Occipital vs other brain regions, PGE vs control

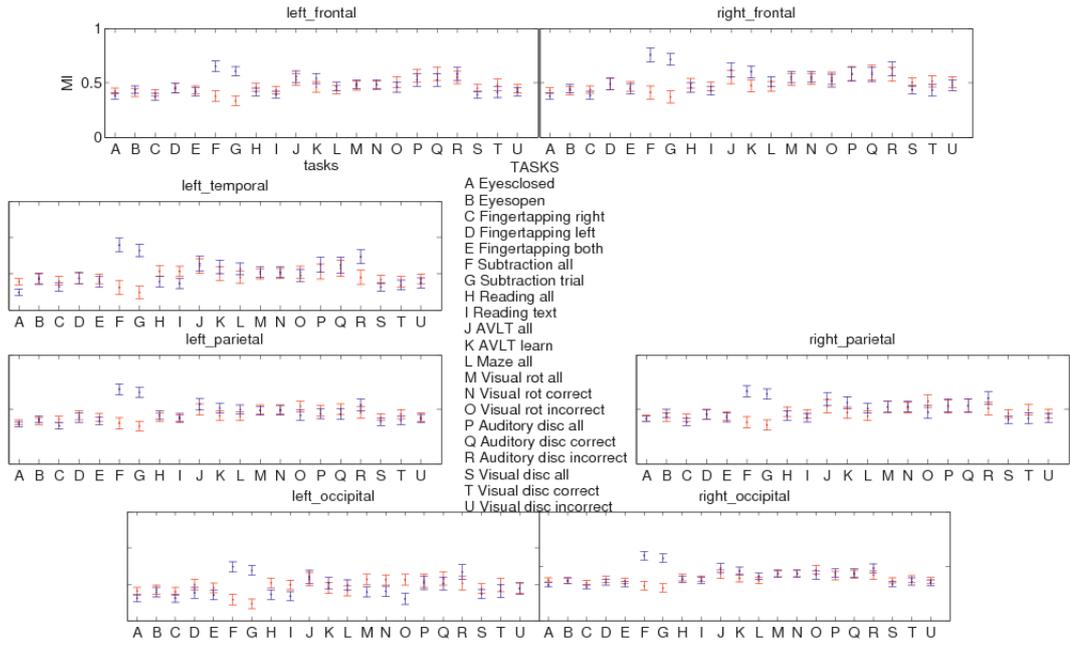


Figure A.6: Task-based mean and SEM of MI of Right Temporal vs other brain regions, PGE vs control

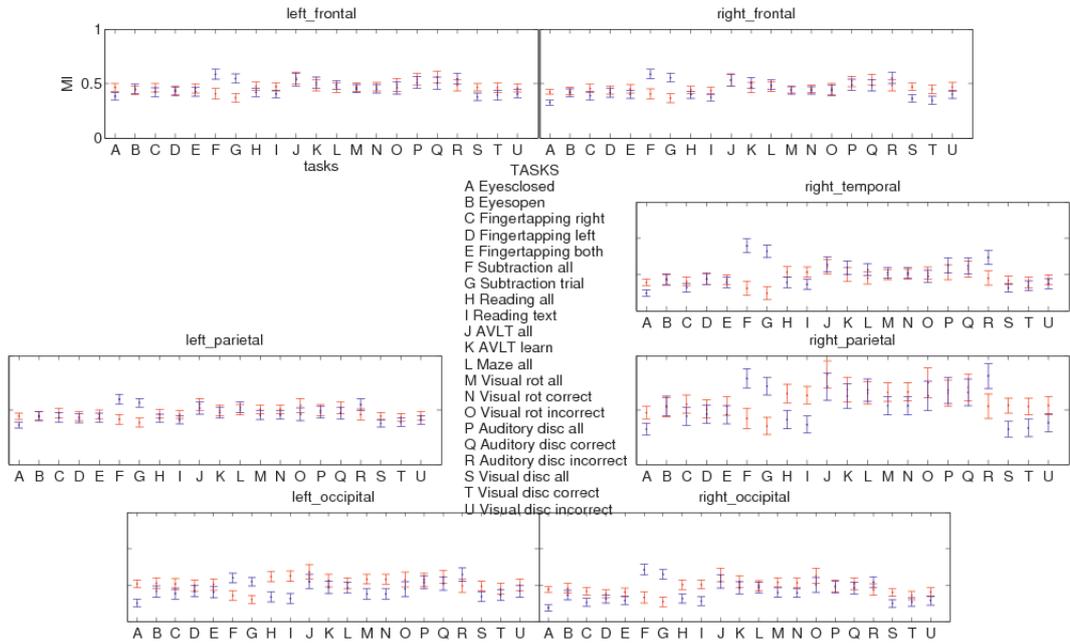


Figure A.7: Task-based mean and SEM of MI of Left Temporal vs other brain regions, PGE vs control

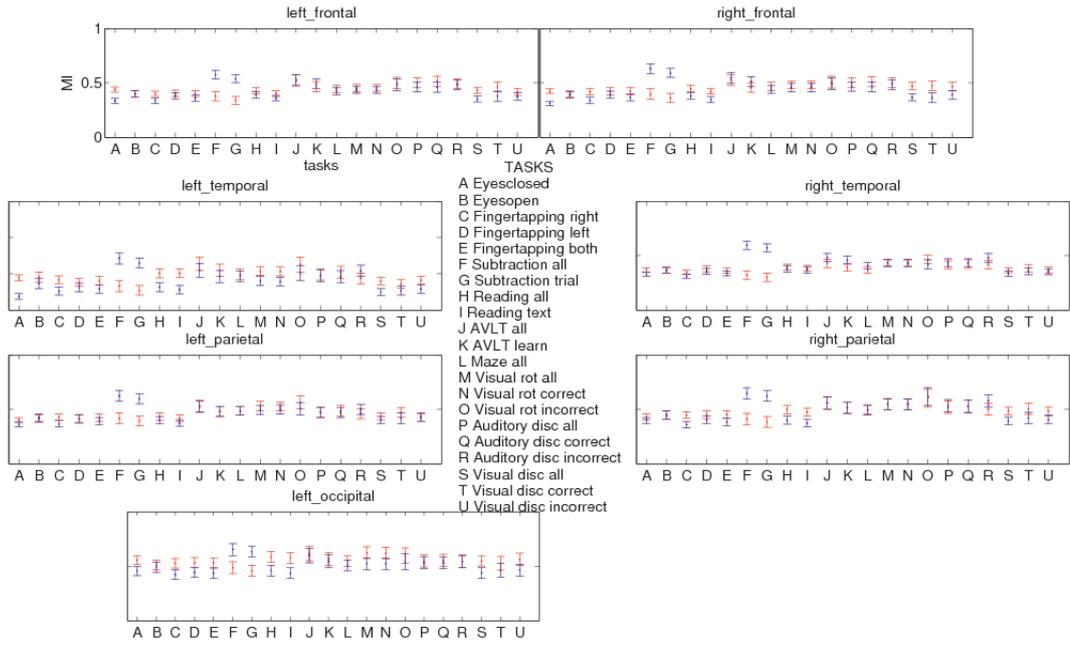


Figure A.8: Task-based mean and SEM of MI of Right Occipital vs other brain regions, PGE vs control

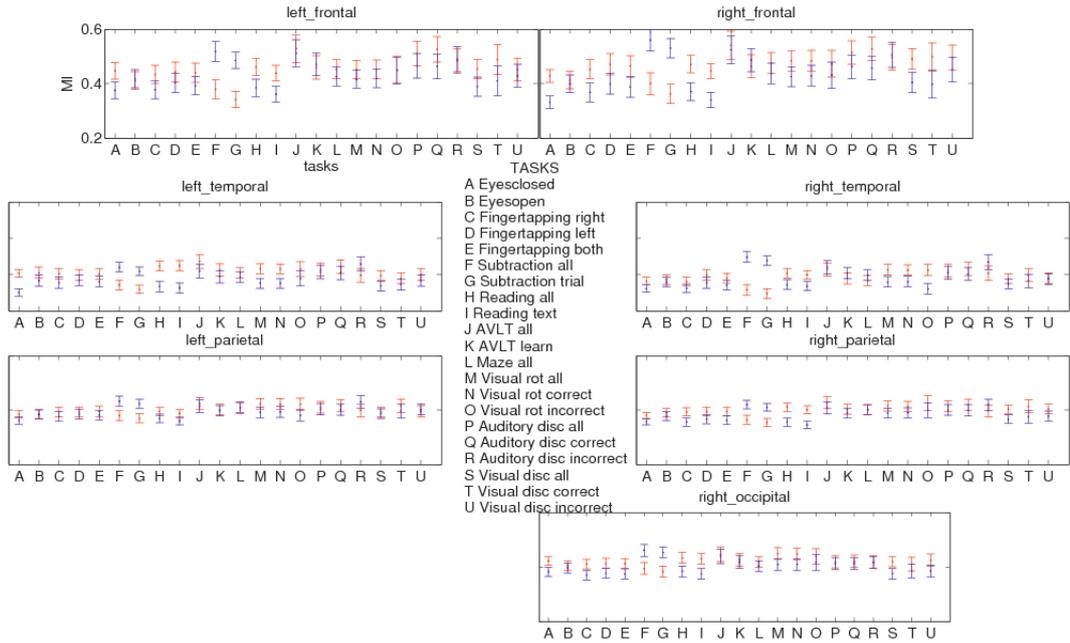


Figure A.9: Task-based mean and SEM of MI of Left Occipital vs other brain regions, PGE vs control

A.2 Appendix B: Neural network choices

When seeking a network for use in estimating impedance from EEG data (chapter 9) we considered using the series/parallel configuration. The main difference is that it uses feedback, which would mean changing the way data were presented to the network for training. During simulation, the impedance estimate is fed back into the inputs. During training, however, the *actual* impedance value would be passed to that input. This means that the network would be able to learn characteristics of the impedance (for example, the impedance doesn't change rapidly, it tends to increase throughout the experiment).

We did not proceed with this configuration for the following reasons:

1. It will tend to make the output of the neural network seem “more certain”. There will be less variation in the output, because the network will base future estimates on past estimates – this will result in a tendency for the network to “wander” from where it should be (a false sense of security).
2. The use of feedback also means that the network is more difficult to conceptualise. The employed network (a basic feed-forward network without memory) simply produces an output that is a direct result of the inputs provided. This is not the case when feedback is employed (since the output is affected by previous outputs) making understanding the response of the network more difficult.
3. Following from point two, the use of feedback would also make it difficult or impossible to establish upon what basis a network made a decision. Contrast this with the feed-forward network used, where we know exactly which historical information was provided.

For these reasons, a simple feed-forward network was chosen for all experiments.

A.3 Appendix C: Human PGE diagnosis

A.3.1 Selection of useful variables

Classification using 30000 variables is very resource-intensive. To reduce computation requirements, we wished to identify the best variables with which to classify the subjects – using a linear process (because of the number of data, it was important that this process was not computationally intensive). As discussed in section 10.4.2, class-based analysis revealed that there are variables that show a strong difference between the two classes, but there are also variables that show no difference. We wanted our classifier to retain the best variables only – those that allowed us to separate the two classes.

Seeking a simple linear-classifier to choose which variables to use for the subject classification is almost nonsensical - for if a linear classifier can choose which electrodes to use, why can't it classify the subject? The answer to this that we are hoping to use a *fast* linear process to reduce the number of data, and then use a neural network to find high-level relationships in the remaining data, and classify from that. The hope is that the data reduction will not remove these high level relationships.

An important consideration when choosing the variables is ensuring that we are not using *a priori* knowledge of the *test subject's* class, because in a true diagnostic scenario, such information would not exist – this is discussed in section A.3.3.

A subject's data were quarantined from the rest of the data. That person represented the test subject, and the remaining subjects were the known subjects. **Only** the remaining subjects' data were then used as the basis for choosing which variables to retain and which to discard – it was hoped that the decisions based upon these data would be useful in the classification of the unknown subject.

Several methods of variable selection were tested in conjunction with clas-

sification, and the method that we found to work best is as follows:

Compare each un-quarantined subject against every other un-quarantined subject. If they are in the same class, find the variables that best show this. If they are in different classes, find the variables that best show this. We use formula A.1 to estimate the distance between the means of each subjects' variables. If the subjects are in the same class, we value variables that have low values of SVF (i.e. the distance between their means is small), and if the subjects are in different classes we choose high values of SVF (so that it's easy to discriminate between the subjects).

$$SVF = \frac{diff(mean([subject1, subject2]))}{sum(var([subject1, subject2])} \quad (A.1)$$

This formula expresses the distance between the means in units of standard deviations (technically the sum of the standard deviations of the two examined subjects). It provides an estimation of the separability of the two classes. The formula is vectorised, so it finds the difference of the subjects' means, and divides this by the sum of their variances.

Iterate through the remaining un-quarantined subjects and, eachfor each iteration, produce a list of variables that are considered "good". After iterating, we find the sum of these records, so that we obtain a vector showing how many times each variable was considered "useful". This was ranked, and we retained the *best* variables.

A.3.2 How many variables should be retained?

As discussed in section 11.4.2, a linear algorithm is used to reduce the number of variables, prior to classification using a neural network. The number chosen is something of a "magic number", but no matter how it was approached

(whether by threshold or absolute number) the choice would always be somewhat arbitrary. I found that while the chosen number of variables did affect classification results, it was not significant. To conclude this, the analysis was run using different numbers of variables (60, 150, 300, 500, 1000). I found that this parameter did affect the results of the classification process – there was always an ability to recognise PGE subjects, but there was variance in the performance. Of course, as computing power and memory increase, I recommend the inclusion of all variables in network training.

A.3.3 The importance of being earnest about data separation

It is absolutely imperative that there is a clear and consistent separation between the data that are to be classified (the *test data*), and the data that are used to train the classifier (the *training data*). If the test data are included with the training data, and used to train the classifier, then when the test data are classified, they will be “recognised” by the classifier. This skews the performance of the classifier, resulting in a much better performance – this is not real, however, it is *cheating*. It is cheating because we are using our prior knowledge of the test subject’s class in the training and classification of that subject. This is something that, were we to implement a system like this in a clinical environment, we could not do, because we would be using the system to classify the subject – we would not have that prior knowledge.

For this reason, great care was taken to ensure that *only* training data were presented to the classifier for training, and that test data were *never* included at this stage, and the only time the classifier saw the test data, was when it was being simulated. However, merely isolating test data from the training process is not sufficient. Test data must not be considered if we perform any preprocessing on the data, prior to neural net training. For example, the selection of variables can not use our fore-knowledge of the test subject’s class.

At no stage can this knowledge be used, because we are trying to develop a tool to estimate the subject’s class when it is unknown.

A.3.4 Averaging multiple neural network simulations

We wished to find the *average* class-estimate across multiple network simulations, however, the output stage of the neural network was a *logsig* function. Averaging could not be performed on the output of the neural network directly, because that was in a *logsig* space (a nonlinear transform) so that simply estimating the mean would not be meaningful. Instead, these results were mapped to a *likelihood space*, where $-\infty$ corresponded to an impossible event, and ∞ represented one that was certain. These results could then be averaged across the simulations, and mapped back to *logsig* space to provide a result in the range of $[0\ 1]$ representing the likelihood that a subject was in a given class.

$$\text{logsig}(n) = \frac{1}{1 + e^{-n}} \quad (\text{A.2})$$

$$\text{logsig}^{-1}(n) = -\ln\left(\frac{1}{n-1}\right) \quad (\text{A.3})$$

Once the data were transformed into a likelihood space, they were ready for the final classification. For each subject, we had results from 15 neural networks, each of which had analysed each instance of the subject. Thus, for each subject we had 15 classification attempts for each of the 100 or so instances. Multiple neural networks were trained and tested because the random allocation of weights during initialisation, and the nonlinear training methods mean that the final performance of networks varies. Thus, we wanted to select for the networks that performed better. This process was quite simple – we selected for networks where there was a large difference between the probability estimate of the two classes. If we look at two fictional sets of data (see table A.1), with ten instances in each we can see an illustration of what I mean by

“separation” of data. In both classes of data, the upper row is generally high in the first 5 instances, and low in the second 5. In the first table, however, the second row varies with the first, whereas in the second table, the second row is low for the first 5 instances, and high for the second 5.

The first table shows data that have relatively poor separation between the rows.

0.64	0.99	0.45	0.78	0.89	0.35	0.27	0.18	0.44	0.57
0.47	0.77	0.38	0.54	0.68	0.44	0.07	0.20	0.19	0.59

The second table shows data that have good separation.

0.64	0.99	0.45	0.78	0.89	0.35	0.27	0.18	0.44	0.57
0.08	0.25	0.01	0.44	0.56	0.78	0.99	0.89	0.75	0.95

Table A.1: Fictional data, illustration separation of results

Please note that these two tables contain fictional data, and have identical upper rows. The data are structured so that there are two classes - the first 5 samples belong to one class, and the second five to another. Both first rows show a difference between the two classes.

The first table’s second row, however, closely tracks the scores in the first row. Hence, the neural network that produced these results will not allow classification as well as the second table’s second row, where the values are complementary to the first row.

Clearly the data in the lower table are much more easily separable into two classes than the data in the upper table.

In the real data, we have 15 tables like each of these, and each table contains approximately 100 instances. To classify the subject, we want to focus on the neural networks that produce good separation between data. A simple method, and the one chosen, was to find the mean of the absolute value of the difference between the rows. A tempting method would be to then only take the neural networks whose $\text{mean}(\text{abs}(\text{diff}(\text{data})))$ values were above a certain threshold, however the selection of this threshold would be arbitrary and would likely be influenced by the results it produced. This would be cheating!

For this reason, the classification by each neural network was multiplied by the $\text{mean}(\text{abs}(\text{diff}(\text{data})))$, which amounts to weighting the results of each neural network by the average separation of the results that it produces. *This is a blind process* - it has no consideration for the known classes of the instances.

References

- [1] *Principal Components and Factor Analysis*. StatSoft Inc.
<http://www.statsoftinc.com/textbook/stfacan.html>, July 2003.
- [2] H.D.I. Abarbanel. *Analysis of Observed Chaotic Data*. Springer, 1996.
- [3] A. Alkan, E. Koklukaya, and A. Subasi. Automatic seizure detection in EEG using logistic regression and artificial neural network. *Journal of Neuroscience Methods*, 148:167–176, 2005.
- [4] Anon. *Wavelet Theory*. Czech Technical University in Prague
http://www.cyber.felk.cvut.cz/gerstner/biolab/bio_web/teach/KP/WaveletTheory.pdf,
Oct, 2003.
- [5] M. Bear, B. Connors, and M. Paradiso. *Neuroscience - Exploring the Brain*. Lippincott Williams and Wilkins, 1996.
- [6] A.J. Bell and T.J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.
- [7] M. Breakspear, J.A. Roberts, J.R. Terry, S. Rodrigues, N. Mahant, and P.A. Robinson. A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis. *Cerebral Cortex*, 16:1296–1313, 2006.
- [8] G.M.I. Chowdhury, Y. Liu, M. Tanaka, T. Fujioka, A. Ishikawa, and S. Nakamura. Cortical spreading depression affects Fos expression in

- the hypothalamic paraventricular nucleus and the cerebral cortex of both hemispheres. *Neuroscience Research*, 2003.
- [9] J. O. Coplien. *Software patterns*, 1996.
- [10] N.C. de Lanerolle and T Lee. New facets of the neuropathology and molecular profile of human temporal lobe epilepsy. *Epilepsy & Behavior*, 7:190–203, 2005.
- [11] P.L.C. Van den Broek, J. Van Egmond, C.M. Ven Rijn, R. Dirksen, A.M.L. Coenen, and L.H.D.J. Booij. The application of a non-linear analysis technique to the monitoring of anesthetic effects in the rat. *Chaos in the brain. Proceedings of the 1999 Workshop*, pages 259–262, 2000.
- [12] I. Dvorak. Takens versus multichannel reconstruction in EEG correlation exponent estimates. *Physics Letters A*, 151(225):225–233, 1990.
- [13] B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, Jan 1979.
- [14] T Eid, M J Thomas, D D Spencer, E Rundén-Pran, J C K Lai, GV-Malthankar, J H Kim, N C Danbolt, O P Ottersen, and N C de Lanerolle. Loss of glutamine synthetase in the human epileptogenic hippocampus: possible mechanism for raised extracellular glutamate in mesial temporal lobe epilepsy. *Mechanisms of disease*, 363:28–37, 2004.
- [15] Thomas C Ferree. Spline interpolation of the scalp EEG. Technical report, Electrical Geodesics, Inc, August 2000.
- [16] Thomas C Ferree and Ramesh Srinivasan. Theory and calculation of the scalp surface laplacian. Technical report, Electrical Geodesics, Inc., August 2000.
- [17] R.M. Frank and G.A. Frishkoff. Automated protocol for evaluation of electromagnetic component separation (APECS): Application of a framework

- for evaluating statistical methods of blink extraction from multichannel EEG. *Clinical Neurophysiology*, 118:80–97, 2007.
- [18] T.D. Frank, A. Daffertshofer, C.E. Peper, P.J. Beek, and H. Haken. Towards a comprehensive theory of brain activity: Coupled oscillator systems under external forces. *Physica D*, 144:62–86, 2000.
- [19] J. Froyland. *Chaos and Coherence*. Institute of Physics Publications, 1992.
- [20] J. Gleik. *CHAOS Making A New Science*. Cardinal, 1989.
- [21] Robert L Harvey. *Neural Network Principles*. Prentice Hall, 1994.
- [22] Simon Haykin. *Neural Networks*. Prentice Hall, 1999.
- [23] D.G. Herrera, D. Maysinger, and M. Goiny. Induction of c-Fos immunoreactivity in the hippocampus following potassium stimulation. *Neuroscience*, 52(2):237–244, 1993.
- [24] J.J. Hiscock, L. Mackenzie, and J.O. Willoughby. Pentobarbitone induces Fos in astrocytes: Increased expression following picrotoxin and seizures. *Experimental Neurology*, 139:115–120, 1996.
- [25] Douglas R. Hofstadter. *Gödel, Escher, Bach: an eternal golden braid*. Penguin Books, 1979.
- [26] Jian-Hui Jiang, Roumiana Tsenkova, and Yukihiro Ozaki. *Principal Discriminant Variate Method for Classification of Multicollinear Data: Principle and Applications*, volume 17. The Japan Society for Analytical Chemistry, 2001.
- [27] Hongkui Jing and Morikuni Takigawa. Low sampling rate induces high correlation dimension on electroencephalograms from healthy subjects. *Psychiatry and Clinical Neurosciences*, 54(4):407–412, 2000.

- [28] H. Kantz and T. Schreiber. *Nonlinear Time Series analysis*. Cambridge University Press, 1997.
- [29] H. Kantz, T. Schreiber, and I. Hoffmann. Nonlinear noise reduction: A case study on experimental data. *Physical Review E*, 48(2):1529–1538, 1993.
- [30] D. T. Kaplan. Model-independent technique for determining the embedding dimension. In Louis M. Pecora, editor, *Chaos in Communications*, pages 236–240. SPIE-The International Society for Optical Engineering, Bellingham, Washington, 98227-0010, USA, 1993.
- [31] J. Kayser and C.E. Tanke. Principal components analysis of laplacian waveforms as a generic method for identifying erp generator patterns: Ii. adequacy of low-density estimates. *Clinical Neurophysiology*, 117:369–380, 2006.
- [32] M.E. Kelly, R.A. Battye, and D.C. McIntyre. Cortical spreading depression reversible disrupts convulsive motor seizure expression in amygdala-kindled rats. *Neuroscience*, 91(1):305–313, 1999.
- [33] J.A.S Kelso. *Dynamic Patterns - The Self Organization of Brain and Behavior*. MIT Press, 1995.
- [34] Wlodzimierz Klonowski. Quantitative measure of complexity of EEG signal dynamics.
- [35] Erwin Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, Inc., 1999.
- [36] D. Kugiumtzis. Test your surrogate data before you test for nonlinearity. *Physical Review E*, 60(3):2808–2816, September 1999.

- [37] D. Kugiumtzis. On the reliability of the surrogate data test for nonlinearity in the analysis of noisy time series. *International Journal of Bifurcation and Chaos*, 11(7):1881–1896, 2000.
- [38] B. Larrosa, J. Pastor, L. López-Aguado, and O. Herreras. A role for glutamate and glia in the fast network oscillations preceding spreading depression. *Neuroscience*, 141:1057–1068, 2006.
- [39] Mirosław Latka, Ziemowit Was, Andrzej Kozik, and Bruce J. West. Wavelet analysis of epileptic spikes, <http://arxiv.org/pdf/physics/0301065>, Oct 2002.
- [40] A.A.P. Leão. Spreading depression of activity in the cerebral cortex. *Neurophysiology*, 7:359–390, 1944.
- [41] K. Lehnertz. Non-linear time series analysis of intracranial EEG recordings in patients with epilepsy an overview. *International Journal of Psychophysiology*, 34:45–52, 1999.
- [42] Klaus Lehnertz, Ralph G. Andrzejak, Jochen Arnhold, Thomas Kreuz, Florian Mormann, Christoph Rieke, Guido Widman, and Christian E. Elger. Nonlinear EEG analysis in epilepsy. *Clinical Neurophysiology*, 18(3):209–222, 2001.
- [43] B.E. Lingwood, K.R. Dunster, G.N. Healy, L.C. Ward, and P.B. Colditz. Cerebral impedance and neurological outcome following a mild or severe hypoxic/ischemic episode in neonatal piglets. *Brain Research*, 969:160–167, 2003.
- [44] E.W. Lothman and R.C. Collins. Kainic acid induced limbic seizures: metabolic, behavioral electroencephalographic and neuropathological correlates. *Brain Research*, (218):299–318, 1981.
- [45] L. Mackenzie, A. Medvedev, JJ Hiscock, KJ Pope, and JO Willoughby. Picrotoxin-induced generalised convulsive seizure in rat: changes in re-

- gional distribution and frequency of the power of electroencephalogram rhythms. *Clinical Neurophysiology*, 113(4):586–96, 2002.
- [46] L Mackenzie, K.J. Pope, and J.O. Willoughby. Physiological and pathological spindling phenomena have similar regional EEG power distributions. *Brain Research*, 1008:92–106, 2004.
- [47] L Mackenzie, K.J. Pope, and J.O. Willoughby. Gamma rhythms are not integral to EEG spindle phenomena. *Clinical Neurophysiology*, 116:861–870, 2005.
- [48] J. Martinerie, C. Adam, M. le van Quyen, M. Baulac, S. Clemencleau, B. Renault, and F.J. Varela. Epileptic seizures can be anticipated by non-linear analysis. *Nature Medicine*, 4(10):1173–1176, 1998.
- [49] A. Medvedev, L. Mackenzie, J. J. Hiscock, and J. O. Willoughby. Kainic acid induces distinct types of epileptiform discharge with differential involvement of hippocampus and neocortex. *Brain Research Bulletin*, 52(2):89–98, 2000.
- [50] A. Medvedev and J.O. Willoughby. Autoregressive modeling of the EEG in systemic kainic acid-induced epileptogenesis. *International Journal of Neuroscience*, 97(3-4):149–167, 1999.
- [51] D. Mewett. Recurrence quantification analysis of electromyogram-force relationships. Master’s thesis, Flinders University, South Australia, 2002.
- [52] I. Midzyanovskaya, V. Strelkov, C. van Rijn, B. Budziszewska, E. van Luijtelaar, and G. Kuznetsova. Measuring clusters of spontaneous spike-wave discharges in absence epileptic rats. *Journal of Neuroscience Methods*, 154:183–189, 2006.
- [53] L. Nyikoc, B. Lasztóczy, K. Antal, R. Kovács, and J. Kardos. Desynchronisation of spontaneously recurrent experimental seizures proceeds with a single rhythm. *Neuroscience*, 121:705–717, 2003.

- [54] T. Olsson, M. Broberg, K.J. Pope, A. Wallace, L Mackenzie, F. Bloomstrand, M. Nilsson, and J.O. Willoughby. Cell swelling, seizures and spreading depression: an impedance study. *Neuroscience*, 140:505–515, 2006.
- [55] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Physical Review Letters*, 45:712–716, 1980.
- [56] J.M. Palacios, J.K. Wamsley, and M.J. Kuhar. High affinity gaba receptors-autoradiographic localization. *Brain Research*, 222(2):285–307., Oct 1981.
- [57] M. Van Quyen, J. Martinerie, M. Baulac, and F. Varela. Anticipating epileptic seizures in real time by a non-linear analysis of similarity between EEG recordings. *Neuroreport*, 10(10):2149–2155, 1999.
- [58] Ira J. Rampil. A primer for EEG signal processing in anesthesia. *Anesthesiology*, 89(4):980–1002, 1998.
- [59] Howard L. Resnikoff and Jr Raymond O. Wells. *Wavelet Analysis – The Scalable Structure of Information*. Springer, 1998.
- [60] B.D. Ripley. *Pattern Classification and Neural Networks*. Cambridge, 1996.
- [61] Osvaldo A. Rosso. Entropy changes in brain function. *International Journal of Psychophysiology*, (64):75–80, 2007.
- [62] M.E Saab and J. Gotman. A system to detect the onset of epileptic seizures in scalp EEG. *Clinical Neurophysiology*, 116:427–442, 2005.
- [63] IE Scheffer and SF Berkovic. Generalized epilepsy with febrile seizures plus. A genetic disorder with heterogeneous clinical phenotypes. *Brain*, 120(3):479–490, 1997.

- [64] G. Scheler. Algorithmicity in neural dynamics - the emergence of conditional loops, <http://www.icsi.berkeley.edu/2003>.
- [65] N.D. Schiff, D.R. Labar, and J.D. Victor. Common dynamics in temporal lobe seizures and absence seizures. *Neuroscience*, 91(2):417–428, 1999.
- [66] S.J. Schiff, D. Colella, G.M. Jacyna, E. Hughes, J.W. Creekmore, A. Marshall, M. Bozek-Kuzmicki, G. Benke, W.D Gaillard, J. Conry, and S.R. Weinstein. Brain chirps: spectrographic signatures of epileptic seizures. *Clinical Neurophysiology*, 111:953–958, 2000.
- [67] T. Schreiber and A. Schmitz. Surrogate time series, 2001.
- [68] Thomas Schreiber and Andreas Schmitz. Improved surrogate data for nonlinearity tests. *Physical Review Letters*, 77(4):635–639, July 1996.
- [69] H.P. Schwan and C.F. Kay. Specific resistance of body tissues. *Circulation Research*, 4:664–670, 1956.
- [70] C. Shannon. A mathematical theory of communication. *The Bell Systems Technical Journal*, 27:379–423, 1948.
- [71] Ramesh Srinivasan. Methods to improve the spatial resolution of EEG. *International Journal of Bioelectromagnetism*, 1(1):102–111, 1999.
- [72] V. Srinivasan, C. Eswaran, and N. Sriraam. Artificial neural network based epileptic detection using time-domain and frequency-domain features. *Journal of Medical Systems*, 29(6):647–660, December 2005.
- [73] O.K. Steinlein, A. Magnusson, J. Stoodt, S. Bertrand, S. Weiland, S.F. Berkovic, K.O. Nakken, P. Propping, and D. Bertrand. An insertion mutation of the chrna4 gene in a family with autosomal dominant nocturnal frontal lobe epilepsy. *Human Molecular Genetics*, 6:943–947, 1997.

- [74] Fernanda Strozzi. Application of nonlinear time series analysis techniques to high frequency currency data. *Serie Metodice quantitative*, 14(99):111–112, 2002.
- [75] C. Tandonnet, B. Burle, T. Hasbroucq, and F. Vidal. Spatial enhancement of EEG traces by surface laplacian estimation: comparison between local and global methods. *Clinical Neurophysiology*, 116:18–24, 2005.
- [76] C.E. Tenke and J. Kayser. Reference-free quantification of EEG spectra: Combining current source density (csd) and frequency principal components analysis (fpca). *Clinical Neurophysiology*, 116:2826–2846, 2005.
- [77] J. Theiler, S. Eubank, A. Longtin, and B. Galdrikian. Testing for nonlinearity in time series: the method of surrogate data. *Physica D*, 58:77–94, 1992.
- [78] M.J. van der Heyden, .D.N. Velis, B.P.T. Hoekstra, J.P.M. Pijn, W. van Emde Boas, C.W.M. van Veelen, P.C. van Rijen, F.H. Lopes da Silva, and J. DeGoede. Non-linear analysis of intracranial human EEG in temporal lobe epilepsy. *Clinical Neurophysiology*, 110:1726–1740, 1999.
- [79] M. van Putten. Proposed link rates in the human brain, 2002.
- [80] T. A. A. Watanabe, C. J. Cellucci, E. Kohegyi, T. R. Bashore, R. C. Josiassen, N. N. Greenbaun, and P. E. Rapp. The algorithmic complexity of multichannel EEGs is sensitive to changes in behavior. *Psychophysiology*, 40:77–97, 2003.
- [81] Ed. Webster, J. *Medical Instrumentation: Application and Design*. John Wiley & Sons, Inc., 1998.
- [82] Eric W. Weisstein. *The CRC Concise Encyclopedia of Mathematics*. CRC Press, 1998.

- [83] E.M. Whitham, K.J. Pope, S.P. Fitzgibbon, T. Lewis, C.R. Clark, S. Lovelless, M. Broberg, A. Wallace, D. DeLosAngeles, P. Lillie, A. Hardy, R. Fronsco, A. Pulbrook, and J.O. Willoughby. Scalp electrical recording during paralysis: quantitative evidence that EEG frequencies above 20 hz are contaminated by emg. 2007.
- [84] J O Willoughby, S P Fitzgibbon, K J Pope, L Mackenzie, A V Medvedev, C R Clark, M P Davey, and R A Wilcox. Persistent abnormality detected in the non-ictal electroencephalogram in primary generalised epilepsy. *Journal of Neurology Neurosurgery and Psychiatry*, 74(1):51–55, 2003.
- [85] J. O. Willoughby, L. Mackenzie, A. Medvedev, and J. J. Hiscock. Fos induction following systematic kainic acid: early expression in hippocampus and later widespread expression correlated with seizure. *Neuroscience*, 77(2):379–392, 1997.
- [86] John O. Willoughby, Lorraine Mackenzie, Marita Broberg, Anna E. Thoren, Andrei Medvedev and Neil R. Sims, and Michael Nilsson. Fluorocitrate-mediated astroglial dysfunction causes seizures. *Journal of Neuroscience Research*, 74:160–166, 2003.
- [87] H. Yang, S. Jiang, J. Yu, J.Wang, Z. Tan, H. Xue, G. Feng, and L. He. Complete genomic sequence of 195 kb of human dna containing the gene gabrg2. *DNA Seq*, 11(5):373–382, 2000.