DYNAMICAL SYSTEM MODELLING OF THE EEG WITH APPLICATION TO DIRECT BRAIN INTERFACING

By

DAVID P. BURKE
B.E. M.ENG.SC.

A thesis presented to
The National University of Ireland, Dublin
in fulfilment of the requirements of
the degree of

PHILOSOPHIAE DOCTOR

in the
DEPARTMENT OF ELECTRONIC AND ELECTRICAL ENGINEERING
FACULTY OF ENGINEERING AND ARCHITECTURE
THE NATIONAL UNIVERSITY OF IRELAND, DUBLIN

March 2003

Supervised by
An tOllamh Annraoi de Paor

Head of Department
Professor T. J. Brazil

ABSTRACT

For some people with severe physical disabilities, brain electrical activity may represent the only feasible channel for communicating with others and for environmental control. A Direct Brain Interface (DBI) operates by harnessing signals arising from processes within the brain, without depending on the brain's normal output pathways of peripheral nerves or muscles. The Electroencephalogram (EEG) is one such signal and this thesis seeks to expand our knowledge of the EEG signal using a substratum of dynamical systems as the primary technical tool. Indeed, characterising the dynamical properties of the EEG has both practical and theoretical implications. Knowledge of the underlying dynamics can assist in the identification of signal processing algorithms for DBI and diagnostic purposes. Secondly, it can put constraints on mathematical properties used to model brain dynamics. While characterisation of the nonlinearity in the EEG as arising from low-dimensional chaos has now all but been discounted in the literature, little has been offered as an alternative. A large focus of this thesis is thus the proposition of such a characterisation, with the aid of an empirical model that combines both deterministic and stochastic behaviour, under the umbrella of dynamical systems theory.

We commence our technical exposition by introducing nonlinear dynamics in the context of Lyapunov theory. We propose a geometric method for synthesising prescribed limit cycle oscillators and subsequently prove asymptotic stability of the limit cycles. The prescribed oscillators are extended by employing Itô calculus to produce dynamical systems exhibiting stochastic limit cycles. A stochastic limit cycle is defined as an invariant set by applying stochastic stability theory. Both purely mathematical- and biophysical-based models of the EEG are created based on stochastic oscillators. A study of the phase space embedding of the actual EEG is undertaken and invariant measures compared and contrasted with those obtained from model data. Next, more detailed systems are considered, employing coupled oscillators to study the interactions of neurons and neural populations as a basis for Evoked Potentials (EP) generation – an important feature of the EEG with application to DBIs. We conclude the thesis with practical applications by presenting novel human-computer interfaces, established with regard to the preceding chapters.

ACKNOWLEDGEMENTS

The work represented within these pages is the culmination of 3.5 years of (mostly) very late nights, early mornings, and busy weekends! Having got this far and survived, it is clear to me that I could not have achieved this unaided.

I have to sincerely thank Harry. Not just for the usual reasons that a research supervisor is thanked by his student but for giving me the freedom, encouragement, and inspiration to fulfil easily my most difficult project to date. Allowing me to venture out into industry so early in this endeavour with your blessing and confidence will not be forgotten – thanks for believing in me and letting me do it my way!

Special thanks to Louise who has seen me through another degree from start to finish. You have amazing tolerance and I promise no more big projects (at least for a while)!

Many thanks to the Voxpilot crew for making it such an energetic and fun 'break' from research. Bringing VoiceXML to Europe was fun! – I hope some of you get to experience postgraduate research also. Thanks to all at the 'lab' and NRH (especially for welcoming me back after prolonged disappearances) – Ted, Jane, Kieran, Ciarán, Deirdre, Yvonne, and all the past generations including Tomás for putting the idea into my head that I might be able to do this! Thanks to Diarmuid for building the analogue computer in Appendix A1 – still impressed it worked first time!

Thanks to Richard for encouragement, understanding, and counsel on both sides of the academia-industry join. Thanks to Dr Philip Nolan and Prof. C. J. Stam for assisting with the specialist EEG data. Thanks to Simon and Ciarán for assisting with the EEG data collection for Chapter VI. Thanks to Conor for helping out in Houston. Thanks to Kevin and Hugh for a fruitful collaboration on fMRI data processing. And finally, thanks to my parents for helping put down the foundations to allow me to pursue postgraduate study.

Dave Burke, March 6th 2003

PUBLICATIONS ARISING FROM THIS THESIS

JOURNAL

BURKE, D., KELLY, S., DE CHAZAL, P., REILLY, R. B., and FINUCANE, C. (2005): 'A Parametric Feature Extraction and Classification Strategy for Brain-Computer Interfacing', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13, pp. 1 – 6

BURKE, D., and DE PAOR, A. M. (2004): 'A stochastic limit cycle oscillator model of the EEG', *Biological Cybernetics*, 91, pp. 221 – 230

BURKE, D., MURPHY, K., GARAVAN, H., and REILLY, R. B. (2004): 'Pattern recognition approach to the detection of single-trial event-related functional magnetic resonance images', *Medical & Biological Engineering & Computing*, 42, pp. 604 – 609

BURKE, D. (2003): 'An extensible C++ framework for solving stochastic differential equations', *C/C++ Users Journal*, 21, pp. 12 – 17

BURKE, D., WARD, T., and DE PAOR, A. M. (2001): 'Image processing used to harness blinking as a channel of communication and control for physically disabled people', *Medical & Biological Engineering & Computing*, 39, pp. 285 – 287

CONFERENCE

FINUCANE, C., **BURKE**, **D**., and DE PAOR, A. M. (2003): 'Investigations into a dual modality P300 based Brain-Computer Interface', in Proceedings of the 10th International Conference on Human-Computer Interaction, Crete, Greece, June

MURPHY, K., **BURKE, D.,** and GARAVAN, H. (2003): 'ARX technique for real-time activation and feedback of event-related fMRI activation in a cognitive task', in

Proceedings of the 9th Annual Meeting of the Organization for Human Brain Mapping (OHMB), New York, June

BURKE, D., and DE PAOR, A. M. (2002): 'An evoked potential model based on the transient synchronisation of coupled oscillators', in Proceedings of the 2nd European Medical & Biological Engineering Conference, EMBEC'02, Vienna, Dec, pp. 1434 – 1437

BURKE, **D**., KELLY, S., DE CHAZAL, P., and REILLY, R. B. (2002): 'A simultaneous filtering and feature extraction strategy for direct brain interfacing', in Proceedings of the 24th Annual International Conference of the Engineering in Medicine and Biology Society (EMBS-BMES), Houston, Texas Oct (on CDROM)

KELLY, S., **BURKE**, **D**., DE CHAZAL, P., and REILLY, R. B. (2002): 'Parametric models and spectral analysis for classification in brain-computer interfaces', in Proceedings of the 14th International Conference on Digital Signal Processing, Greece, July, pp. 307 – 310

KELLY, S., **BURKE**, **D**., DE CHAZAL, P., and REILLY, R. B. (2002): 'Parametric models and classification for direct brain interfaces', in Proceedings of the Irish Signals and Systems Conf., Cork, June (paginated)

BURKE, D. and DE PAOR, A. M. (2001): 'Signal prediction based on a chaotic attractor model of the electroencephalogram', in Proceedings of the 3rd International Conference on Measurement, Slovak Republic, May (paginated)

BURKE, **D.** (2001): 'Assistive technology for educational use by a young girl with total Amelia', Spring Meeting of the British Society of Rehabilitation Medicine, Belfast Waterfront Hall, June (abstract only)

BURKE, D. (2000): 'Image processing used to harness blinking as a channel of communication and control for physically disabled people', Irish Association of Rehabilitation Medicine, Belfast, Sept (abstract only)

TABLE OF CONTENTS

CHAPTER I

BACKGROUND

1.1 Introduction

The Electroencephalogram, or EEG as it is commonly known, refers to the tiny electrical potentials that may be recorded from the surface of the scalp. The EEG is the result of the dynamic behaviour of neural populations in the brain operating in a cooperative and synchronous fashion (SPECKMANN and ELGER, 1999). For some people with very severe physical disabilities (such as C3 lesions), the minute potentials of the EEG may represent the only feasible channel for communicating with others and for environmental control[1]. This application forms the primary motivation of this thesis, that is, to facilitate communication and control by harnessing the EEG signal in a Direct-Brain Interface (DBI). The human brain consists of approximately 10^{10} neurons realised in many identified and as yet unidentified structures and hence it is not surprising that the current understanding of the resulting EEG is still in its infancy (STERIADE, 1999). The central goal of this work is to explore the EEG and related signals using a substrate of dynamical system theory as the primary technical tool. Creating models of the EEG encourages one to form hypotheses for the generation of the signal, affords one the ability to perform experiments and test hypotheses otherwise impossible to undertake in an intact brain, and may predict new properties of the system suggesting further experimental study. Moreover, an understanding of the underlying dynamics can help motivate the choice of tools in signal processing applications, for example in the creation of a Direct-Brain Interface. A particularly salient point here is choosing a tool on the basis of an interpretation of the EEG as a manifestation of a chaotic attractor or as some equilibrium under random perturbation.

The EEG has found many applications in both research and clinical settings since its discovery by Hans Berger in 1929 (NIEDERMEYER and LOPES DA SILVA, 1999). Much of the benefit of the EEG signal is derived from its non-invasiveness and its excellent time resolution, in particular when compared to alternative techniques such as

[1] Obviously many sensitive and contentious issues arise with severe physical disabilities such as C3 lesions (FIRSCHING, 1998) but there appears to often be a role for alternative communication devices

Functional Magnetic Resonance Imaging or Positron Emission Tomography. Good time resolution is an important requirement for creating DBIs and hence helps explain the popularity of the EEG for this purpose. Table 1.1 summarises some common applications of the EEG partitioned into both clinical and research domains.

TABLE 1.1: APPLICATIONS OF THE EEG

Clinical Domain	Research Domain
Sleep analysis	Direct Brain Interface
Epilepsy detection	Behavioural studies
Spinal cord monitoring	Epilepsy seizure prediction
Cerebral monitoring	Brain function studies
Anaesthesia depth measurement	
Biofeedback	

Modelling the EEG is a difficult task. Much of what is known about neurons has been learned from isolated nerve cells and most of our knowledge of the cortex has come from isolated slices maintained outside the brain. However, these slices do no oscillate and do not process information. The choice of scale in the modelling task is also paramount, and might vary from the single neuron level, to the neural population, to communities of neural populations. Freeman uses the phrases microscopic, mesoscopic, and macroscopic respectively to represent these scales (FREEMAN, 2000a). A wide variety of modelling techniques exist, all ultimately reducible to the common language used in this thesis: dynamical systems. We review some of the more popular models from the literature in Section 1.3. Such is the state of knowledge regarding the EEG that even the choice of interpretation of the resultant signal as the result of chaotic effects or stochastic effects is uncertain and it is this aspect of modelling the EEG that holds our focus in later chapters.

1.2 Neurophysiological Basis of Brain Electrical Activity

This section outlines some of the basic neurophysiological principles underlying brain electrical activity and the EEG. Here space restrictions necessitate a rather cursory overview, in particular nervous system and cortical structure details have been omitted. However, it should be stressed that only through a thorough understanding of the relevant neurophysiological details are engineers, mathematicians, and physicists truly able to contribute to the problems of uncovering the workings of the nervous system. There are countless books on physiology and neurophysiology – the author

has found GUYTON (1991) indispensable throughout the work for this thesis. For general information on electroencephalography, including its neurophysiological principles, the book by NIEDERMEYER and LOPES DA SILVA (1999) is highly recommended.

1.2.2 The Neuron

The neuron is the elemental unit from which the nervous system is composed. The central nervous system (i.e. brain and spinal cord) contains an enormous 10^{11} neurons. There are three main components to a neuron: the cell body or *soma*, the *axon*, and the *dendrites*. The dendrites form repeatedly branching structures corresponding to the receptive portion of the neuron, which continue to grow during the lifetime of the brain. The input received at the dendrites (from other neurons) is converted to an output by the single axon, which emerges from the cell body. The axon carries the neural output over long distances and may have branched offshoots known as *collaterals*, which spatially distribute the neuron's output (sometimes with *recurrent collaterals* that transmit information to neurons in its immediate vicinity). Two types of structures are seen. In some cases, the axon only has a small bush of terminals thus providing a topographic mapping from a transmitting population of neurons to a target population. In other cases, the axon branches and diverges repeatedly onto the target and the target therefore performs spatial integration.

Synapses connect neurons to each other (see Section 1.2.2.2). Located primarily on the dendrites (but also the soma), the synapses provide the electromotive force for the current whose preferred path is along the dendritic shaft toward the cell body. The direction of the current depends on whether the synapse is excitatory or inhibitory but by Kirchhoff's current law, the current flowing across the cell membrane is in the opposite direction to the direction at the synapse – see Figure 1.1. In particular, the currents sum at the initial segment of the axon known as the *trigger zone*. A sufficiently excited neuron may respond by "firing" i.e. producing an electrical pulse called an *action potential* that propagates down its axon where it impinges upon other neurons. The action potential is a simple and elegant mechanism for communication. The signal is a binary response – there is no graded output or negative response and the direction is always from the soma to the axon extremities.

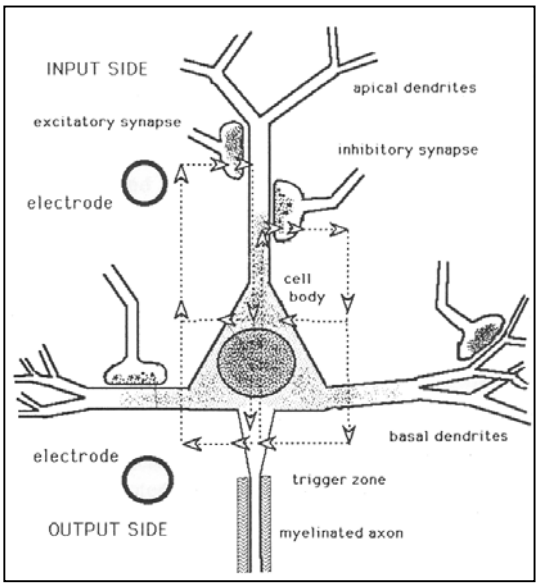

Figure 1.1: Loop currents flow in the extracellular (and intracellular) space driven by the EMF produced from thousands of synapses. The current flowing through the extracellular space sums with currents produced from neighbouring neurons resulting in the local mean field potential recorded by electrodes as the EEG. Reproduced from FREEMAN (1992)

The transmission is an active mechanism with the energy being provided locally by the axon's excitable membrane. This is in contrast to the passive properties of the dendrite, which obeys linear superposition and may be conveniently modelled as a linear ODE. In summary, neurons convert the sum of their current inputs into a pulse train where the pulse frequency is proportional to the current input.

1.2.2.1 Membrane Potential

The cell *membrane* provides a boundary separating the internal workings of the cell from its external environment. The membrane, consisting of phospholipid molecules, is selectively permeable thus permitting the passage of some materials while restricting others. The membrane wall contains water filled pores and protein-lined pores called *channels* which allow the passage of specific molecules. In particular, the cell membrane acts as a barrier to the flow of water. Both the intracellular and extracellular environments include an aqueous solution of dissolved NaCl and KCl salts, which dissociate into Na^+, K^+, and CL^- ions. Molecules may be transported across the cell membrane by passive or active processes. An active process is defined as one that requires the expenditure of energy, while a passive process results solely from the inherent, random movement of molecules. Some small molecules can pass through the membrane by diffusion, while large molecules require *carrier-mediated*

diffusion where the molecule is bound to a carrier molecule that moves readily through the membrane. Concentration differences are maintained by active mechanisms that use energy to pump ions against their concentration gradients. One of the most important examples is the Na^+-K^+ pump, which uses the energy stored in ATP molecules to pump Na^+ out of the cell and K^+ into the cell. The Na^+-K^+ pump moves 3 sodium ions out of the cell for every 2 potassium ions pumped in, resulting in a high concentration of sodium outside the cell and a high concentration of potassium inside the cell. Maintaining the intracellular ionic concentrations at their correct levels is required by the cell to control its volume, as the membrane is unable to withstand any hydrostatic pressure that would occur if the pumping machinery was disabled.

Measuring the intracellular potential with respect to the extracellular region displays a negative potential. To understand the reasons for this, consider a membrane only permeable to potassium. Since there is a greater intracellular concentration of potassium ions, the result is a net outward flow of ions. This outflow of positive ions will cause the electrical charge inside the cell to drop. Eventually the negative potential will become strong enough to prevent the positively charged potassium ions from leaving the cell. The equilibrium condition of the potential gradient offsetting the concentration gradient results in the *Nernst* potential (valid for a single ion species)

$$V_s = \frac{RT}{zF} \ln\left(\frac{[S]_e}{[S]_i}\right)$$

where $[S]_e$ is the extracellular concentration, $[S]_i$ is the intracellular concentration, R is the universal gas constant, T is the absolute temperature, z is the charge on the ion S, and F is Faraday's constant. The Nernst potential in large mammalian nerve fibre for potassium is -94 mV. Now consider a membrane permeable only to sodium ions. The greater extracellular concentration of sodium will result in an inward flow of ions yielding a net increase in potential inside the cell. The Nernst potential for sodium ions in large mammalian nerve fibre is +61 mV. Since the nerve membrane is 100 times more permeable to potassium than sodium, we might expect the actual resting potential to be quite negative. This is the case, with the combined effects yielding a membrane potential of -86 mV. However, the Na^+-Cl^+ pump, which causes a net removal of positive charge from the interior of the cell, further decreases the resting potential to yield a value of -90 mV.

1.2.2.2 Synapses

Transmission of information between neurons occurs via a special junction called the synapse. A typical neuron has thousands of synapses located on its dendrites and soma. A chemical synapse may be classified as *excitatory*, *inhibitory*, or *modulatory*. An excitatory synapse causes current to flow into the dendrite at the synapse, along the dendritic cable away from the synapse, outwardly across the membrane, and back to the synapse in the space outside the membrane. An inhibitory synapse, conversely, causes current to flow outwardly at the synapse and inwardly elsewhere. The minimal input to a synapse is a single action potential impulse lasting about one millisecond. The corresponding current response seen in the dendrite (assuming an excitatory synapse) rises monotonically to a maximum in about 1 to 3 milliseconds and decays exponentially with a decay time of 5 to 10 milliseconds. The current flow causes an increase of potential, in particular at the trigger zone of the soma and is called an *excitatory postsynaptic potential* (EPSP); an inhibitory synapse has exactly the opposite effect resulting in an *inhibitory postsynaptic potential* (IPSP). The third kind of synapse, called modulatory, operates indirectly to modify the strength of neighbouring excitatory and inhibitory synapses, and is the subject of much contemporary neurophysiology.

The actual mechanism of a chemical synapse involves, as the name implies, chemicals known as *neurotransmitters*. An action potential arriving at the presynaptic zone will initiate the release of tiny packets of neurotransmitter chemicals called *vesicles*, which diffuse rapidly across the gap between the nerve axon and postsynaptic cell called the *synaptic cleft*. The neurotransmitters bind to receptors in the postsynaptic cell and initiates changes in its membrane potential. The neurotransmitter substance is subsequently removed from the synaptic cleft by diffusion and hydrolysis. The type of postsynaptic potential is determined in part by the species of neurotransmitter at the synapse. There are over 40 different types of synaptic transmitters - some of the more common ones include acetylcholine (ACh), gamma-aminobutyric acid (GABA), norepinephrine, dopamine, and serotonin.

1.2.2.3 Hodgkin-Huxley Model

Experimental data shows that when we place a pair of electrodes, one inside a neuron and the other outside, and pass a weak current pulse across the membrane, we get an

immediate change in transmembrane potential followed by an exponential decay back to the baseline. This may be modelled as a simple RC circuit as illustrated in Figure 1.2.

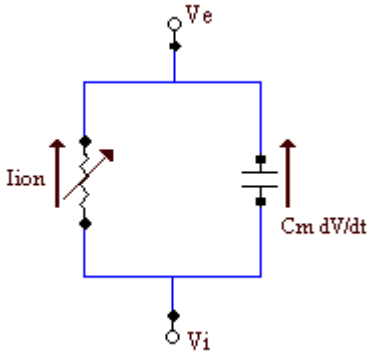

Figure 1.2: Electrical circuit model of cell membrane

The corresponding equation is

$$C_m \frac{dV}{dt} + I_{ion}(V,t) = 0 \tag{1.1}$$

where $V = V_i - V_e$. For a small applied voltage V, we may assume an ohmic resistance for the membrane and write the transmembrane current as

$I_S = g(V - V_S)$

where V_S is the Nernst potential and g is the membrane conductance. In reality, the membrane resistance is anything but ohmic over a large range of membrane voltage. Alan Hodgkin and Andrew Huxley developed the first quantitative model for the propagation of an electrical signal along a squid giant axon (HODGKIN and HUXLEY, 1952), which won them the Nobel Prize in physiology and medicine in 1963. Their model has since been extended and found applicability in modelling a wide range of excitable cells. The Hodgkin-Huxley model is arguably the most important model in the study of brain electrical activity and we briefly review it here.

In the squid giant axon, as in many neural cells, the principal ionic currents are the sodium current and potassium current. When the sodium and potassium channels are open, the I-V curves are approximately linear and we may write (1.1) as

$$C_m \frac{dV}{dt} = -g_{Na}(V - V_{Na}) - g_k(V - V_k) - g_L(V - V_L) + I_{app} \tag{1.2}$$

where other small ionic channels are lumped into a leakage current written with the subscript L, and I_{app} is an externally applied current. The Hodgkin-Huxley quantitative

model is based on replacing the channel conductances with variable conductances that are a function of both voltage and time and are a result of modelling channel gating. The Hodgkin-Huxley model restates (1.2) as

$$C_m \frac{dV}{dt} = -\bar{g}_{Na} m^3 h(V - V_{Na}) - \bar{g}_k n^4 (V - V_k) - \bar{g}_L (V - V_L) + I_{app}$$

where the conductances have been rescaled and new variables introduced. The variable m is called the *sodium activation* as it is initially small and first increases. The variable h is called the *sodium inactivation* because h inactivates or shuts down the sodium current. Similarly the variable n is the *potassium activation*. The raised powers of the activation and inactivation functions were chosen by Hodgkin and Huxley to fit their data (obtained using impressive ingenuity that contributed to their winning of the Nobel Prize). The activation functions satisfy the following ODEs:

$$\frac{dm}{dt} = \alpha_m (1 - m) - \beta_m m$$
$$\frac{dn}{dt} = \alpha_n (1 - n) - \beta_n n$$
$$\frac{dh}{dt} = \alpha_h (1 - h) - \beta_h h$$

The specific functions α and β are functions of voltage V involving exponentials. For the sake of brevity we will omit them and refer the interested reader to the excellent text of KEENER and SNEYD (1998). Figure 1.3 illustrates an action potential and gating variables using the original values of HODGKIN and HUXLEY (1952).

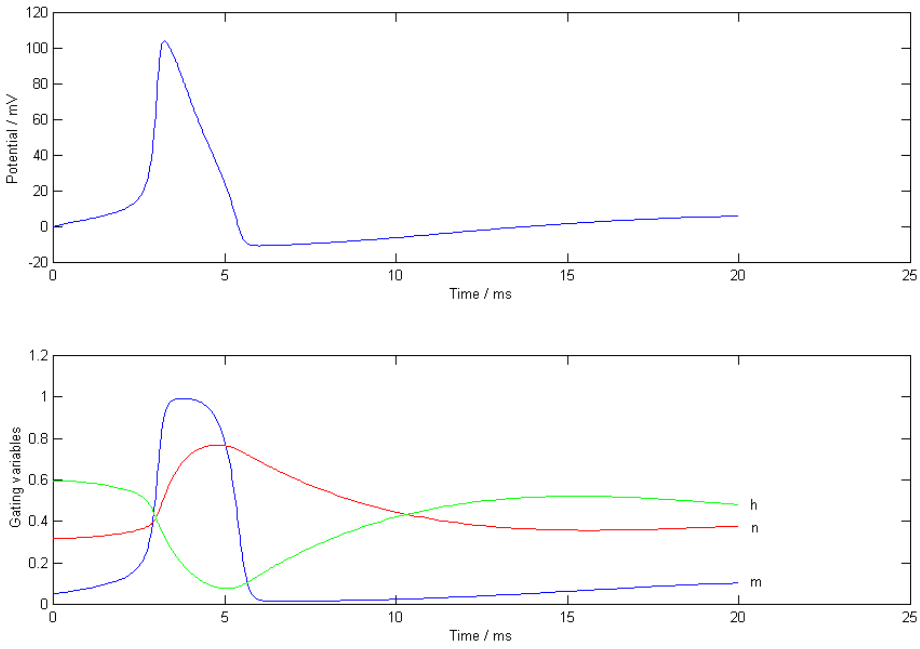

Figure 1.3: Action potential and gating variables for the Hodgkin-Huxley equations with a constant $I_{app} = 5$

1.2.3 Neural Populations

Each neuron in the human cerebral cortex receives 1,000 to 10,000 synapses and transmits to a correspondingly large number of other neurons. Neural populations do not exist as simply large numbers of neurons driven by a common input in parallel, rather they are characterised by virtue of their feedback connections – see Figure 1.4 (FREEMAN, 1975). Cooperative feedback occurs when an excitatory pool excites another and tends to couple neurons to fire more or less together. Cooperative feedback is used by the brain to coordinate the activity of neurons. With competitive feedback, one inhibitory population inhibits another and vice versa. Locally excited cells can inhibit their neighbours and in turn cause them to cease their inhibitory effect on the excited cells. This results in the excited cells becoming even more active and is called *disinhibition* – a mechanism that can provide spatial contrast processing.

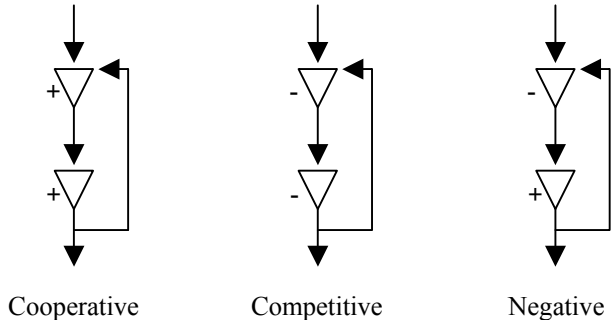

Figure 1.4: Different patterns of connections among neurons.
Reproduced from FREEMAN (1975)

With negative feedback, there is interaction between an excitatory pool and an inhibitory pool. This type of configuration is common in the brain and is the basis for oscillation. Oscillatory dynamics and synchronisation of neural activity are hypothesised to be of significant functional relevance to information processing in the brain (STURM and KÖNIG, 2001; WARD, 1998). The large number of neurons in the cortex is arranged in layers. It is now well established that pools of neurons in local neighbourhoods tend to share activity leading to the concept of cortical columns as the macroscopic entity for behaviourally related function (FREEMAN, 2000a). The column is conceived to extend the full thickness of the cortex and thus the basic dynamics of the cortex may be represented as a two dimensional layer of interconnected local populations with parallel input and output lines entering and leaving each column. Columns are interconnected more densely over short distances and more sparsely over longer distances. Finally, unlike certain cardiac circuits, the oscillations in the brain

are not thought to be driven by pacemaker cells and arise naturally from the interaction of excitatory and inhibitory populations of neurons (FREEMAN, 2000a).

1.2.4 Field Potentials and the EEG

Extracellular potentials recorded outside of the central nervous system are generally known as *field potentials*. Because of the time course of various membrane potential fluctuations, the postsynaptic potentials are known to contribute primarily to the generation of the extracellular potentials (SPECKMANN and ELGER, 1999). The primary transmembrane currents produced at the synapses generate secondary ionic currents along the cell membranes in both the intracellular and extracellular regions as illustrated in Figure 1.1. The portion of these currents that flow through the extracellular space is directly responsible for the field potentials. In addition to the neurons, the interspersed *glial* cells may also play a role in the generation of extracellular field potentials (SPECKMANN and ELGER, 1999). These cells are not excitable, i.e. do not produce action potentials and the magnitude and polarity of their membrane potential approximates the potassium equilibrium potential. An increase in extracellular potassium, for example as would occur if neighbouring neurons were repetitively firing, will cause a depolarisation of the local glial cells. Thus the glial cells are thought to have a magnifying effect on the field potentials.

Location of neurons and synapses relative to recording electrodes has a profound effect on polarity and morphology of the resulting waveforms. Figure 1.5 illustrates an excitatory synapse positioned at a superficial region of a dendrite and at a deep region of the dendrite.

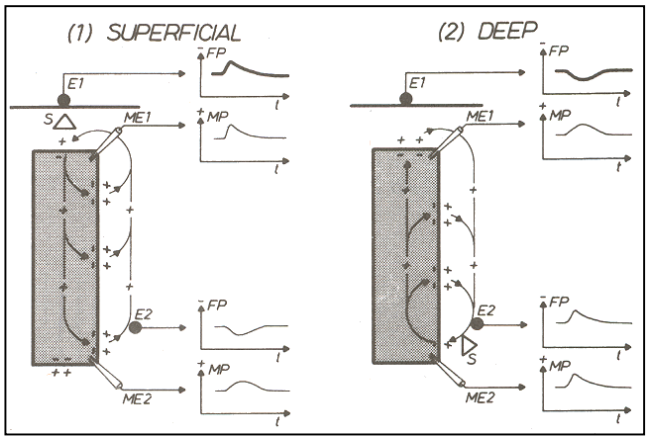

Figure 1.5: Generation of extracellular field potentials – see text for details. Reproduced from SPECKMANN and ELGER (1999)

The membrane potential of the dendritic area (shaded) is recorded at both ends by the microelectrodes ME_1 and ME_2. The extracellular field potential is recorded from the surface by the E_1 electrode as well as in the neighbourhood of ME_2 by the E_2 electrode. Excitatory synapses are marked by triangles. Consider first the superficial synapse. The inward current at S generates an EPSP that appears in ME_1 and also a smoothed and delayed version in ME_2 (the transmission may be described by the cable equation (KEENER and SNEYD, 1998)). Due to the direction of the current flow, E_2 exhibits a positive deflection (negativity drawn upwards) and E_1 exhibits a negative deflection. When the synapse is located at a deep location, the current flow direction is reversed. The records produced by ME_1 and ME_2 are swapped, as are E_1 and E_2. Finally, if we were to consider an inhibitory synapse then exactly the reverse would happen.

We conclude this section by summarising that the EEG potentials recorded from the surface of the scalp may be characterised as aperiodic oscillations with amplitude histograms that are near Gaussian (amplitudes typically range from 10 – 100 μV), autocorrelation functions that rapidly go to zero, log spectra that decrease approximately linearly with log frequency (*1/f*), and with intermittent bursts of oscillation having spectral peaks in certain preferred bands: 0.1-3.5 Hz (*delta*). 4-7.5 Hz (*theta*), 8-13 Hz (*alpha*), 14-30 Hz (*beta*) and >30 Hz (*gamma*). For a detailed clinical interpretation of the various EEG rhythms see NIEDERMEYER and LOPES DA SILVA (1999).

1.3 EEG Models from the Literature

Models of the EEG come in a variety of flavours, some of which are based on ordinary language (e.g. those used by biologists), others are comprehensive physiological constructs, whilst others still might be termed empirical models. It is often the case that empirical models (sometimes referred to as phenomenological) are the consequence of attempting to fit a simple formula to capture an ordinary language model. Further classification (related to scale) may be sought on the basis of whether the model is distributed as is usually the case with compartmental modelling techniques - for a review see KOCH and SEGEV (1998) - or lumped where for example a collection of neurons are modelled as a single process. The goal of this section is to briefly review some of the more popular models of EEG dynamics and phenomena

from the literature; selected models are subsequently reviewed in more detail for their relevance to the present thesis.

Possibly one of the earliest models of the EEG dates back to JASPER (1936) who employed a relaxation oscillator to show phenomena common to both relaxation oscillators and the EEG, namely that larger amplitudes are accompanied by lower frequencies and that oscillations are quenched for a large driving input (which in the context of the brain he termed "cortical excitatory state"). WILSON and COWAN (1972) proposed a simple and elegant model for the dynamics of tightly coupled excitatory and inhibitory neural populations. Phase plane analysis demonstrated the model to be capable of reproducing hysteresis, impulse responses, and limit cycle behaviour. The model was later extended to a 2-dimensional topology in WILSON and COWAN (1973). FREEMAN (1975) proposed a popular physiological based model of the olfactory cortex and has developed this over time to include, among others, chaotic phenomena – see FREEMAN (2000a) for a review. A lumped model of the alpha rhythm based on the interactions of the thalamus and cortex was proposed by LOPES DA SILVA (1974) and has been developed to take spatial-temporal characteristics into account by ROTTERDAM *et al* (1982). The Lopes da Silva model has appeared frequently as the basis of others' work, for example in WRIGHT and LILEY (1996) and for modelling evoked potentials in JANSEN *et al* (1993). A popular mathematical model of the EEG is its interpretation as filtered white noise. One of the earliest such proposal was by PRAST (1949), suggested as a model for the alpha rhythm. ZETTERBERG (1977) modelled the EEG with three filtered, independent noise sources corresponding to different rhythms. An equivalent formalisation of a white noise driven filter is the autoregressive model (see PARDEY *et al* (1996) for a review with application to the EEG). NUNEZ (1995) suggests an alternative mathematical approach where the alpha rhythm is modelled as a standing wave of activity on a spherical model of the cortex; a significant correlation between head size and alpha frequency is cited as experimental evidence. Nunez has also combined Rotterdam's local model with his global model. Computational neuroscience methods (KOCH and SEGEV, 1998) have been successfully employed in biologically realistic models on both large and small scales. WILSON and BOWER (1989) created a large-scale model of the piriform cortex (a part of the mammalian olfactory system) by employing compartment modelling to produce EEG-like behaviour (see also PROTOPAPAS *et al* (1998) and PROTOPAPAS and BOWER

(1998)). Using similar techniques, LILEY *et al* (1999) created a large-scale model of pyramidal neurons that was capable of producing alpha rhythm activity. GOLOMB (1996) *et al* have proposed small-scale models for the generation of 7-14 Hz spindle activity that occurs during the onset of light sleep. Other important EEG phenomena such as that resulting from epilepsy have been modelled by TRAUB (1982). More recently, interpretations of the EEG as resulting from a nonlinear and/or chaotic source, in particular during epilepsy have emerged (BABYLOYANTZ, 1985; BABYLOYANTZ and DESTEXHE, 1986).

1.3.1 Wilson-Cowan Model

The WILSON and COWAN (1972) model is based on the derivation of a set of coupled nonlinear differential equations representing spatially localised populations of excitatory and inhibitory neurons. The motivation of modelling at the coarser neural population scale is based on the observation that cognitive functions such as pattern recognition is in some sense a global process, and that modelling at the single neuron scale quickly leads to intractable problems while not providing much additional insight since local neuron-to-neuron interactions appear random anyway. Now, assuming the neurons are all in close proximity and their interconnections are random but dense enough such that there is at least one path between any two neurons, the central idea of the model may be stated thus:

Let,

E(t) be the proportion of excitatory cells firing per unit time at instant *t*

I(t) be the proportion of inhibitory cells firing per unit time at instant *t*

The value of these functions at time $t+\tau$ is equal to the proportion of cells that receive at least threshold stimulation *and* the number of cells that are not refractory (i.e. sensitive) at time *t*.

Let *C(w)* be the synaptic distribution function (number of synapses per cell). If all cells have the same threshold θ and *x(t)* is the average excitation per synapse, then cells with at least $\theta/x(t)$ number of synapses will be expected to fire. Define the population response function *S*, giving the expected proportion of cells in a population receiving at least threshold excitation per unit time as a function of average levels of excitation, as

$$S(x) = \int_{\theta/x(t)}^{\infty} C(w)dw \tag{1.3}$$

We require a function for each subpopulation type (i.e. excitatory or inhibitory). Under the assumption of homogenous cell types, *C(w)* can be expected to be a unimodal distribution and hence (1.3) assumes a sigmoidal form. The qualitative results obtained by Wilson and Cowan are independent of the exact form of (1.3) except that it should attain a minimum of 0 and maximum of 1 as $x \to -\infty$ and $x \to +\infty$ respectively. Next the proportion of cells that are sensitive (e.g. for the excitatory population) may be written as

$$1 - \int_{t-r}^{t} E(\tau) d\tau \tag{1.4}$$

where *r* is taken to be the absolute refractory period in ms. In general there will be a correlation between the level of excitation of a cell and the probability that it is sensitive but for a richly interconnected population of cells and under the assumption of spatial and temporal fluctuations in the average level of excitation Wilson and Cowan argued that the correlation could be neglected. The resulting integrodifferential equations can be written as

$$\begin{aligned} E(t+\tau) &= \left[1 - \int_{t-r}^{t} E(\tau) d\tau\right] S_e\left(\int_{-\infty}^{t} \alpha(t-\tau)[c_1 E(\tau) - c_2 I(\tau) + P(\tau)] d\tau\right) \\ I(t+\tau) &= \left[1 - \int_{t-r}^{t} I(\tau) d\tau\right] S_i\left(\int_{-\infty}^{t} \alpha(t-\tau)[c_3 E(\tau) - c_4 I(\tau) + Q(\tau)] d\tau\right) \end{aligned} \tag{1.5}$$

where α models the stimulation decay of a cell to its summed input, and *P* and *Q* are external afferent inputs to the excitatory and inhibitory populations respectively. To make equation (1.5) more tractable mathematically, it is necessary to remove the temporal integrals. Wilson and Cowan do this successfully by replacing the integrals with their moving time averages, and performing a Taylor expansion around $\tau = 0$ to yield

$$\begin{aligned} \tau_e \frac{dE}{dt} &= -E + (k_e - r_e E) S_e(c_1 E - c_2 I + P) \\ \tau_i \frac{dI}{dt} &= -I + (k_i - r_i I) S_i(c_3 E - c_4 I + Q) \end{aligned} \tag{1.6}$$

The state $E = 0$, $I = 0$ corresponds to low-level background activity (thus negative values can be admitted). In the absence of input, $P = 0$ and $Q = 0$ is required to be a stable equilibrium and can be achieved if $S(0) = 0$. This is satisfied for any sigmoid by subtracting $S(0)$ from the original function but the result of this must be accounted for, namely that the maximum of the response functions are less than unity, and hence the

introduction of k in (1.6). The Wilson and Cowan model exhibits a number of interesting phenomena: hysteresis loops constituting multiple stable equilibria driven by external inputs (P or Q), damped sinusoidal response to impulses, and limit cycle behaviour for steady state input P with frequency proportional to the input. Functional significance may be associated to these phenomena: short-term memory, evoked responses, and EEG activity respectively. Figure 1.6 illustrates multiple equilibria with hysteresis.

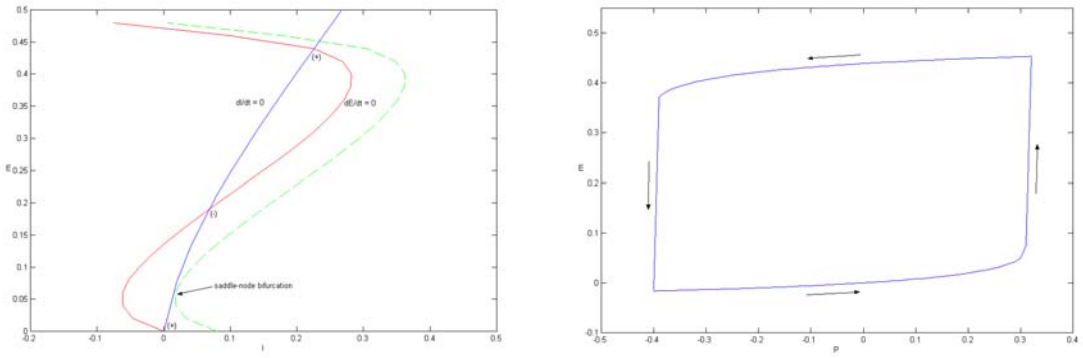

Figure 1.6: Multiple equilibria and hysteresis produced by (1.6). Parameters: $c_1 = 12$, $c_2 = 4$, $c_3 = 13$, $c_4 = 11$, $a_e = 1.2$, $a_i = 1$, $\theta_e = 2.8$, $\theta_i = 4$, $r_e = 1$, $r_i = 1$, $\tau = 8$, $P = 0$, $Q = 0$

For $P = 0$, the nullclines intersect to form three equilibrium points: two stable and one unstable (Figure 1.6, left). As P is increased, the $dE/dt = 0$ nullcline moves to the right and for approximately $P = 0.32$, a stable and an unstable equilibrium coalesce and disappear in a saddle-node bifurcation (GUCKENHEIMER and HOLMES, 1983). Figure 1.6 right illustrates the hysteresis effect of the steady state values of E as P is varied. Figure 1.7 illustrates limit cycle behaviour for a certain set of parameters.

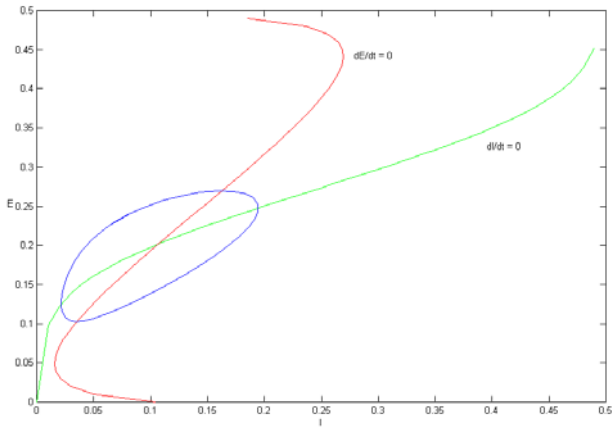

Figure 1.7: Limit cycle produced by (1.6). Parameters: $c_1 = 16$, $c_2 = 12$, $c_3 = 15$, $c_4 = 3$, $a_e = 1.3$, $a_i = 2$, $\theta_e = 4$, $\theta_i = 3.7$, $r_e = 1$, $r_i = 1$, $\tau = 8$, $P = 1.25$, $Q = 0$

Finally, Figure 1.8 illustrates a damped oscillatory response to a brief stimulating pulse.

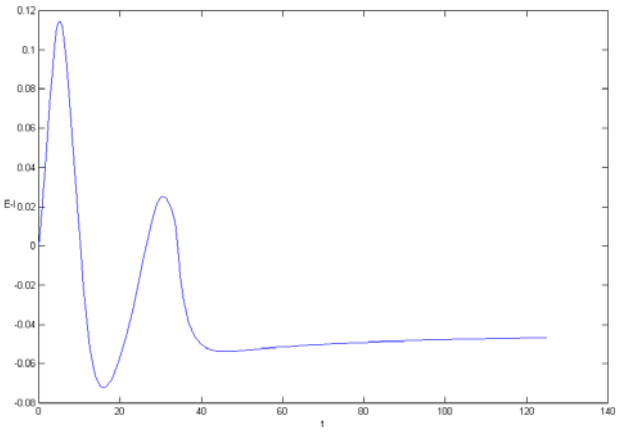

Figure 1.8: Impulse response produced by (1.6).
Parameters: $c_1 = 15$, $c_2 = 15$, $c_3 = 15$, $c_4 = 3$, $a_e = 1$, $a_i = 2$, $\theta_e = 2$, $\theta_i = 2.5$, $r_e = 1$, $r_i = 1$, $\tau = 10$, $Q = 0$

1.3.2 Freeman Model

Freeman has used linear systems coupled with a nonlinear gain to model dynamic neural interactions occurring in the olfactory bulb and olfactory cortex in animals. The olfactory cortex is studied on the basis that it is the simplest and best-known part of the cerebral cortex (see FREEMAN (2000a) for a recent review). The basic premise of Freeman's work is that pools of neurons in local neighbourhoods share activity, leading to the concept of the cortical column as the macroscopic entity for behaviour related function and indeed modelling. A distinction between modelling at the single neuron and at the neural population scale is made (FREEMAN, 1975). Neurons have two main state variables: dendritic potential (wave amplitude) and axon pulse frequency. The dendritic wave amplitude is converted to pulse frequency at the neuron trigger zones and pulse frequency conversion to current amplitude occurs at the synapses. Figure 1.9 illustrates the conversion characteristics for a single neuron.

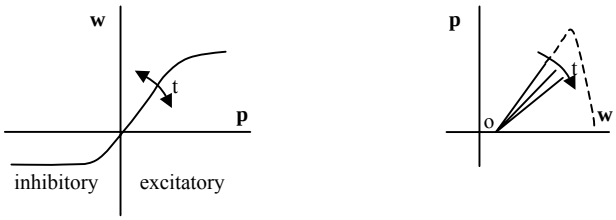

Figure 1.9: Pulse-wave and wave-pulse conversion characteristics for a single neuron

The pulse-wave characteristic is time varying and nonlinear and the wave-pulse characteristic curve is also time-varying but linear above a threshold. The dashed line corresponds to the maximum pulse rate restricted by the absolute refractory period at which cathodal block occurs (FREEMAN, 1975).

For neural populations, Freeman indicates the relevant state variables as the dendritic current density over the population and the axonal pulse density over the population – see Figure 1.10. The pulse-wave interaction is constrained to a small-signal linear region because its domain of input is bounded by the preceding wave-pulse conversion in feedback loops; the gain varies with learning. The wave-pulse characteristic is sigmoid in shape and displaced to the left indicating that for no input some neurons on average are firing; the gain is increased by arousal.

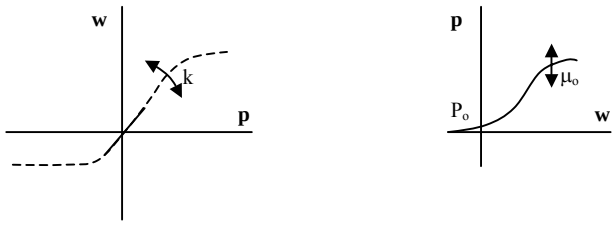

Figure 1.10: Pulse-wave and wave-pulse conversion characteristics for a neural population

Freeman characterises neural populations by the type of internal synaptic connection. A K0 set is one without internal connections e.g. sensory receptor cells. A KI set refers to a population with mutual synaptic connections all having the same sign thus a KI_e is an excitatory population and a KI_i is an inhibitory population. Freeman suggests that negative feedback leads to sustained oscillations and denotes a reciprocally connected KI_e and KI_i set as a KII set.

To create a model of the EEG, Freeman uses a second order linear equation for each population given by

$$\begin{aligned} &\frac{1}{B}\left(\frac{d^2}{dt^2}x(t) + A\frac{d}{dt}x(t) + Bx(t)\right) = P(t) \\ &A = a + b, \\ &B = ab \end{aligned} \tag{1.7}$$

for some constants a, b and where P is an external input. Neural populations are coupled via the product of a constant term (linear gain in Figure 1.10 left) and a

sigmoid term (Figure 1.10 right). The sigmoid curve was found experimentally to give the best fit with data (FREEMAN, 1979) and is given by

$$
\begin{aligned}
&P = P_o(1 + 1 - \exp[-(e^v - 1)/Q_m]) && v > -u_o \\
&P = 0 && v \le -u_o \\
&u_o = -\ln[1 - Q_m \ln(1 + 1/Q_m)]
\end{aligned} \tag{1.8}
$$

The parameter Q_m is used to fit (1.8) to experimental data and is reduced in sleep and anaesthesia, decreasing to 0 for very deep anaesthesia, and is increased by waking and emotional arousal. For very deep anaesthesia, the system may be considered open loop and thus synaptic interactions are suppressed. Under these conditions the impulse response subject to an electrical stimulus may be measured experimentally and (1.7) may be fitted to the data. To study oscillatory interactions, a KII set may be constructed comprising an excitatory and inhibitory population each described by (1.7) coupled with a nonlinear function (1.8) scaled by a constant term (pulse-wave gain)

$$
\begin{aligned}
&\frac{1}{B}\left(\frac{d^2}{dt^2}m(t) + A\frac{d}{dt}m(t) + Bm(t)\right) = P(t) + K_{gm}Q(g(t)), \quad K_{gm} < 0 \\
&\frac{1}{B}\left(\frac{d^2}{dt^2}g(t) + A\frac{d}{dt}g(t) + Bg(t)\right) = K_{mg}Q(m(t)), \quad K_{mg} > 0
\end{aligned} \tag{1.9}
$$

In Freeman's models, *m* denotes excitatory mitral cells (KI_e) of the olfactory cortex and *g* denotes the inhibitory granule cells (KI_i). Equation (1.9) is capable of producing damped and sustained oscillations. More complex topologies e.g. FREEMAN (1993) are capable of exhibiting chaotic time series. Chaotic activity has been proposed by FREEMAN (1992) as a mechanism for information processing in the olfactory cortex. For a system (e.g. the cortex) to adapt to rapid changes in the environment, it must be intrinsically unstable and subject to sudden transitions, i.e. microscopic sensory events should be able to trigger macroscopic patterns. A proposed method for pattern recognition is suggested based on a global chaotic attractor where a recognised sensory input constrains the system to a wing of the attractor. In this way, access to and from the attractor wing is rapid and does not require a large expenditure of energy. Freeman also indicates a role for chaotic attractors consisting of multiple patterns in Hebbian[2] training individual synapses in response to repeated stimuli.

[2] The Hebbian paradigm refers to a neural network learning strategy: an increase in synapse strength occurs as a result of correlated pre- and postsynaptic activity (BOWER and BEEMAN, 1998)

1.3.3 Zetterberg Model

ZETTERBERG (1969) introduced a popular empirical model of the EEG based on the concept of flat spectrum noise feeding a filter network. This model may be conveniently reformulated into a parametric model such as the Auto Regressive (AR) or Auto Regressive Moving Average (ARMA) model (LJUNG, 1987). Empirical models of this nature for the EEG possess a number of desirable features (ZETTERBERG, 1977):

1) Achieves a precise description of essential properties with few parameters
2) Accurate spectral estimation is possible
3) Changes in spectral properties may be tracked
4) Transients and nonstationary events may be detected
5) Classification of the EEG may be performed (see Chapter VI)

Figure 1.11 illustrates Zetterberg's original proposal, which is the ARMA model.

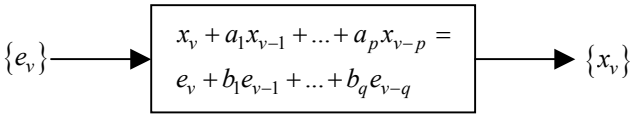

Figure 1.11: ARMA model of the EEG where e_v are uncorrelated with mean 0 and variance σ^2

A popular simplification occurs when $b_i = 0$, resulting in the AR model. Many algorithms for calculating the coefficients efficiently exist (ZETTERBERG, 1977; LJUNG, 1987; PARDEY *et al,* 1996). Rephrased in the frequency domain, the spectral density of the ARMA model may be written ($B(z) = 1$ for the AR model) as

$$R_x(\omega) = \sigma^2 \left[\frac{B[\exp(-j\omega T)]}{A[\exp(-j\omega T)]} \right]^2$$

where

$$A(z) = 1 + a_1 z + \ldots + a_p z^{-p}$$
$$B(z) = 1 + b_2 z + \ldots + b_q z^{-q}$$

and we have used $z = e^{j\omega T}$ where T is the sampling interval.

Zetterberg went one step further by imposing some structure about the underlying stochastic process (ZETTERBERG, 1969; WENNBERG and ZETTERBERG, 1971). It is based on the decomposition of the spectral density into three components consisting of

the delta, alpha, and beta rhythms denoted by $R_\delta(f)$, $R_\alpha(f)$, and $R_\beta(f)$ respectively, i.e.

$$R(f) = R_\delta(f) + R_\alpha(f) + R_\beta(f)$$

This decomposition is illustrated in Figure 1.12. A first order model is used to produce the delta component, while the alpha and beta components each require a second order model. Each component is defined by its centre frequency *f*, bandwidth σ, and power *G* defined by

$$G = \int_{-\infty}^{\infty} R(f)df$$

and thus a fifth order spectrum corresponds to eight parameters. The procedure for determining the parameters in Figure 1.12 from a decomposition of the ARMA model spectrum is called Spectral Parameter Analysis (SPA). The SPA technique has been employed in a number of experimental settings e.g. WENNBERG and ZETTERBERG (1971) and ISAKSSON and WENNBERG (1975).

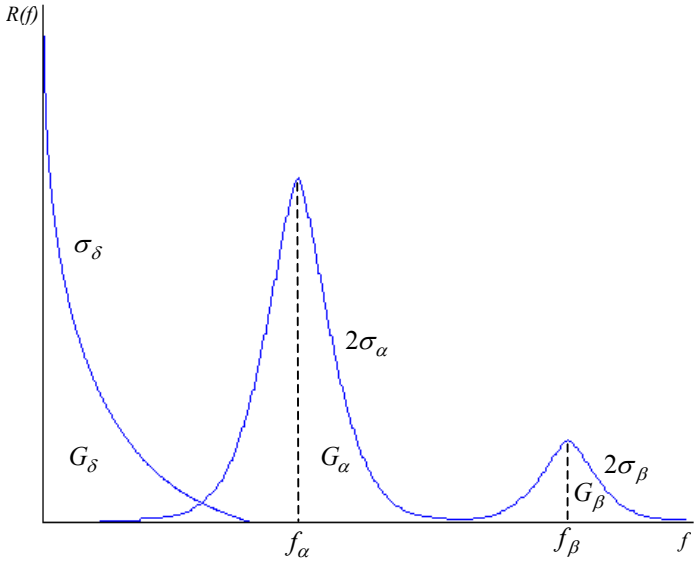

Figure 1.12: Spectral decomposition of the EEG into delta, alpha, and beta components resulting in eight parameters

1.4 Chaotic vs. Stochastic Dynamics

A pertinent concern in both modelling and analysing the EEG is whether the observed aperiodic activity is a manifestation of a deterministic, possibly chaotic attractor or some equilibrium state under random perturbation. Since neuronal systems involve nonlinear mechanisms at the microscopic level, it is natural to hypothesise that the unpredictable nature of macroscopic activity observed in the EEG might be ascribed to

these same underlying nonlinear effects. Characterising the dynamical properties of the EEG has both practical and theoretical implications. Knowledge of the underlying dynamics can assist in the identification of signal processing algorithms for direct brain interfacing and diagnostic purposes. Secondly, it can put constraints on mathematical properties used to model brain dynamics. Nonlinear time series analysis (see ABARBANEL (1996) or KANTZ and SCHREIBER (1997) for a comprehensive introduction) was first applied to EEG data by BABLOYANTZ (1985). Correlation dimension and Lyapunov exponents were calculated for the awake state and various sleep stages. For sleep stage 2 and stage 4, the author found evidence for nonlinear determinism. BABLOYANTZ and DESTEXHE (1986) carried out a dimensional analysis of petit mal epileptic EEG and found a low dimension, which in addition to a positive Lyapunov exponent, they suggested was indicative of chaos. These early nonlinear time series studies of the EEG helped spur a subsequent plethora of related publications - see BAŞAR (1990), PRITCHARD and DUKE (1992), and LEHNERTZ *et al* (1999) for publications dedicated to the subject.

Robust evidence for deterministic nonlinearity, however, has not been shown conclusively for the normal EEG. A comparison of values for invariant measures such as attractor dimension have shown little consistency over values reported by different labs and this has been confounded by the fact that many early studies did not specify under what conditions the measurements were made (ELBERT *et al*, 1994). More recent studies have imposed robust statistical frameworks to aid investigation. When testing for nonlinear structure in the EEG, a popular null hypothesis is that the signal is a realisation of a linear Gaussian random process. Unlike deterministic data, which contains information pertaining to time direction, linear Gaussian random processes do not and may be characterised fully by their power spectrum. It is possible to generate surrogate data exhibiting the same spectrum as the data in question but with the deterministic information destroyed by taking the Fourier transform, randomising the phases, and returning the data to the time domain by an inverse Fourier transform (THEILER *et al*, 1992). Employing these more objective tests, and accounting for correlations in over-sampled data yielding spurious low dimension measures, THEILER and RAPP (1996) concluded that they could find no evidence for low-dimensional chaos. In fact, SCHREIBER (1999) suggests that the problem of whether the observed irregularity in the EEG can be ascribed to intrinsic instability, the large number of

neurons, or noise, may not be decidable based on time series data alone. The author presents a number of spurious (Type I) rejections of the null hypothesis from artificially generated data. STAM *et al* (1999) found the null hypothesis for the eyes-closed relaxed EEG could only be rejected in 1.25% of cases measured over 60 subjects using nonlinear time series prediction. Analysis of pathophysiological data, however, appears to exhibit stronger deterministic tendencies. While PIJN *et al* (1991) could not differentiate normal EEG data from phase-randomised surrogates for data from rats, epileptic data in many cases could. CASDAGLI *et al* (1997) found statistically significant nonlinearities in invasive EEG recordings from patients with temporal lope epilepsy. Similarly, VAN PUTTEN and STAM (2001) concluded that the interictal EEG activity seen in West syndrome could not be described accurately by a Gaussian linear stochastic process.

On the basis of the literature over the last 15 years on nonlinear dynamical analysis of the EEG, we are probably now at a point where we can say with reasonable confidence that the EEG shows little or no evidence of exhibiting chaotic signatures. The signal does appear to exhibit weakly nonlinear characteristics under certain conditions (nonlinear determinism is a necessary but not sufficient condition for chaos) but is unlikely to be due to noise-free low dimensional chaos dynamics (ROMBOUTS *et al*, 1995; STAM *et al*, 1999). FREEMAN (2000b) suggests the name "stochastic chaos" for aperiodic brain activity while SCHREIBER (1999), HERNÁNDEZ (1996), and this author prefer the interpretation of aperiodicity as stochastic limit cycle behaviour. Chapters II, III, and IV of this thesis are concerned with the development and study of a stochastic limit cycle model of the EEG, essentially a hybrid of deterministic and stochastic mathematical properties. It should be stressed that interpretation of the EEG as a stochastic signal does not exclude the use of nonlinear techniques - only their interpretation. A particular case in point is LEHNERTZ and ELGER (1998) who interpret the correlation dimension as an "operational definition" and show its capability in predicting epileptic seizures several minutes prior to the event. We review stochastic interpretations of some of the results from nonlinear (deterministic) time series analysis in Chapter IV. Finally, although an optimum description of the EEG might involve statistical terms, these need not imply that the biophysical processes underlying the EEG are random, only that they may have such a high degree of complexity that only a statistical description is possible (LOPES DA SILVA *et al*, 1999).

1.5 Summary

We commence our technical exposition in Chapter II by introducing nonlinear dynamics and oscillations in the context of Lyapunov theory. We propose a geometric method for synthesising prescribed limit cycle oscillators and subsequently prove asymptotic stability of the limit cycles. In Chapter III, the prescribed oscillators are extended to the stochastic case by employing Itô calculus to produce stochastic limit cycle oscillators. A stochastic limit cycle is defined as an invariant set by applying stochastic stability theory. Both purely mathematical- and biophysical-based models of the EEG are created based on stochastic oscillators. Chapter IV undertakes a study of the phase space embedding of the actual EEG, and invariant measures are compared and contrasted to those obtained from model data. In Chapter V, more detailed systems are considered, employing coupled oscillators to study the interactions of neurons and neural populations as a basis for Evoked Potentials (EP) generation – an important feature of the EEG with application to DBIs. Chapter VI investigates some practical applications of novel direct brain interfaces. Finally, we conclude the thesis with a discussion in Chapter VII.

It is probably fair to say that the work encompassed in a PhD thesis is never complete, with new avenues presenting themselves for study at each juncture. This thesis is no different, and the final sections of Chapters II – VI inclusive are concerned with areas that the author feels warrants further research but for which there was insufficient time to study in detail in the present scope of work.

CHAPTER II

DYNAMICAL SYSTEMS AND PRESCRIBED OSCILLATORS

2.1 Introduction

This part of the thesis provides the technical foundations for later chapters. Nonlinear dynamical systems are introduced in the context of Lyapunov theory (MERKIN, 1997; LA SALLE and LEFSCHETZ, 1961). We propose a novel method for synthesising geometrically prescribed limit cycle oscillators. Asymptotic stability is subsequently proved by invoking La Salle's extension to *Lyapunov's Direct Method* (LA SALLE and LEFSCHETZ, 1961). Examples are applied to well known planar curves to demonstrate the flexibility of the technique. Finally, a method for approximating a certain class of periodic time series via limit cycle oscillators is presented.

Nonlinear dynamics is a complex field. Analytic solutions are uncommon and complete theories rare; we resort to a combination of geometric, numerical, and perturbation techniques. A number of excellent texts exist on the topic of nonlinear dynamics and oscillations including LA SALLE and LEFSCHETZ (1961), ANDRONOV *et al* (1966), MERKIN (1997), and GUCKENHEIMER and HOLMES (1983).

2.2 Basic Dynamical System Theory

We will regard a dynamical system as

$$\dot{x} = F(x) \tag{2.1}$$

where $x = x(t) \in R^n$ is a vector valued function of time and $F : U \to R^n$ is a smooth function defined on some subset $U \subseteq R^n$. We say that the *vector field* F generates a *flow* $\phi_t : U \to R^n$ where $\phi(x,t)$ is defined for all x in U and t in some interval $I = (a,b) \subseteq R$ and ϕ satisfies

$$\frac{d}{dt}(\phi(x,t))\,|_{t=\tau} = F(\phi(x,\tau)) \tag{2.2}$$

for all $x \in U$ and $\tau \in I$. Note that we have restricted t to an interval I since it is possible for solutions to have a *finite escape time* e.g. $dx/dt = 1 + x^2$ has solutions $x(t) = \tan(t+c)$ that tend to infinity on unrestricted time intervals. In this chapter, we are

interested in synthesising systems exhibiting a prescribed periodic orbit $\phi(x_o,t)$ subject to some initial condition $\phi(x_o,0) = x_o$.

2.2.1 Stability Theory

Fixed points are defined by the vanishing vector field

$$F(x_e) = 0$$

Without loss of generality we shall consider the fixed point x_e to be at the origin. We wish to study the stability of this equilibrium point subject to small perturbations. Writing $x(t) = (x_1,x_2,\ldots,x_n)$ and denoting initial conditions at time t_0 as x_{0j} we introduce the concept of stability *in the sense of Lyapunov*: If for any positive value ε no matter how small, one can find another positive value δ such that at time $t = t_0$, for all initial perturbations x_{0j} satisfying

$$\sum x_{0j}^2 \leq \delta$$

the inequality

$$\sum x_j^2 < \varepsilon$$

holds, then the unperturbed motion is stable; otherwise it is unstable. A geometric interpretation ($n = 2$) is illustrated in Figure 2.1. Given an arbitrarily small sphere of radius $\sqrt{\varepsilon}$ (called the ε-sphere) if we can find a smaller sphere of radius $\sqrt{\delta}$ (the δ-sphere) such that a point starting within (or on) the δ-sphere never reaches the ε-sphere then the system is stable.

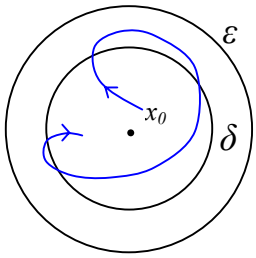

Figure 2.1: Geometric interpretation of stability in the sense of Lyapunov

If $x_j(t) \to 0$ as $t \to \infty$ then the equilibrium is called *asymptotically stable.* If we can choose x_{0j} arbitrarily and asymptotic stability holds then we say the system *is stable in the large* or *globally asymptotically stable*.

Before introducing Lyapunov's Direct Method, we need to introduce definite functions, which are used in the method's definition and supplied proof. We consider

real, single-valued, continuous functions $V(x) = V(x_1,x_2,\ldots,x_n)$. A *definite* function (positive or negative) is a function that maintains one sign and vanishes only at the origin. An *indefinite* function is one that vanishes at point(s) other than the origin. To provide a geometric interpretation of V and its time derivative, we construct the surface $V(x) = c$ (c positive) and choose an arbitrary point M on it. We evaluate the gradient at M (which lies perpendicular to the surface $V(x) = c$) as (where e_i are unit vectors along the axis x_i)

$$\text{grad } V = \frac{\partial V}{\partial x_1} e_1 + \frac{\partial V}{\partial x_2} e_2 + \ldots + \frac{\partial V}{\partial x_n} e_n$$

Calculating the time derivative of V

$$\dot{V} = \frac{\partial V}{\partial x_1}\dot{x}_1 + \frac{\partial V}{\partial x_2}\dot{x}_2 + \ldots + \frac{\partial V}{\partial x_n}\dot{x}_n \tag{2.3}$$

we can write

$$\dot{V} = U \cdot \text{grad } V \tag{2.4}$$

where U is the velocity of the point M (whose components are dx_i/dt). From (2.4) we can see that if $dV/dt < 0$, the angle between U and *grad* V is obtuse, if $dV/dt > 0$ the angle is acute, and if $dV/dt = 0$ the angle between the two vectors is a right angle. Figure 2.2 provides an illustration of this.

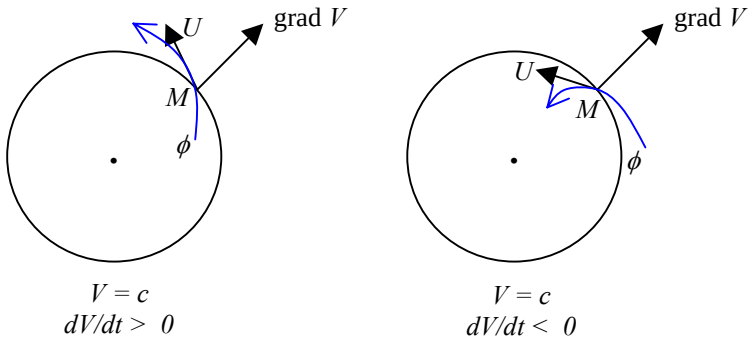

Figure 2.2: Geometric interpretation of (2.4)

Analysis of stability can often be performed by Lyapunov's Direct Method, which can provide sufficient conditions for the stability of the system. The Direct Method is concerned with generalisations of energy functions V with stability being proven for systems where V is shown to decrease along solution curves. The real benefit of this approach is that it avoids the need to obtain explicit solutions of the system in question (which is rarely possible in the nonlinear case). We present Lyapunov's Direct Method in the following two sections.

2.2.1.1 Lyapunov's Theorem of Motion Stability

Theorem 2.1 If one can find a definite function V for the set of differential equations (2.1) such that the time derivative dV/dt is either identically zero or semi-definite with opposite sign, then the motion is stable.

Proof

Construct an arbitrary small sphere of radius $\sqrt{\varepsilon}$ containing a smaller sphere of radius $\sqrt{\delta}$.

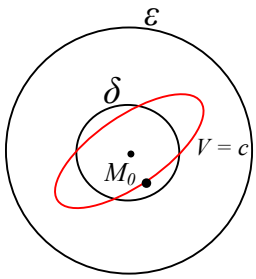

Figure 2.3: Theorem 2.1 – see text for details

We select a Lyapunov contour $V = c$ as illustrated in Figure 2.3 such that we can start a point M in motion at M_0 on the contour which is within the δ-sphere. To prove the theorem we show that the point M_0 never reaches the ε-sphere. At the instant we set M in motion we note that, from the theorem definition, $dV/dt \leq 0$. From (2.2) we see that if $dV/dt < 0$ the angle between the velocity vector U and grad V (normal to the $V = c$ surface) is obtuse and thus the M must be directed inwards. When $dV/dt = 0$ the angle is perpendicular and M moves tangent to the Lyapunov contour. Therefore, M can never leave the closed area bounded by $V = c$ and thus never reach the ε-sphere. This proves the theorem.

2.2.1.2 Lyapunov's Theorem of Asymptotic Stability

Theorem 2.2 If one can find a definite function V for the set of differential equations (2.1) such that the time derivative dV/dt is definite with opposite sign, then the motion is *asymptotically* stable.

Proof

Following from the proof of Theorem 2.1 we note that if the time derivative of V is strictly less than zero, then M must always be directed inwards. Selecting a sequence

of decreasing contours $V = c_j$, where $c_j < c_k$, $k < j$, we verify that in the limit of $c_j \to 0$, $M \to 0$ and the system is thus asymptotically stable. This concludes the proof.

In practice the conditions for asymptotic stability laid out in Theorem 2.2 are quite strict, requiring definiteness in both V and dV/dt. These requirements can be relaxed with the use of Krasovsky's theorem, which we state next without proof.

2.2.1.3 Krasovsky's Theorem of Asymptotic Stability

Theorem 2.3 If one can find a definite function V for the set of differential equations (2.1) such that the time derivative:

- $dV/dt < 0$ outside of K
- $dV/dt = 0$ on K

where K is a manifold not containing whole trajectories of the system for $0 \leq t < \infty$, then the motion is asymptotically stable.

Figure 2.4 illustrates an example of Krasovsky's Theorem. The system is asymptotically stable, however dV/dt is only strictly negative outside K.

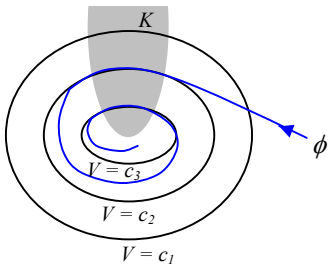

Figure 2.4: Krasovsky's Theorem

While Lyapunov's Theorem requires dV/dt to be a monotonic decreasing function, here we only require that it decrease in a piecewise fashion. A useful generalisation of the theorem is stated next.

2.2.1.4 Barbashin-Krasovsky's Theorem of Asymptotic Stability in the Large

Theorem 2.4 If one can find a definite function V for the set of differential equations (2.1) where V is radially unbounded, that is

$$\lim_{x_1, x_2 \to \infty} V(x_1, x_2) \to \infty$$

(interpreted as at least one $x_i \to \infty$ along any path) such that the time derivative:

- $dV/dt < 0$ outside of K
- $dV/dt = 0$ on K

where K is a manifold not containing whole trajectories of the system for $0 \leq t < \infty$, then the motion is *globally* asymptotically stable.

2.2.2 Stability of Periodic Motion

Before discussing the stability of periodic motion we will introduce some concepts related to asymptotic behaviour of flows. A discussion on asymptotic behaviour and stability in the sense of Lyapunov for periodic orbits precedes an extension of Lyapunov's Direct Method, which facilitates analysis of convergence to dynamical behaviours more general than equilibrium, i.e. limit cycles.

2.2.2.1 Asymptotic Behaviour

Definition 2.1 An *invariant set S* for a flow ϕ_t is a subset $S \subset \mathbb{R}^n$ such that $\phi_t(x) \in S$ for $x \in S$ for all $t \in \mathbb{R}$.

In other words, every trajectory that starts in S remains in S for all future time. An invariant set is thus a generalisation of the equilibrium point, for example.

Definition 2.2 A point p is an *ω-limit point* of x if there are points $\phi_{t_1}(x), \phi_{t_2}(x), \ldots$ on the orbit of x such that $\phi_{t_i}(x) \to p$ as $t \to \infty$. A point q is an *α-limit point* of x if there are points $\phi_{t_1}(x), \phi_{t_2}(x), \ldots$ on the orbit of x such that $\phi_{t_i}(x) \to q$ as $t \to -\infty$. The *α(x)* and *ω(x) limit sets* are naturally defined as the sets of α and ω limit points of x for dynamical behaviours more general than equilibrium

A limit set describes where the flow tends to with infinite time. If *x(t)* is bounded, then its *ω-limit set* is a nonempty, compact, invariant set (LA SALLE and LEFSCHETZ, 1961).

Definition 2.3 A point p is *nonwandering* for the flow ϕ_t if for any neighbourhood U of p, there exists arbitrarily large t such that $\phi_t(U) \cap U \neq \varnothing$. A set of all such points is a *nonwandering set*.

For two-dimensional planar flows, all possible nonwandering sets fall into three classes (ANDRONOV *et al*, 1966)

(i) Fixed points

(ii) Closed orbits

(iii) Unions of fixed points and the trajectories joining them

With reference to the third case, a trajectory joining two different fixed points is called a *heteroclinic* trajectory and one connecting the point to itself is called a *homoclinic* trajectory. A useful result (two dimensional flows only) for determining the existence of closed orbits is the *Poincaré-Bendixson* theorem (stated here without proof)

Theorem 2.5 *(Poincaré-Bendixson)* For two-dimensional flows, a nonempty compact α- *or* ω- limit set of planar flow, which contains no fixed points, is a closed orbit.

To verify the existence of a stable periodic orbit, one may determine two nested contours C_1 and C_2 such that trajectories cross C_1 outwards and C_2 inwards.

2.2.2.2 Stability in the sense of Lyapunov

In addition to the existence of a limit cycle, we generally require information regarding its stability. Analogous to the study of the stability of fixed points, we can apply Lyapunov stability theory to the study of period orbits (ANDRONOV *et al*, 1966). We observe a representative point circulating the closed phase path and enclose the point in a small sphere of radius $\sqrt{\varepsilon}$, which moves with the representative point. If for arbitrarily small ε, we can find a smaller sphere of radius $\sqrt{\delta}$ such that every representative point situated within the δ -sphere at the initial instant never reaches the ε-sphere, then the period motion is stable in the sense of Lyapunov. More specifically, consider (2.1) for a planar system exhibiting periodic motion $x_1 = \phi_1(t)$, $x_2 = \phi_2(t)$ with period T. Stability is assured if for every $\varepsilon > 0$, we can find a $\delta(\varepsilon)$ such that for any *other* motion $x_1(t)$, $x_2(t)$ satisfying

$$\left|x_1(t_0) - \phi_1(t_0)\right| \le \delta \quad and \quad \left|x_2(t_0) - \phi_2(t_0)\right| \le \delta$$

the inequalities

$$\left|x_1(t) - \phi_1(t)\right| < \varepsilon \quad and \quad \left|x_2(t) - \phi_2(t)\right| < \varepsilon$$

are satisfied for all $t > t_0$.

We conclude this section by stating with proof, an extension to Lyapunov's Direct Method to allow for negative semi-definite functions $\dot{V}$ while still facilitating conclusions to be drawn regarding asymptotic stability (Theorem VI, pp. 58, LA SALLE and LEFSCHETZ, 1961). By doing this, we will be able to consider stability of more complex dynamic behaviours, in particular periodic motions.

Theorem 2.6 *(Theorem VI, pp. 58, LA SALLE and LEFSCHETZ, 1961)* Let *V(x)* be a scalar function with continuous first partial derivatives. Let Ω_l designate the region where $V(x) < l$. Assume that Ω_l is bounded and that within Ω_l:

$V(x) > 0 \quad for \quad x \neq 0,$
$\dot{V}(x) \leq 0$

Let *R* be the set of all points within Ω_l where $\dot{V}(x) = 0$ and let *M* be the largest invariant set in *R*. Then every solution *x(t)* in Ω_l tends to *M* as $t \to +\infty$.

Proof

The conditions on *V* imply that it is non-increasing for $t \to +\infty$ and thus every solution *x(t)* starting in Ω_l must remain in it. Therefore *V(t)* has a limit l_0 as $t \to +\infty$ and $l_0 < l$. By continuity, one concludes that $V(x) = l_0$ is on the ω-limit set of *x*. Hence the ω-limit set is in Ω_l and $\dot{V} = 0$ is on the ω-limit set. Consequently, the ω-limit set is in *R*, and since it is an invariant set it is in *M*. Since *x(t)* remains in Ω_l it is bounded for $t \geq 0$, and $x(t) \to M$ as $t \to +\infty$. If in addition *V* is radially unbounded and $\dot{V}(x) \leq 0$ everywhere, then all solutions globally asymptotically converge to *M* as $t \to +\infty$. This concludes the proof.

2.3 Prescribed Limit Cycle Oscillators

2.3.1 Method I

We are now ready to present our main result for this chapter, which is a novel method for synthesising two-dimensional oscillators exhibiting a prescribed limit cycle contour. Specifically, we seek a system of the form

$$\frac{dx_1}{dt} = f(x_1, x_2)$$
$$\frac{dx_2}{dt} = g(x_1, x_2) \tag{2.5}$$

such that it exhibits an asymptotically stable limit cycle oscillation tracing out a specified contour in the (x_1,x_2) phase plane. A natural system to use as a starting point is the Hamiltonian system (GUCKENHEIMER and HOLMES, 1983)

$$\dot{x}_1 = \frac{\partial H}{\partial x_2}$$
$$\dot{x}_2 = -\frac{\partial H}{\partial x_1} \tag{2.6}$$

where $H(x_1,x_2)$ is the Hamiltonian or energy function whose critical points are the fixed points of (2.6). Since

$$\frac{dH}{dt} = \frac{\partial H}{\partial x_1}\dot{x}_1 + \frac{\partial H}{\partial x_2}\dot{x}_2 = 0$$

the level curves $H(x_1,x_2) = constant$ are solutions of (2.6). A Hamiltonian system is *conservative* (linearising such a system yields only centres or saddles; no sinks or sources). Put another way, the periodic motion is not asymptotically stable in the sense of Lyapunov (ANDRONOV *et al*, 1966) and a small perturbation results in periodic motion that will not return to the original motion. By introducing additional terms to the right hand side of (2.6) we will show that it is possible to produce asymptotically stable limit cycle behaviour described by the contour $H = 0$. The proposed system takes the form of:

$$\dot{x}_1 = \frac{\partial H}{\partial x_2} - \lambda H \frac{\partial H}{\partial x_1}$$
$$\dot{x}_2 = -\frac{\partial H}{\partial x_1} - \lambda H \frac{\partial H}{\partial x_2}, \quad \lambda > 0 \tag{2.7}$$

Lemma 2.1

Consider a closed, smooth contour $H(x_1,x_2) = 0$ centred on the origin and let $V(x_1, x_2) = H^2(x_1, x_2)$. Then $\left(\partial H / \partial x_1\right)^2 + \left(\partial H / \partial x_2\right)^2 = 0$ is only satisfied at maxima of V.

Proof

Since $\left(\partial H/\partial x_1\right)^2+\left(\partial H/\partial x_2\right)^2=0$ only occurs when the partial derivatives are identically 0, it remains to show that $\partial H/\partial x_i=0$, *i = 1, 2* occur at maximum points of *V*. The critical points of *V* are obtained by letting

$$\frac{\partial V}{\partial x_1}=2H\frac{\partial H}{\partial x_1}=0$$
$$\frac{\partial V}{\partial x_2}=2H\frac{\partial H}{\partial x_2}=0$$

which occurs on the contour *H = 0* and when $\partial H/\partial x_i=0$, *i = 1, 2*. The function *V* is a nonlinear measure of distance from the contour *H = 0* and thus only possesses one minimum on the contour itself and no saddle points. Thus, the other critical points must be maxima. This concludes the proof.

Theorem 2.7 *(Prescribed Limit Cycle Oscillators)* Consider the closed contour defined by $H(x_1,x_2) = 0$ centred on the origin. Assuming the first partial derivatives of *H* exist, then equation (2.7) yields asymptotically stable limit cycle behaviour described by the contour $H(x_1,x_2) = 0$.

Proof

Consider a Lyapunov function, which is a nonlinear measure of distance from the limit cycle

$$V(x_1,x_2)=H^2(x_1,x_2)$$

Applying Theorem 2.6, we choose a bounded region Ω_l that surrounds $H(x_1,x_2)$. The time derivative of *V* is written as

$$\dot{V}=\frac{\partial V}{\partial x_1}\frac{dx_1}{dt}+\frac{\partial V}{\partial x_2}\frac{dx_2}{dt}=2H\left[\frac{\partial H}{\partial x_1}\frac{dx_1}{dt}+\frac{\partial H}{\partial x_2}\frac{dx_2}{dt}\right]$$
$$=-2\lambda H^2\left\{\left(\frac{\partial H}{\partial x_1}\right)^2+\left(\frac{\partial H}{\partial x_2}\right)^2\right\} \qquad (2.8)$$

and is thus negative semi-definite. The region R for which $\dot{V} = 0$ consists of $\left(\partial H / \partial x_1\right)^2 + \left(\partial H / \partial x_2\right)^2 = 0$ and the limit cycle defined by $H(x_1,x_2) = 0$, both of which are invariant sets. Thus M, the largest invariant set in R, is simply the union of these sets. By Theorem 2.6, every point starting in Ω_l tends to M as $t \to \infty$. By Lemma 2.1, since $\left(\partial H / \partial x_1\right)^2 + \left(\partial H / \partial x_2\right)^2 = 0$ correspond to local *maxima* of V, a small deviation in any direction will drive trajectories away from it, and thus are unstable equilibria. In contrast, trajectories on the limit cycle $H(x_1,x_2) = 0$ correspond to a local minimum of V and are stable. Therefore, all trajectories starting from anywhere (excluding the unstable equilibrium points) must converge to the limit cycle and hence it is asymptotically stable. Since V is radially unbounded (it is a nonlinear measure of distance)

$$\lim_{x_1,x_2 \to \infty} H^2(x_1,x_2) \to \infty$$

and if

$$\dot{V}(x_1,x_2) \le 0$$

over the whole state space then the limit cycle is *globally* asymptotically stable. This concludes the proof.

The parameter λ determines the stability of the limit cycle. For λ > 0, the limit cycle is stable, for λ < 0 the limit cycle is unstable, and for λ = 0 it is marginally stable (this is just a reduction to the Hamiltonian system equivalent). For positive λ, its magnitude determines the degree of attraction to the limit cycle, an important factor for models comprised of coupled oscillators. In the sequel, we will refer to (2.7) as *Method I.*

2.3.1.1 Examples of Prescribed Oscillators

To illustrate the flexibility of our method, we present examples of oscillators exhibiting limit cycle contours of some well known planar curves (REKTORYS, 1969). In all cases, we have employed a Runge-Kutta (4,5), variable step-size, explicit solver.

Example 2.1: *Ellipse Limit Cycle Oscillator*

Let

$$H(x_1,x_2) = \frac{x_1^2}{a^2} + \frac{x_2^2}{b^2} - 1 = 0$$

be the standard equation of the ellipse. The parameters a and b are the lengths of the semi-axes. The special case for the circle is obtained by setting $a = b$. The corresponding differential system is

$$\dot{x}_1 = \frac{2x_2}{b^2} - \lambda\left[\frac{x_1^2}{a^2} + \frac{x_2^2}{b^2} - 1\right]\frac{2x_1}{a^2}$$
$$\dot{x}_2 = -\frac{2x_1}{a^2} - \lambda\left[\frac{x_1^2}{a^2} + \frac{x_2^2}{b^2} - 1\right]\frac{2x_2}{b^2} \qquad (2.9)$$

Figure 2.5 illustrates a phase portrait for the simulated system with $a = 2$, $b = 1$, $\lambda = 1$. We see that the *nullclines* intersect at (0,0) to form an *unstable focus.* This singularity is bounded by a stable elliptic limit cycle as initially prescribed.

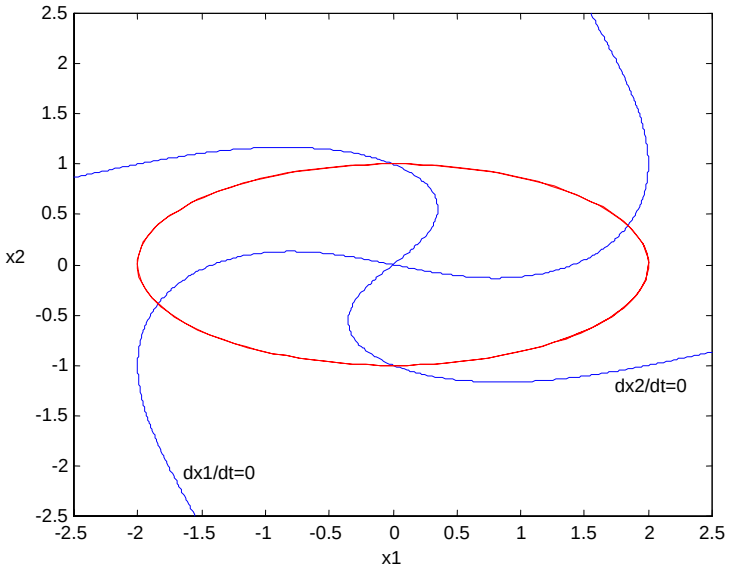

Figure 2.5: Ellipse limit cycle (2.9) with $a = 2$, $b = 1$, $\lambda = 1$

Example 2.2: *Square Limit Cycle Oscillator*

Let

$$H(x_1, x_2) = |x_1| + |x_2| - r^2 = 0$$

describe a square contour circumscribed by a circle of radius r. In calculating the partial derivatives, we write

$$\frac{\partial |x_i|}{\partial x_i} = \mathrm{sgn}(x_i)$$

where sgn denotes the signum function. The corresponding differential system is

$$\dot{x}_1 = \mathrm{sgn}(x_2) - \lambda\left[|x_1| + |x_2| - r^2\right]\mathrm{sgn}(x_1)$$
$$\dot{x}_2 = -\mathrm{sgn}(x_1) - \lambda\left[|x_1| + |x_2| - r^2\right]\mathrm{sgn}(x_2) \qquad (2.10)$$

Figure 2.6 illustrates a phase portrait for the simulated system with $r = 1$, $\lambda = 1$.

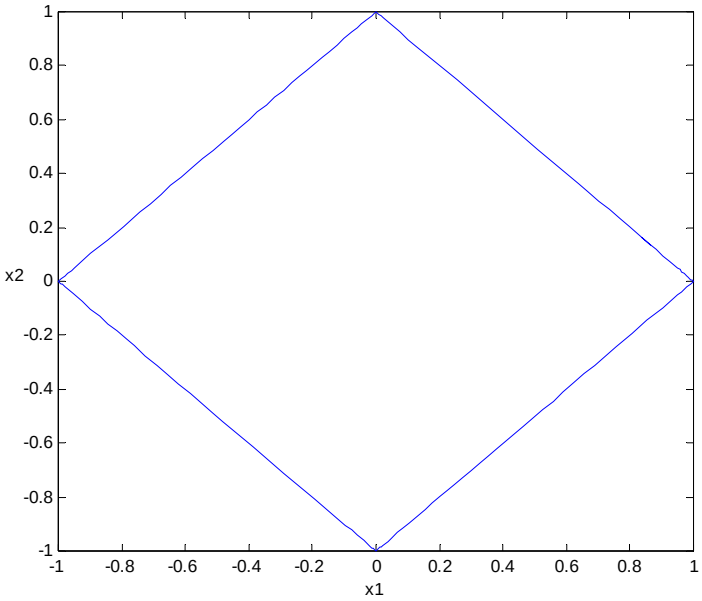

Figure 2.6: Square limit cycle (2.10) with $r = 1$, $\lambda = 1$

Example 2.3: *Epicycloid and Hypocycloid Limit Cycle Oscillators*

A *simple epicycloid* / *hypocycloid* is formed by every point on a circle of radius r rotating over the exterior / interior of a circle of radius m (REKTORYS, 1969). If the distance d between the generating point and the centre of the rotating circle is greater / less than r, the adjective *prolate* / *curate* is invoked. For $r = 0.5m$ we have an example of a simple epicycloid (*nephroid*):

$$H(x_1, x_2) = (x_1^2 + x_2^2 - 4r^2)^3 - 108r^4 x_2^2 = 0 \tag{2.11}$$

Figure 2.7 (red) illustrates a phase portrait for the simulated system with $r = 1$, $\lambda = 1$. Setting $r = m$ yields a simple (cartoid) epicycloid and its phase portrait is also illustrated in Figure 2.7 (blue) with $r = 1$, $\lambda = 1$.

$$H(x_1, x_2) = \left(x_1^2 + x_2^2 - 2rx_1\right)^2 - 4r^2(x_1^2 + x_2^2) = 0 \tag{2.12}$$

Setting $r = m$ and $d > r$ yields the (prolate) epicycloid (*the limaçon of Pascal*)

$$H(x_1, x_2) = \left(x_1^2 + x_2^2 - 2dx_1\right)^2 - 4r^2(x_1^2 + x_2^2) = 0 \tag{2.13}$$

Figure 2.7 (green) illustrates a phase portrait for the simulated system with $r = 1$, d = 2, $\lambda = 1$.

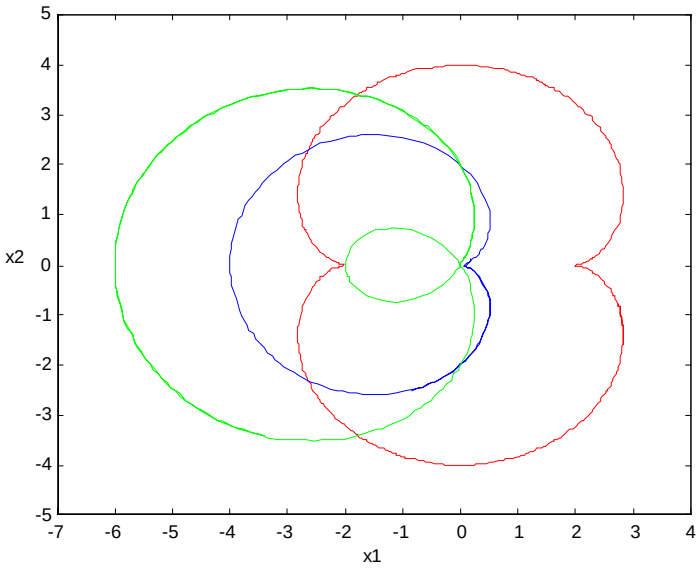

Figure 2.7: Epicycloid and hypocycloid limit cycles (2.11), (2.12), (2.13) with $r = 1$, $\lambda = 1$ in all cases

Note that the trajectories of the limaçon of Pascal limit cycle intersect at the origin to form a saddle node. Finally, a simple hypocycloid (*Steiner's hypocycloid*) follows by setting $r = m/3$ in the following equation

$$H(x_1, x_2) = \left(x_1^2 + x_2^2\right)^2 + 8rx_1\left(3x_2^2 - x_1^2\right) + 18r^2\left(x_1^2 + x_2^2\right) - 27r^4 = 0 \qquad (2.14)$$

and is illustrated in Figure 2.8.

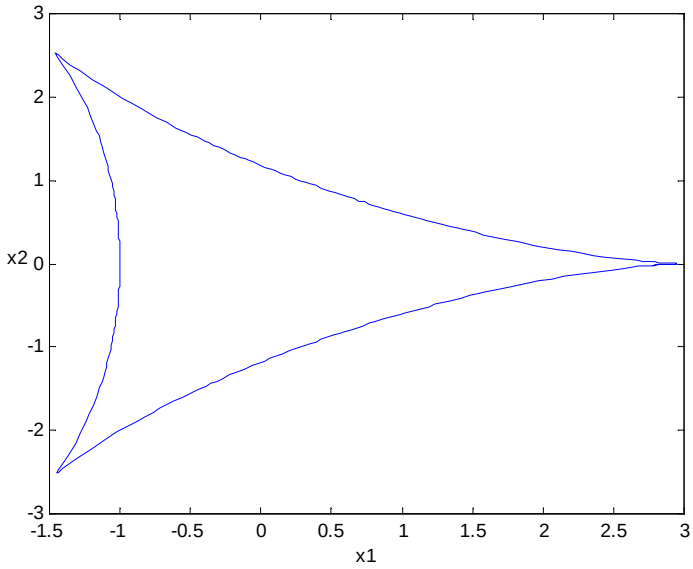

Figure 2.8: Steiner's hypocycloid limit cycle (2.14) with $r = 1$, $\lambda = 1$

Example 2.4: *Cassinian Ovals Limit Cycle Oscillator*

Let

$$H(x_1, x_2) = \left[(x_1 - b)^2 + x_2^2\right]\left[(x_1 - c)^2 + x_2^2\right] - k = 0 \tag{2.15}$$

be the equation for the Cassinian Ovals (REKTORYS, 1969). Figure 2.9 illustrates a phase portrait for the simulated system with $b = 1$, $c = -1$, $\lambda = 1$, $k = 0.5$ (green), $k = 1.0$ (blue), $k = 2.0$ (red).

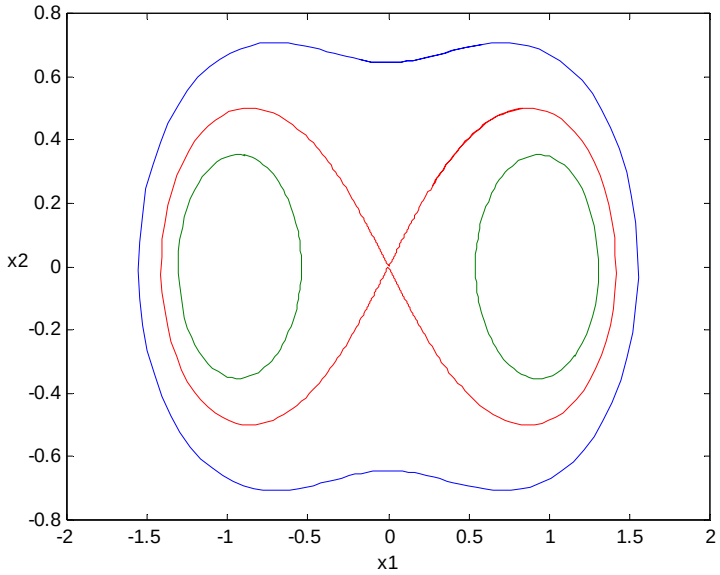

Figure 2.9: Cassinian ovals limit cycle (2.15) with $b = 1$, $c = -1$, $\lambda = 1$, $k = 0.5$ (green), $k = 1.0$ (blue), $k = 2.0$ (red)

Note that for $k = 1.0$ the trajectories cross at a saddle point. As the oscillator traverses in the clockwise direction from the left lobe, the trajectory will either continue on that lobe via a *homoclinic* trajectory, or continue on the outside of the right lobe. The trajectories do not actually pass through the saddle node since the equilibrium forms a separatrix. See Appendix A1 for an electronic circuit realisation of the oscillator.

2.3.1.2 Nested Limit Cycles and Black Holes

It is possible to prescribe second order systems with multiple limit cycles as Example 2.4 demonstrates (for the case where $k = 0.5$). Nested limit cycle oscillators can also be created with multiple stable and unstable limit cycles. Oscillators exhibiting a stable limit cycle surrounding an unstable limit cycle have been successfully employed in modelling cardiac arrest by DE PAOR (1994). In this model, trajectories crossing the unstable separatrix enter a "black hole" and converge to a stable equilibrium. In this

section we show how simple nested limit cycles can be obtained using the prescribed oscillator technique.

Consider two nested circles (for simplicity) given by the equation

$$H = \left(x_1^2 + x_2^2 - r_1^2\right)\left(x_1^2 + x_2^2 - r_2^2\right) = 0 \tag{2.16}$$

with partial derivatives

$$\frac{\partial H}{\partial x_1} = 4x_1\left(x_1^2 + x_2^2 - \frac{r_1^2 + r_2^2}{2}\right)$$

$$\frac{\partial H}{\partial x_2} = 4x_2\left(x_1^2 + x_2^2 - \frac{r_1^2 + r_2^2}{2}\right) \tag{2.17}$$

Inserting (2.16) and (2.17) into (2.7) yields the desired result: an oscillator with nested limit cycles. Figure 2.10 illustrates a simulation for initial conditions inside the inner limit cycle and outside the outer limit cycle. Note the different directions of the flow for the different limit cycles.

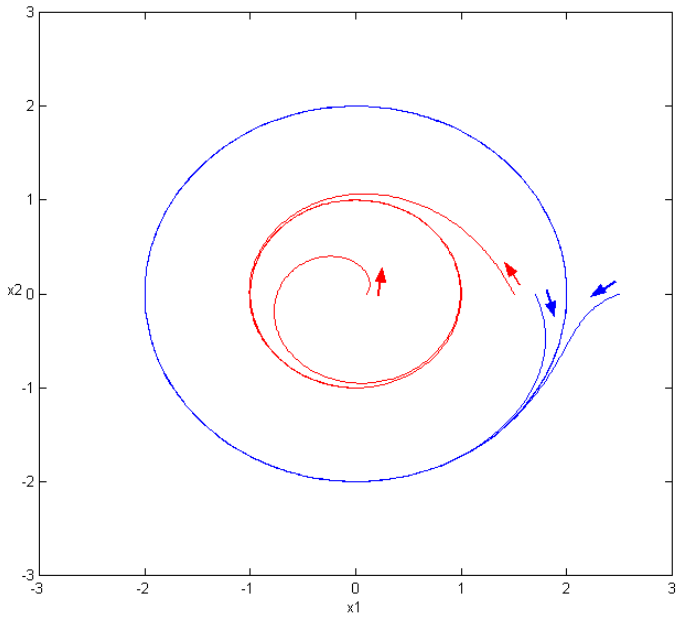

Figure 2.10: Nested (stable) limit cycles
with $r_1 = 1$, $r_2 = 2$, $\lambda = 0.2$

The question naturally arises of where the separatrix lies between the two limit cycles. The equation for the separatrix is obtained from

$$\left(\partial H / \partial x_1\right)^2 + \left(\partial H / \partial x_2\right)^2 = 0$$

and by Lemma 2.1 corresponds to an unstable equilibrium. From (2.17) it can be easily seen that the separatrix for (2.16) is given by a circle of radius $\sqrt{r_1^2 + r_2^2/2}$. Equation

(2.16) can be modified to create a "black hole" by changing the sign of r_1^2. Trajectories inside the separatrix now converge to the origin as illustrated in Figure 2.11.

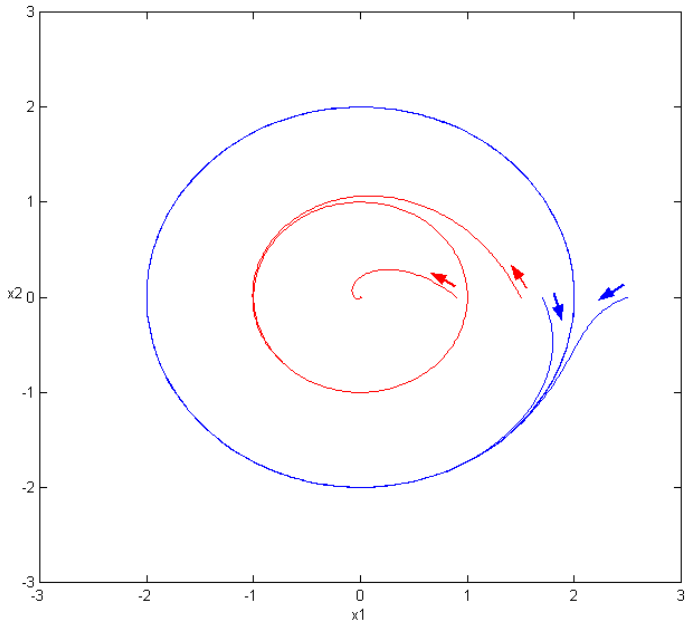

Figure 2.11: Nested (stable and unstable) limit cycles
with $r_1 = 1$, $r_2 = 2$, $\lambda = 0.2$

2.3.2 Method II

For our second method of prescribing limit cycle oscillators, we are interested in approximating a given periodic time series. We require a technique to encode time information into the oscillator. After reviewing the concepts of topological equivalence (GUCKENHEIMER and HOLMES, 1983), we propose a method for specifying a contour by a Fourier series representation of its radius as a function of angle, expressed in Cartesian coordinates. By placing a constraint on the first state variable of the oscillator, it is possible to relate the limit cycle to a contour generated by a limited class of periodic data. We find a system capable of approximating the contour by numerical optimisation and conclude the chapter with an example.

2.3.2.1 Topological Equivalence and Conjugacy

Consider two dynamical systems given by

$$\frac{dx}{dt} = f(x), \quad x \in \mathbb{R}^n$$

$$\frac{dy}{dt} = g(y), \quad y \in \mathbb{R}^n$$

We say they are *topologically equivalent* if there is a homeomorphism (continuous mapping with continuous inverse) $h:\mathbb{R}^n \leftrightarrow \mathbb{R}^n$ taking each orbit $\Phi_t(x)$, $x \in \mathbb{R}^n$ of the first system to an orbit $\Psi_t(y)$, $y \in \mathbb{R}^n$ of the second. By topological equivalence, we have not required that parameterisation by time be preserved, i.e. for any *x* and *t*, there is a t_1, which could differ from *t*, such that

$$h(\Phi_t(x)) = \Psi_{t_1}(h(x))$$

If $t = t_1$, i.e. time parameterisation is preserved, then the equivalence is called *conjugacy*. The time parameterisation of a single state variable in a two-dimensional system is evidently not uniquely defined by the topology of its attractor. To see this consider a two dimensional system defined by

$$\frac{dx_1}{dt} = f(x_1, x_2)$$
$$\frac{dx_2}{dt} = g(x_1, x_2)$$

At any point in state-space, the slope of a trajectory is given by

$$\frac{dx_2}{dx_1} = \frac{g(x_1, x_2)}{f(x_1, x_2)} \tag{2.18}$$

Now, referring to the trajectory illustrated in Figure 2.12, we wish to derive Δt, the (approximate) time taken to go from state *P* to state *Q*.

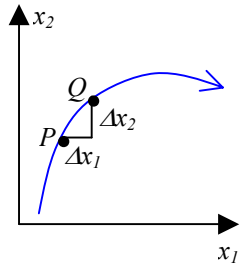

Figure 2.12: Parameterisation by time
– see text for details

By defining,

$$\bar{f}(x_1, x_2) = \frac{f(x_1, x_2)\big|_P + f(x_1, x_2)\big|_Q}{2}$$
$$\bar{g}(x_1, x_2) = \frac{g(x_1, x_2)\big|_P + g(x_1, x_2)\big|_Q}{2}$$

we may write the approximation for the time taken for the trajectory to move from point *P* to point *Q* as

$$\Delta t = \frac{1}{2}\left[\frac{\Delta x_1}{\overline{f}(x_1, x_2)} + \frac{\Delta x_2}{\overline{g}(x_1, x_2)}\right] \tag{2.19}$$

It is evident that there are a large number of *f* and *g* combinations that might yield the same limit cycle contour (2.18) but different time parameterisation (2.19). Each of these systems is topologically equivalent but not conjugate.

2.3.2.2 Contours with a Fourier Series Prescribed Radius

In this section, we outline a method for defining a contour in terms of a Fourier series representation of its radius as a function of angle, expressed in Cartesian coordinates. We shall limit ourselves to the case where the radius is a single valued function of angle (hence only a limited set of periodic time series can be synthesised). Given

$$r = f(\theta)$$

we may write *f* in terms of its Fourier coefficients

$$f(\theta) = a_0 + \sum_{k=1}^{N} a_k \cos k\theta + b_k \sin k\theta$$

where we note that a truncated series fits in the least-squares sense thus allowing us to select the degree of approximation of *f* by the number of coefficients in the expansion. The coefficients are obtained from

$$a_k = 2\,\mathrm{Re}(c_k)$$
$$b_k = -2\,\mathrm{Im}(c_k), \quad k \neq 0$$

where

$$c_k = \frac{1}{T}\int_0^T f(\theta) e^{-jkw_0 t}\, dt$$

and *f* is periodic with period $2\pi/w_0$. The coefficient $a_0 = c_0$ is simply the average of *f* over one period. In the phase plane representation, where x_1 is the abscissa and x_2 is the ordinate, we obtain

$$\sin\theta = \frac{x_2}{\sqrt{x_1^2 + x_2^2}}$$
$$\cos\theta = \frac{x_1}{\sqrt{x_1^2 + x_2^2}} \tag{2.20}$$

By invoking De Moivre's theorem,

$$\cos n\alpha + j\sin n\alpha = (\cos\alpha + j\sin\alpha)^n = \sum_{k=0}^{n}\binom{n}{k}\cos^k\alpha(j\sin\alpha)^{n-k}$$

and equating real and imaginary parts, we can write the Fourier series as a polynomial of sine and cosine functions of fundamental frequency. Finally, by using (2.20) we obtain $r = f(x_1, x_2)$.

2.3.2.3 Approximating Periodic Waveforms

Before proposing a two dimensional system capable of approximating a given periodic waveform, we need to apply some kind of constraint to the system to enforce conjugacy. We may interpret a periodic waveform in terms of its *derivative coordinates embedding* (KANTZ and SCHREIBER, 1997), that is by considering the periodic signal of interest $s(t)$ and its first derivative given by the approximation

$$\dot{s}(t) \approx \frac{\left(s(t+\Delta t) - s(t-\Delta t)\right)}{2\Delta t}$$

This generates a reconstructed phase space (see Chapter IV), which may now be directly related to a dynamical system of the form

$$\begin{aligned}\frac{dx_1}{dt} &= x_2\\ \frac{dx_2}{dt} &= \Gamma(x_1, x_2)\end{aligned} \qquad (2.21)$$

for some function Γ. In other words, assuming that (2.21) exhibits the same limit cycle contour as the reconstructed phase space, the system described by (2.21) is topologically conjugate to the system producing this embedding, and $x_1(t)$ will approximate $s(t)$. Of course the system describing $s(t)$ may be of a higher order than two and hence (2.21) corresponds to a two dimensional projection. In general, we are interested in contours that can be described by

$$g(x_1, x_2) = x_1^2 + x_2^2 - f^2(x_1, x_2) = 0 \qquad (2.22)$$

where the radius f is generated from a Fourier series representation of a single-valued function of angle as described in the previous section. We propose a novel method of finding Γ in (2.21) that produces the limit cycle described by (2.22) based on constrained optimisation. Consider, for (2.21) the system

$$\begin{aligned}\frac{dx_1}{dt} &= x_2\\ \frac{dx_2}{dt} &= -a_0 x_1 - a_1 x_2 - h(x_1, x_2)\end{aligned} \qquad (2.23)$$

where h describes the prescribed contour. In matrix notation we write (2.23) as

$$\frac{dX}{dt} = AX + \begin{bmatrix} 0 \\ -h(X) \end{bmatrix}$$

where

$$A = \begin{bmatrix} 0 & 1 \\ -a_0 & -a_1 \end{bmatrix}$$

We specify the quadratic Lyapunov function candidate

$$V = X^T LX$$

where L is a given symmetric positive definite matrix. Differentiating V along the system trajectories gives

$$\frac{dV}{dt} = X^T \left[A^T L + LA \right] X - 2h(x_1, x_2)\left[l_{12}x_1 + l_{22}x_2 \right] \tag{2.24}$$

Now, let Q be defined by the Lyapunov equation

$$A^T L + LA = -Q \tag{2.25}$$

and let

$$h(x_1, x_2) = \lambda g(x_1, x_2)\left[l_{12}x_1 + l_{22}x_2 \right]$$

where g is some contour perhaps defined by (2.22) and $\lambda > 0$. Using the last two expressions, we can write (2.24) as

$$\frac{dV}{dt} = -X^T QX - 2\lambda g(x_1, x_2)\left[l_{12}x_1 + l_{22}x_2 \right]^2 \tag{2.26}$$

Equation (2.26) tells us that trajectories far from the contour g will approach inwards towards the contour since $dV/dt < 0$ and trajectories near the origin will approach the contour outwards. By Theorem 2.5 (*Poincaré-Bendixson*), a limit cycle must exist in the annulus (assuming no fixed points). Given the contour g, we wish to find the values for the free parameters that results in the closest approximation of the simulated limit cycle to the ideal contour. We generate an error function based on the distance between the simulated limit cycle and the ideal contour measured over one full rotation. The total error is given by (for a discretisation $0 = \theta_1 < \theta_2 < ... < \theta_N = 2\pi$)

$$e = \frac{1}{N}\sum_{i=1}^{N} \left| r(\theta_i) - r_1(\theta_i) \right| \tag{2.27}$$

where r is the radius of the simulated limit cycle and r_1 is the radius of the ideal contour as illustrated in Figure 2.13.

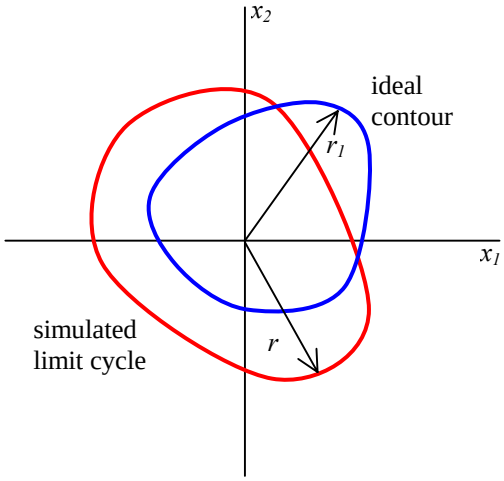

Figure 2.13: Calculating the error between a simulated limit cycle and an ideal contour

Our goal is to minimise (2.27) and we may state this as a constrained nonlinear optimisation problem including the following steps:

1) Specify a symmetric positive definite matrix Q (by using Sylvester's Criterion to constrain values of the elements of Q)
2) Specify a_0, a_1
3) Solve the Lyapunov equation (2.25). If the solution does not exist or is not unique then go to step (1)
4) Solve (2.23) allowing for an initial transient
5) Generate the error between the simulated limit cycle and ideal contour given by (2.27).

2.3.2.4 Examples

The optimisation procedure described in the previous section was applied to a number of different contours. In all cases, it was found that the system (2.23) tended to the simpler one described by

$$\frac{dx_1}{dt} = x_2$$
$$\frac{dx_2}{dt} = -x_1 - \lambda g(x_1, x_2)x_2 \qquad (2.28)$$

where $\lambda >> 1$. This results in a stiff set of equations that nonetheless produces the desired result: parameterisation by time. We conclude the chapter with an example of an ellipse and show the effect of large λ by singular perturbation analysis.

Example 2.5: *Ellipse Oscillator (Method II)*

Consider the equation for the ellipse

$$g(x_1, x_2) = \frac{x_1^2}{a^2} + \frac{x_2^2}{b^2} - 1 = 0 \tag{2.29}$$

Inserting (2.29) into (2.28), we observe that the simulated limit cycle approaches the ideal (2.29) as the magnitude of λ increases. We can explain this in terms of singular perturbation analysis. Scaling the x_1 axis and time we may write (2.28) as

$$\begin{aligned} \frac{dx_1}{dt} &= x_2 \\ \varepsilon \frac{dx_2}{dt} &= -x_1 - g(\frac{x_1}{\sqrt{\varepsilon}}, x_2)x_2 \end{aligned} \tag{2.30}$$

where

$$\varepsilon = \frac{1}{\lambda^2}$$

is a small parameter. Since one of the derivatives is being multiplied by a small parameter, the system is *singularly perturbed*.

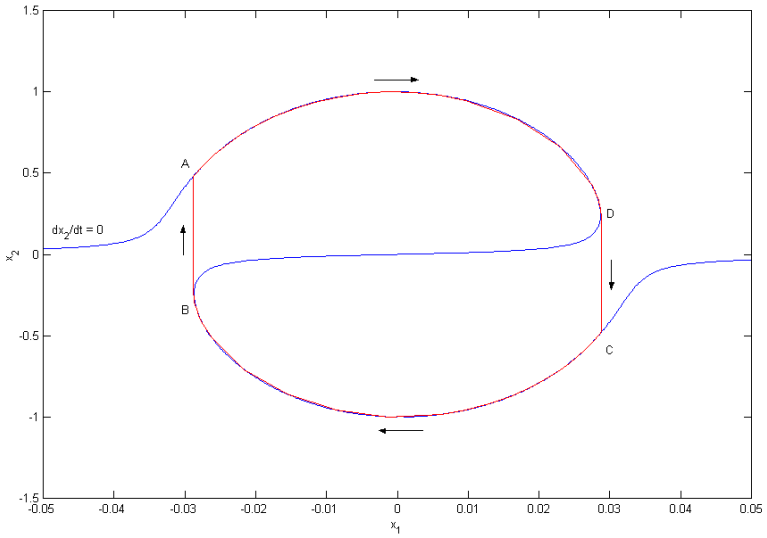

Figure 2.14: Limit cycle from system (2.30) with contour specified by (2.29). Parameters: a = 10, b = 1, ε = 0.00001

Figure 2.14 illustrates a simulation of (2.30) using (2.29) where a = 10, b = 1, ε = 0.00001. Except along the dx_2/dt nullcline where

$$x_1 \approx -g(\frac{x_1}{\sqrt{\varepsilon}}, x_2)x_2$$

x_2 changes rapidly by $O(1/\varepsilon)$ along BA and DC. The appropriate independent variable here is $\tau = t/\varepsilon$, where τ is a fast time (as opposed to the slow time t). Making this transformation (and as $\varepsilon \to 0$), the first equation of the system (2.30) becomes

$$\frac{dx_1}{d\tau} = \varepsilon x_1 \Rightarrow x_1 \approx const.$$

which is what occurs along BA and DC. As λ is made very large (conversely ε is made very small), the distance over which x_2 is changing very rapidly decreases. Figure 2.15 illustrates the case for $\varepsilon = 10^{-8}$.

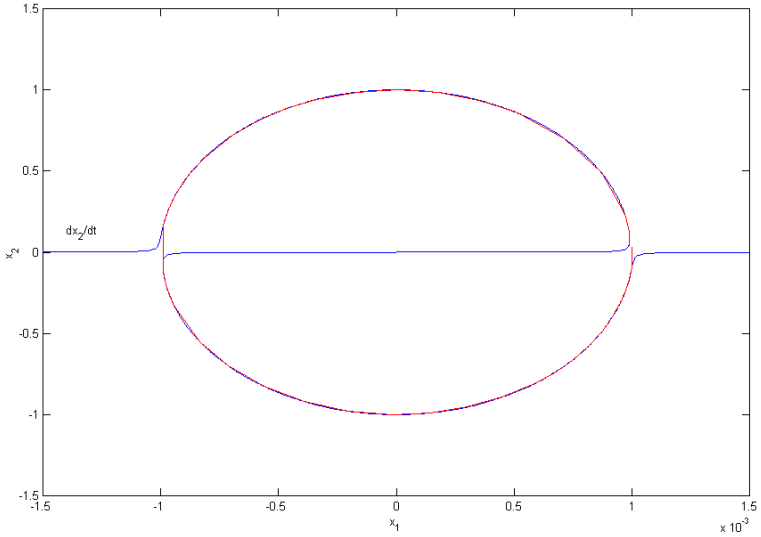

Figure 2.15: Limit cycle from system (2.30) with contour specified by (2.29). Parameters: a = 10, b = 1, ε =10^{-8}

Thus, increasing λ gives a good approximation to the ideal contour at the expense of making the system equations stiff. In the limit as $\lambda \to \infty$ the nullclines for (2.28) are given by the x_1 axis and the union of the x_1 axis and the ideal contour.

2.4 Summary

In this chapter, a solid technical foundation for creating prescribed limit cycle oscillators has been introduced for use in modelling tasks later in the thesis. Based on an extension to the conservative Hamilitonian system, the addition of dissipative terms produces asymptotically stable limit cycle behaviour. A wide variety of complex topologies were employed to demonstrate the flexibility of the method. It is expected that this novel construction will have applications in modelling, signal generation, and education. A method for approximating a certain class of periodic time series was also presented. In the next chapter, we will extend the Method I construction to produce *stochastic* limit cycle oscillators, which will serve as the basis for a novel model of the EEG signal.

CHAPTER III

MODELLING WITH STOCHASTIC OSCILLATORS

3.1 Introduction

The human electroencephalogram (EEG) potentials are aperiodic unpredictable oscillations with amplitude distributions that are near Gaussian, with intermittent bursts of oscillations having spectral peaks in certain preferred bands: 0.1-3.5 Hz (*delta*), 4-7.5 Hz (*theta*), 8-13 Hz (*alpha*), 14-30 Hz (*beta*) and >30 Hz (*gamma*) (FREEMAN, 1992). From a dynamical systems point of view, the bursts of oscillations suggest the existence of multiple limit cycles. Power spectral analysis of the EEG reveals an approximate underlying *1/f* slope, which is indicative of noise (BARLOW, 1993). We hypothesise that the ensuing activity can be described by limit cycle oscillations under continuous modulation by noise. In this chapter we develop the mathematical techniques to create prescribed stochastic limit cycle oscillators, and apply them to modelling of the EEG.

The term *Stochastic Differential Equations* (SDE) generally refers to a formulation by Itô in 1942 (GARD, 1988) with publications in the modelling applications literature only really starting to appear in the last 20 years. Analysis of SDEs differs significantly from ODEs due to peculiarities of stochastic calculus. A thorough technical exposition on the subject is beyond the scope of this thesis but we make an attempt to introduce the basic theory sufficient for our purposes; the interested reader is referred to the excellent monographs of GARD (1988), ØKSENDAL (1998), KLOEDEN and PLATEN (1999), and SOONG (1973) for a detailed introduction. After describing a new numerical toolbox for solving SDEs developed in this research, we extend our method for geometrically prescribing oscillators in the previous chapter to produce a mathematical method for prescribing *stochastic* limit cycle oscillators. An analysis of the oscillator is performed resulting in the definition of a stochastic limit cycle as an invariant set. We present two applications of SDEs for synthesising brain electrical activity. The first of these employs a purely mathematic approach and the second is based on a biophysical model of WILSON and COWAN (1972). The chapter is

concluded with a study of how complex behaviour may arise from simple stochastic limit cycle oscillators operating close to a bifurcation.

3.2 Probability Spaces, Random Variables, Stochastic Processes

To present SDEs and the applications in this chapter requires a basic understanding of formal probability theory, which is built on measure theory. We present the bare minimum of fundamental concepts (without proof) in this section.

Definition 3.1 A *probability space* is represented as an ordered triple (Ω, A, P) consisting of a sample space of outcomes Ω, a σ-algebra A of subsets of Ω called events, and a function P defined on A called probability. A σ-algebra A of Ω is a family of subsets such that

(i) the empty set $\varnothing \in A$

(ii) $\mathrm{A}_1 \in A \Rightarrow \mathrm{A}_1^c \in A$

(iii) if $\mathrm{A}_1, \mathrm{A}_2, \ldots$ are in A, then $\mathrm{A}_1 \cup \mathrm{A}_2 \cup \ldots \in A$

The outcomes $\omega \in \Omega$ are called the *simple events* and the subsets A are called the *observable events* to which we assign the actual probability measure defined by:

(i) If $\mathrm{A}_1 \in A$, $P(\mathrm{A}_1) \geq 0$, $P(\Omega) = 1$

(ii) If $\{\mathrm{A}_n\} \subseteq A$ then $P(\cup \mathrm{A}_n) = \sum P(\mathrm{A}_n)$ whenever $\mathrm{A}_i \cap \mathrm{A}_j = \varnothing$ if $i \neq j$

Given the measurable spaces (Ω_1, A_1) and (Ω_2, A_2), we can define a mapping $X : \Omega_1 \to \Omega_2$. We say X is *measurable* if $X^{-1}(\mathrm{A}) \in A_1$ for each $\mathrm{A} \in A_2$, i.e. there is an event for each occurrence x of X. The *σ-algebra generated by X* is the smallest σ-algebra that contains X i.e. the collection $\left\{X^{-1}(\mathrm{A}) : \mathrm{A} \in A_2\right\}$. Of particular importance is the Borel σ-algebra B, the σ-algebra generated by the open sets (or equivalently closed sets) of real numbers. Elements of the Borel σ-algebra are called Borel sets denoted by B. Taking $(\Omega_2, A_2) = (\Re, B)$ we define a *random variable* as $X : \Omega_1 \to \Re$, i.e. a random variable is simply an abstraction of information from the possible outcomes. We can generalise to higher dimensions by considering random vectors and random matrices, where each component is a random variable.

Definition 3.2 The *distribution function* $F_X(x)$ of a random variable is defined by $F_X(x) = P(X \le x)$. If $F_X(x) = \int_{-\infty}^{x} f_X(y)dy$ exists, f_X is called the *density* of X.

Definition 3.3 The *expectation* of a random variable X is the Lebesgue integral (assuming it exists) $\int_{\Omega} XdP$

Note, for a discrete random variable (and in the limit, the continuous random variable), the Lebesgue integral is defined as

$$\int_{\Omega} XdP = \sum_{i=1}^{\infty} x_i P\{A_i\}$$

where $A_i = \{\omega \in \Omega : X(w) = x_i\}$

A very powerful extension to Definition 3.3 is *conditional expectation*. Consider the random variable X defined on the probability space (Ω, A, P). Let H be a sub-σ-algebra of A i.e. $H \subset A$. Thus, although X is a random variable on A, it is not necessarily so on H, in other words it is not necessarily measurable with respect to H. We define the conditional expectation $E(X \mid H)$ as the random variable Y on the space (Ω, H, P) satisfying

$$\int_{\mathrm{H}} YdP = \int_{\mathrm{H}} XdP, \quad \mathrm{H} \in H$$

The existence and uniqueness of Y is guaranteed by the Radon-Nikodym theorem (ØKSENDAL, 1998). Since $E(X \mid H)$ represents the portion of information carried by the sub-σ-algebra H, we have two extremes:

$E(X \mid H) = X$

if X is $H-$measurable and

$E(X \mid H) = E(X)$

if X and H are independent. In other words, $E(X \mid H)$ is our best estimate of X given H. The latter can be seen by noting (for $\mathrm{H} \in H$ and where $I_{\mathrm{A}}(\omega)$ is the indicator function i.e. $I_{\mathrm{A}}(\omega) = 1$ if $\omega \in \mathrm{A}$ and 0 otherwise)

$$\int_{\mathrm{H}} XdP = \int_{\Omega} X.I_{\mathrm{H}}dP = \int_{\Omega} XdP.\int_{\Omega} I_{\mathrm{H}}dP = E(X).P(\mathrm{H})$$

Finally, to build intuition, it is instructive to also consider the conditional expectation definition of a random variable of X conditioned on another discrete random variable Y and is given by

$$E\left(X \mid Y\right) = \sum_{i=1}^{\infty} x_i P(X = x_i \mid Y)$$

Definition 3.4 A *stochastic process* is a parameterised collection of random variables $\{X_t\}_{t\in T}$ defined on a probability space $\left(\Omega, A, P\right)$.

The usual notation suppresses the ω but a stochastic process is written formally as $X(t,\omega),\ \omega \in \Omega,\ t \in T.$

For each fixed $\omega \in \Omega$, $X(\cdot,\omega)$ is the *realisation of the process* or *sample path*. It is useful to think of t as time and ω as an individual experiment and thus $X(t,\omega)$ can be interpreted as the result of an experiment ω at time t. Note that t may be discrete or continuous. A stochastic process is defined by its joint distributions (sometimes called its *probability law*) of its random variables written as $F_{X(t_1)\ldots X(t_n)}(x_1,\ldots,x_n) = P(X(t_1) \le x_1,\ldots,X(t_n) \le x_n).$ Not all joint distributions constitute a stochastic process; the Kolmogorov Extension Theorem (GARD, 1988) defines when they do. A Gaussian process is a stochastic process whose joint distributions are Gaussian.

Definition 3.5 A family A_t of σ-algebra fields on the probability space (Ω, A, P) parameterised by $t \in T$ is called a *filtration* if

$A_s \subset A_t \subset A$

for any $s,t \in T$ such that $s < t.$

If *X(t)* is A_t measurable for each $t \in T,$ we say that *X(t)* is *adapted* to A_t. A filtration tells us that our knowledge of a stochastic process increases with time i.e. A_t represents all the events that would have happened at time *t*.

Definition 3.6 A stochastic process X_t adapted to a filtration A_t is called a *martingale* if the conditional expectation $E\left(X(t)\middle|A_s\right) = X(s)$.

A martingale X_t is sometimes described as a "fair game" process. Interpreting X_s as the winnings at time s, a game is fair if the expected winnings for a future time $t+s$ given the game history up to the current time s are equal to the current winnings. A super-(sub-) martingale replaces the equality sign in Definition 3.6 with a $\leq$ ($\geq$). For a positive, super-martingale, Doob's maximal martingale inequality (ØKSENDAL, 1998) may be written as

$$P\left(\sup_{t_0 \leq t \leq T} X_t \geq \varepsilon\right) \leq \frac{1}{\varepsilon^p} E\left(X_{t_0}^p\right)$$

In fact, martingale convergence theorems guarantee the existence of the random variable X_∞ for super- and sub-martingales for which $X_t \to X_\infty$ with probability one and $E\left(X_t\right) \to E\left(X_\infty\right)$ as $t \to \infty$. We use the concept of super-martingales in the study of stochastic stability later in this chapter.

Definition 3.7 A stochastic process *X(t)* is (*strictly*) *stationary* if its probability distributions are invariant under an arbitrary time translation. Specifically, if for each n, t, then $F_{X(t_1)\ldots X(t_n)}(x_1,\ldots,x_n) = F_{X(t_1+\tau)\ldots X(t_n+\tau)}(x_1,\ldots,x_n)$

Definition 3.8 A stationary stochastic process is ergodic relative to G if, for every $g[X(t)] \in G$, the time average $\overline{g}[x(t)] = E\{g[x(t)]\}$

where we define

$$\overline{g}[x(t)] = \lim_{T\to\infty} (1/2T)\int_{-T}^{T} g[x(t+\tau)]d\tau$$

As an example, $g[X(t)] = X(t)$ yields the definition for *ergodic in the mean*. Similar definitions can be made for higher order moments and joint moments.

3.2.1 Wiener Process

In 1828, the Scottish botanist Robert Brown observed that pollen grains suspended in a liquid performed random motion and was subsequently explained later to be the result of the random collisions with molecules of the liquid. Norbert Wiener gave the first rigorous treatment of Brownian motion as a stochastic process describing the position of particle ω at time t. This stochastic process is known as the *Wiener process* (GARD, 1988) and is defined as a Gaussian process with independent increments $\left[W(t) - W(s)\right]$, $t > s$ such that the following properties are observed

$$W(0) = 0$$
$$E(W(t)) = 0$$
$$E\left([W(t) - W(s)]^2\right) = |t - s| \tag{3.1}$$

i.e. the variance of sample paths increase without bound while the mean remains 0. For computational purposes, it is convenient to consider discretised Wiener processes. Consider a time interval $[0, T]$ with a discretisation of $0 = \tau_0 < \tau_1 < \ldots < \tau_N = T$ with time step $\delta = T/N$.

Let,

$$\Delta W_n = W_{\tau_{n+1}} - W_{\tau_n}$$

This can be approximated by

$$\Delta W_n \approx \sqrt{\delta} N(0,1).$$

Figure 3.1 illustrates an example of a simulated sample path of a Wiener process.

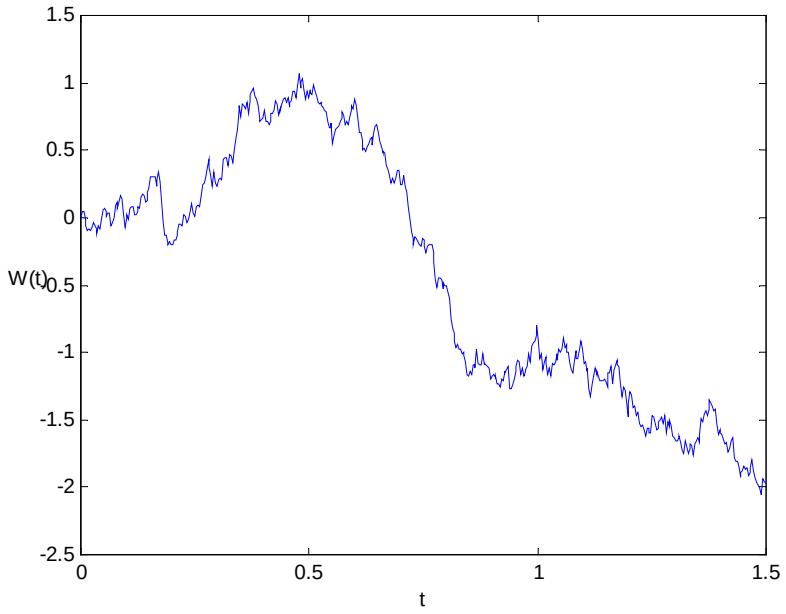

Figure 3.1: Sample path of a Wiener process where $W_j = W_{j-1} + \Delta W_j$ with $N = 500$, $T = 1.5$

The importance of the Wiener process for the sequel stems from its relationship with white noise. The covariance of the Wiener process, following from (3.1), is

$$C_W(t,s) = \min(t,s)$$

Although the Wiener process possesses continuous sample paths, it is nowhere differentiable. Nevertheless, we might consider the formal derivative below and the corresponding covariance

$$C_{\dot{W}}(t,s) = \partial^2 C_W(t,s) / \partial t \partial s$$

which evaluates to

$$C_{\dot{W}}(t,s) = \partial^2 C_W(t,s)/\partial t \partial s = \partial U(t-s)/\partial t = \delta(t-s)$$

where U is the Heaviside unit step and δ is the Dirac delta function. Thus the covariance of the formal derivative of the Wiener process is the covariance of white noise. The key step we use in the following sections is to replace white noise contributions with the Wiener process in differential equations.

3.3 Stochastic Differential Equations

Stochastic Differential Equations arose from early attempts by Langevin to describe Brownian Motion in terms of a deterministic or averaged drift term perturbed by a noisy diffusive term (KLOEDEN and PLATEN, 1999). The symbolic equation of the form

$$dX = a(t,X)dt + b(t,X)\eta(t)dt$$

is typically used to define an SDE where *a* is the drift term, *b* is the diffusion term, and $\eta(t)$ represents Gaussian white noise. This equation is interpreted as an integral equation of the form

$$X(t) = X(t_0) + \int_{t_0}^{t} a(s,X(s))ds + \int_{t_0}^{t} b(s,X(s))\eta(s)ds$$

Replacing the white noise term $\eta(t)$ with the time derivative of the Wiener process (as discussed in Section 3.2) we arrive at

$$X(t) = X(t_0) + \int_{t_0}^{t} a(s,X(s))ds + \int_{t_0}^{t} b(s,X(s))dW(s) \tag{3.2}$$

3.3.1 Itô Stochastic Differential Equations

Evaluating (3.2) by conventional means presents a problem. We might expect to interpret the second integral as the Riemann-Stieltjes integral (where just for simplicity we consider *b(s,X(s)) = b(s)*, i.e. additive noise):

$$\sum_{n=1}^{N} b(s'_n)\{W(\tau_{n+1}) - W(\tau_n)\} \qquad s'_n \in [\tau_n, \tau_{n+1}) \tag{3.3}$$

A problem arises in that the Wiener process is of unbounded variation and nowhere differentiable. Different values of s'_n within the interval above result in the integral converging to different values (in contrast to a regular Riemann-Stieltjes integral). Itô showed that the mean-square limit of (3.3) exists and is unique if we take $s'_n = \tau_n$ and the integral defined by

$$\sum_{n=1}^{N} b(\tau_n)\{W(\tau_{n+1}) - W(\tau_n)\}$$

is known as the *Itô integral*. It is easy to see that the expectation of an Itô integral is 0, a result which will prove useful later. When discussing stochastic differential equations of the form

$$dX = a(t, X)dt + b(t, X)dW(t)$$

$$X(t) = X(t_0) + \int_{t_0}^{t} a(s, X(s))ds + \int_{t_0}^{t} b(s, X(s))dW(s) \tag{3.4}$$

we are implying the use of Itô integrals. There is another alternative definition of a stochastic integral known as the Stratonovich integral and involves taking s'_n at the midpoint of the interval in (3.3). From a mathematical viewpoint both definitions are correct and we can map easily from one definition to the other (sometimes one approach can be more advantageous than the other for a particular application).

3.3.1.1 Existence and Uniqueness

When considering solutions to (3.4), we may deal with *strong solutions* or *weak solutions*. A strong solution is where the Wiener process *W* is specified. Weak solutions deal with the case where the coefficients *a* and *b* are specified but the actual Wiener process is not. As with ODEs, SDEs can, but typically do not have explicit solutions. The ability of a numerical method to calculate strong/weak solutions is called the *strong*/*weak order of convergence* and is discussed in Section 3.4. The question of whether an SDE actually possesses a unique solution for a numerical technique to find is answered next. This is really an extension of the deterministic case.

Theorem 3.1 *(Existence and uniqueness of stochastic differential equations)*

Equation (3.4) possesses a pathwise unique solution defined on $[0,T]$ if for all $t \in [0,T]$

- *a(t,X)* and *b(t,X)* are measurable with respect to *t* and *x*
- $|a(t, X_1) - a(t, X_2)| + |b(t, X_1) - b(t, X_2)| \le K|X_1 - X_2|$ (Lipschitz condition)
- $|a(t, X)|^2 + |b(t, X)|^2 \le K^2(1 + |X|^2)$ (growth condition)
- $X(t_0)$ is independent of *W(t)*, for $t > 0$, and $E\{X(t_0)^2\} < \infty$

By pathwise unique, we mean for two equivalent solutions of (3.4) $X_1(t)$ and $X_2(t)$ with continuous sample paths, we have

$$P\left(\sup_{0\le t\le T}\left|X_1(t)-X_2(t)\right|>0\right)=0$$

The solution of (3.4) subject to Theorem 3.1 is a stochastic process and is called an *Itô diffusion*. More accurately, the *family* of solutions $X_t^{0,x}$ (where $X_0 = x$) $\forall x \in \Re$ is an Itô diffusion.

3.3.1.2 Itô Formula

Stochastic integrals do not transform under the usual rules of the chain rule. The stochastic analogue of the chain rule for Itô integrals is known as the *Itô formula*. Consider

$$Y_t = f(t, X_t)$$

where X_t is a solution of (3.4) and f has continuous partial derivatives

$$\frac{\partial f}{\partial t}, \frac{\partial f}{\partial x}, \text{ and } \frac{\partial^2 f}{\partial x^2}.$$

By taking the Taylor expansion of Y_t around t and X_t we obtain

$$\begin{aligned} Y(t+\Delta t, X+\Delta X_t) = Y(t, X_t) \\ &+\left\{\frac{\partial f}{\partial t}\Delta t+\frac{\partial f}{\partial x}\Delta x\right\} \\ &+\frac{1}{2}\left\{\frac{\partial^2 f}{\partial t^2}(\Delta t)^2+2\frac{\partial^2 f}{\partial t\partial x}\Delta t\Delta x+\frac{\partial^2 f}{\partial x^2}(\Delta x)^2\right\} \\ &+ \ldots \end{aligned} \tag{3.5}$$

The usual chain rule can be deduced by taking infinitesimal increments and incorporating only the first derivatives. In the stochastic case, noting that $E((dW)^2) = dt$ an extra term enters though the last term on the RHS of (3.5) yielding the Itô formula (inserting (3.4) and taking infinitesimal increments)

$$dY_t = \left\{\frac{\partial f}{\partial t}+a\frac{\partial f}{\partial x}+\frac{1}{2}b^2\frac{\partial^2 f}{\partial x^2}\right\}dt + b\frac{\partial f}{\partial x}dW_t \tag{3.6}$$

3.3.1.3 Stable in Probability

Lyapunov's Direct Method (see Chapter II) can be extended for the stochastic case by introducing the definition for stability by HAS'MINSKII (1980). Consider X_t the solution

of (3.4), we say it is *stable in probability* or *stochastically stable* if for any $\varepsilon > 0$ and $t_0 \geq 0$

$$\lim_{x_0 \to 0} P\left(\sup_{t \geq t_0} \left| X_t^{t_0, x_0} \right| \geq \varepsilon \right) = 0$$

Furthermore, the solution X_t is *stochastically asymptotically stable* if

$$P\left(\sup_{t \to \infty} \left| X_t^{t_0, x_0} \right| \to 0 \right) = 1$$

As with the deterministic case, Lyapunov's method is concerned with time derivatives of the Lyapunov function *V(x)* (we'll consider the 1-dimensional case initially). By invoking the Itô formula we may write

$$dV_t = \left\{ a \frac{\partial V}{\partial x} + \frac{1}{2} b^2 \frac{\partial^2 V}{\partial x^2} \right\} dt + b \frac{\partial V}{\partial x} dW_t \tag{3.7}$$

Denoting the first term as

$$LV = \frac{\partial V}{\partial t} + a \frac{\partial V}{\partial x} + \frac{1}{2} b^2 \frac{\partial^2 V}{\partial x^2}$$

we may state the stochastic version of Lyapunov's Direct Method by simply replacing *dV/dt* with *LV*.

Theorem 3.2 *(Stochastic Version of Lyapunov's Direct Method)* If one can find a definite function *V* for the set of differential equations (3.4) such that *LV* is either identically zero or semi-definite with opposite sign, then the motion is stable.

Sketch of Proof

Consider (3.4) with a steady state solution $X_t=0$ (and hence $a(t,0)=0$ and $b(t,0)=0$). Writing (3.7) in integral form yields

$$V(t, X_t) - V(t_0, X_0) = \int_{t_0}^{t} LV(\tau, X_\tau) d\tau + \int_{t_0}^{t} b(\tau, X_\tau) \frac{\partial V}{\partial x} dW_\tau \tag{3.8}$$

Assuming the inequality

$LV \leq 0$

allows us to write (3.8) as

$$V(t, X_t) - V(t_0, X_0) \leq \int_{t_0}^{t} b(\tau, X_\tau) \frac{\partial V}{\partial x} dW_\tau \tag{3.9}$$

By taking expectations of both sides (and noting that the expectation of an Itô integral is 0) we get

$$E(V(t, X_t)\big| A_{t_0}) \le V(t_0, X_{t_0})$$

which states that the Lyapunov function evaluated along the solution curves of (3.4) is a super-martingale (i.e. it tends to decrease) and the maximal martingale inequality

$$P\left(\sup_{t \ge t_0} V(t, X_t) \ge \varepsilon \right) \le \frac{1}{\varepsilon} V\left(t_0, X_{t_0}\right)$$

holds. We have avoided the technical difficulty that since V is not necessarily smooth on $X = 0$, the Itô formula (see Section 3.3.1.2) cannot be used directly. This problem can be avoided by minor changes to the proof – see GARD (1988) for a more complete version. Finally, for the d-dimensional case we may generalise LV to

$$LV = \sum_{i=1}^{d} a^i \frac{\partial V}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^{d} \left(bb^T \right)^{i,j} \frac{\partial^2 V}{\partial x_i \partial x_j} \tag{3.10}$$

3.4 Numerical Solution of SDEs

Deterministic approaches for numerically solving deterministic differential equations do not carry over directly to stochastic differential equations due to the different calculus involved and hence a different approach is required. Moreover, and although there has been much development in the recent two decades, the techniques for solving SDEs are still in their infancy.

3.4.1 Strong vs. Weak Convergence

To compare different solution techniques, it is instructive to introduce a measure of accuracy between a given solution and the actual solution of an SDE. Unlike the simpler deterministic case, we need to consider two kinds of approximations namely those related to sample paths and those corresponding to distributions. Consider the actual solution of an SDE as X_t and an approximated solution as Y_i, discretised over N steps. The *strong order of convergence* γ is then given evaluated at the final time instant via

$$E\left(\left| X_T - Y_N^{\delta} \right| \right) \le K \delta^{\gamma}$$

for all step sizes $\delta \in (0,1)$ and K is a finite positive constant. In other situations, close pathwise approximations may not be required (for example maybe just the mean or

higher order moments) and the *weak order of convergence* β might be more appropriate

$$\left|E\left(g\left(X_T\right)\right)-E\left(g\left(Y_N^\delta\right)\right)\right| \le K_g \delta^\beta$$

where g is any polynomial. We will be only concerned with the strong order of convergence measure since we will be working with direct simulation of trajectories for comparison with real data.

3.4.2 Stochastic Taylor Series

The Taylor series expansion is used as the basis for creating deterministic numerical solver techniques by including higher order terms for increased accuracy. A stochastic version of the Taylor series (KLOEDEN and PLATEN, 1999) in integral form can be derived by taking the integral form of (3.4) and repeatedly applying the Itô formula:

$$X_t = X_{t_0} + a(X_{t_0})\int_{t_0}^{t} d\tau + b(X_{t_0})\int_{t_0}^{t} dW_\tau$$
$$+b(X_{t_0})b'(X_{t_0})\int_{t_0}^{t}\int_{t_0}^{\tau_2} dW_{\tau_1} dW_{\tau_2} + R \qquad (3.11)$$

where R is a remainder containing higher order terms. Equation (3.11) contains multiple Itô integrals as building blocks and their approximation is a considerable challenge for numerical techniques.

3.4.3 Numerical Solution Algorithms

In this section we shall consider a collection of numerical solution algorithms of increasing strong order accuracy and complexity namely the *Euler-Maruyama* (strong order 0.5), *Milstein* (strong order 1.0), and *Runge-Kutta* (strong order 1.5) methods. We will consider approaches that avoid the need for derivatives to simplify implementation (see next section). The monograph by KLOEDEN and PLATEN (1999) serves as the 'recipe book' for this and the following sections.

3.4.3.1 Euler-Maruyama

The Euler-Maruyama method is probably the simplest solver for SDEs and is an extension of the Euler method used for deterministic differential equations. The scheme follows from the integral form of (3.4) if we consider a uniform discretisation of N equal time steps from time *0* to time *T*. This yields a time step of

$$\Delta t = \tau_{n+1} - \tau_n = T/N$$

The scheme may be written as

$$Y_{n+1} = Y_n + a(Y_n)\Delta t + b(Y_n)\Delta W_n \tag{3.12}$$

where Δt is the time step and

$$\Delta W_n = W(\tau_{n+1}) - W(\tau_n)$$

For $b = 0$, we obtain the usual deterministic Euler scheme. The random variables in (3.12) may be obtained from a Gaussian distribution of mean 0 and variance $\sqrt{\Delta t}$ .The multi-dimensional case with dimension *d* for the drift vector and *d* x *m* for the diffusion matrix is given by

$$Y_{n+1}^k = Y_n^k + a^k \Delta t + \sum_{j=1}^{m} b^{k,j} \Delta W^j$$

where ΔW^j refers to the j^{th} component of the increments of an *m*-dimensional Wiener process (each component is independent).

3.4.3.2 Milstein

Adding an extra term from the stochastic Taylor series results in the Milstein scheme of strong order 1.0

$$Y_{n+1} = Y_n + a(Y_n)\Delta t + b(Y_n)\Delta W_n + \frac{1}{2\sqrt{\Delta t}}\left\{b(\tau_n, \hat{Y}_n) - b(Y_n)\right\}\left\{(\Delta W)^2 - \Delta t\right\} \tag{3.13}$$

where the supporting $\hat{Y}$ is given by

$$\hat{Y}_n = Y_n + a\Delta t + b\sqrt{\Delta t}$$

This allows us to replace the derivative of *b* with its equivalent deterministic Taylor expansion (assuming negligible higher order terms). For generalising to the multi-dimensional case, a useful simplification will be made. We consider only diagonal noise, that is each component of the Itô process X_t is disturbed by the corresponding component of the Wiener process W_t. This implies that the diffusion matrix *b* is replaced with a diagonal matrix i.e. $b_{i,j} = 0$ for $i \neq j$. The multi-dimensional case with dimension *d* for the drift vector and *d* x *d* for the diffusion matrix is then given by

$$\begin{aligned} Y_{n+1}^k &= Y_n^k + a^k \Delta t + b^{k,k} \Delta W^k \\ &+ \frac{1}{2\sqrt{\Delta t}}\left\{b^{k,k}(\tau_n, \hat{Y}_n^k) - b^{k,k}\right\}\left\{(\Delta W^k)^2 - \Delta t\right\} \end{aligned} \tag{3.14}$$

where the supporting $\hat{Y}$ is given by

$$\hat{Y}_n^k = Y_n + a\Delta t + b^k\sqrt{\Delta t}$$

3.4.3.3 Runge-Kutta

By adding further terms from the stochastic Taylor series, and again assuming diagonal noise, the resultant multi-dimension 1.5 strong order scheme is given by

$$\begin{aligned}
Y_{n+1}^k &= Y_n^k + b^{k,k}\Delta W^k \\
&+ \frac{1}{2\sqrt{\Delta t}}\sum_{j=1}^{d}\left\{a^k(\hat{Y}_+^j) - a^k(\hat{Y}_-^j)\right\}\Delta Z^k \\
&+ \frac{1}{4}\Delta t\sum_{j=1}^{d}\left\{a^k(\hat{Y}_+^j) - \frac{2(d-2)}{d}a^k + a^k(\hat{Y}_-^j)\right\} \\
&+ \frac{1}{4\sqrt{\Delta t}}\left\{b^{k,k}(Y_+^k) - b^{k,k}(Y_-^k)\right\}\left\{(\Delta W^k)^2 - \Delta t\right\} \\
&+ \frac{1}{2\Delta t}\left\{b^{k,k}(Y_+^k) - 2b^{k,k} + b^{k,k}(Y_-^k)\right\}\left\{\Delta W^k\Delta t - \Delta Z^k\right\} \\
&+ \frac{1}{4\Delta t}\left\{b^{k,k}(\Phi_+^k) - b^{k,k}(\Phi_-^k) - b^{k,k}(Y_+^k) + b^{k,k}(Y_-^k)\right\} \\
&\times\left\{\frac{1}{3}(\Delta W^k)^2 - \Delta t\right\}\Delta W^k
\end{aligned} \tag{3.15}$$

where the supporting values are given by

$$\hat{Y}_\pm^k = Y_n + \frac{1}{d}a\Delta t \pm b^k\sqrt{\Delta t}$$

$$\Phi_\pm^k = \hat{Y}_+^k \pm b^k(\hat{Y}_+^k)\sqrt{\Delta t}$$

Note that an extra simplification of assuming an autonomous system *b* matrix was made in (3.14) to reduce the complexity and increase the efficiency of implementations. Finally note that the ΔZ term in (3.14) is an approximation of

$$\int_{\tau_n}^{\tau_{n+1}}\int_{\tau_n}^{s_2} dW_{s_1} ds_2$$

and is related to ΔW by considering two normally distributed random variables *N(0,1)* U_1 and U_2 in

$$\Delta W = U_1\sqrt{\Delta t}$$

$$\Delta Z = \frac{1}{2}\Delta t^{3/2}\left(U_1 + \frac{1}{\sqrt{3}}U_2\right)$$

3.4.4 SDESolver: A Toolbox for the Numerical Solution of SDEs

Although there has been an accelerating development of algorithms for the numerical solution of stochastic differential equations in the last decade or so, the availability of software toolboxes is still very rare. To facilitate application of the theory presented in this chapter, it quickly became apparent that a **new software toolbox** needed to be created. SDESolver is an extensible toolbox for performing numerical solution of SDEs and is used in the simulation of all the stochastic differential systems found in this thesis. An Object-Oriented (O-O) framework was designed to facilitate new solvers to be "plugged-in". The core classes are written in ANSI C++ for portability. The toolbox is further packaged for Linux in "GNU style", facilitating a suite of command line arguments and the ability to load dynamical systems at runtime from shared libraries thus obviating the need to recompile the main binary when changes are required. Three solvers are supplied with the toolbox:

- Euler-Maruyama (strong order 0.5)
- Milstein (strong order 1.0)
- Runge-Kutta (strong order 1.5)

For simplicity we have assumed the use of Itô integrals and explicit solvers that do not require the derivatives of the system to be solved. The source code can be found in Appendix A2 (CD-ROM).

3.4.4.1 O-O Framework

The object-oriented architecture of the toolbox is illustrated in Figure 3.2 using Unified Modelling Language (FOWLER and SCOTT, 2000). Solver implementations are derived from the base class CSolver and instantiated via the Parameterised Factory CSolverFactory (GAMMA *et al*, 1994). The CSolverSettings class aggregates the solver independent settings. To allow the solvers to be developed independently of the dynamical system to be solved (and vice versa), a generic interface CDynSystem is employed to represent a dynamical system. This class has a method called GetDriftAndDiffusion(), which returns the a and b components. This is noteworthy as it differs from traditional (deterministic) solvers, which often employ a function to return the entire right hand side of the equation to be solved. Different algorithms for solving SDEs on the other hand manipulate the drift and diffusion components independently and hence need to be obtained from the dynamical system separately.

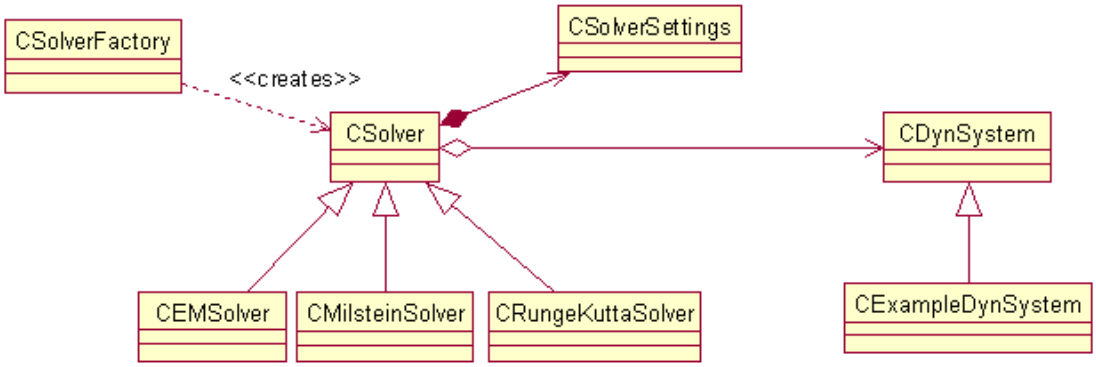

Figure 3.2: UML class diagram of the toolbox

To ease the development of new solvers, the Template Method pattern (GAMMA *et al*, 1994) is employed so that CSolver defines a skeleton of the solution procedure, deferring the specific details of advancing the solution to subclasses. Figure 3.3 illustrates a sequence diagram of the solution procedure.

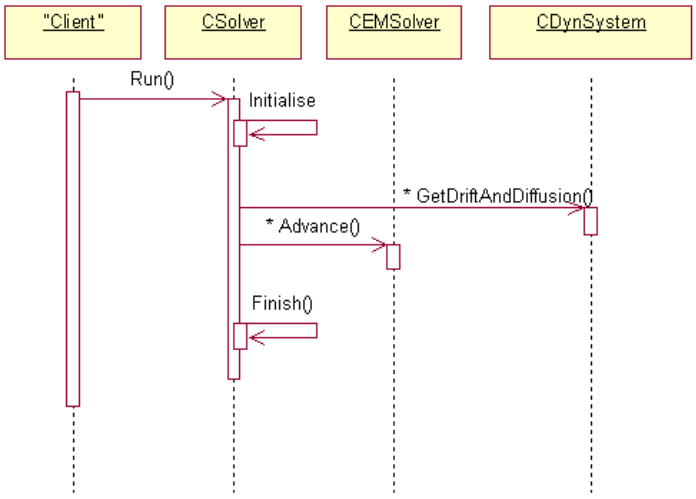

Figure 3.3: UML sequence diagram of the solution procedure for an example solver (CEMSolver)

A solver is invoked by calling the Run() method. This calls Initialise(), which seeds the random number generator (either by a user supplied value or one derived from the system time), creates the Wiener process realisation, and initialises memory for the solution. If specified, the Initialise() method writes the Wiener process realisation to disk also. For each time step, the drift and diffusion components of the system to be solved are retrieved via the GetDriftAndDiffusion() and the Advance() method of the particular solver is invoked. The solution process terminates with a call to the Finish() method, which writes the solution to disk. Figure 3.4 illustrates the usage.

```
CSolverSettings settings;
// … set member variables of settings as appropriate…

// Create solver from factory
CSolver* pSolver = CSolverFactory::CreateSolver(CSolverFactory::MILSTEIN, settings);
if (NULL != pSolver)
{
        // Register dynamical system to be solved
        CDynSystem* pSystem = new CGBMDynSystem();
        pSolver->RegisterDynSystem(pSystem);

        // Run solver
        if (!pSolver->Run())
        {
                fprintf(stderr, "%s\n", pSolver->GetLastError().c_str());
        }
}
```

Figure 3.4: Usage for SDESolver core classes (C++)

3.4.4.2 Random Number Generation

Gaussian distributed deviates are required for modelling a Wiener process and thus the stochastic contribution in SDEs. The CSolver base class creates these random variables during the Initialise() method by invoking an efficient Box-Muller transformation of a uniformly generated sequence of random numbers (BOX and MULLER, 1958). The latter are generated by using a high quality random number generator, the "Mersenne Twister" (MATSUMOTO and NISHIMURA, 1998). It should be noted that the random generator provided in the standard C library generally performs very poorly when employed in stochastic modelling settings. Figure 3.5 illustrates a typical distribution of the Wiener increments created by this method. To simplify the implementation and to prevent the introduction of errors, the Wiener process is calculated at m_nStepSize increments and the solver computes the solution at increments of m_nStepSize*m_nMultiplier, thus avoiding the need for interpolation.

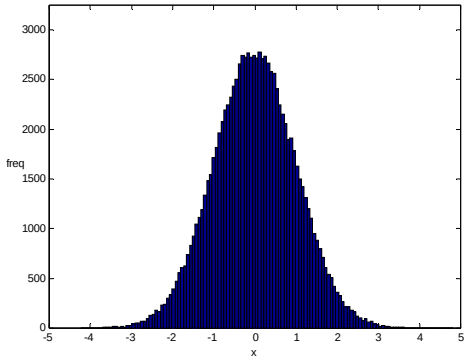

Figure 3.5: High quality *N(0,1)* for 10,000 points generated by SDESolver using the Mersenne Twister and the Box-Muller algorithm

3.4.4.3 The SDESolver Toolbox in Action

To demonstrate SDESolver, we'll solve an SDE from the famous Black-Scholes model for pricing derivatives (BLACK and SCHOLES, 1973). The equation, used to model the value of shares on the stock market, is also known as Geometric Brownian Motion and is given by

$$dS = \mu S dt + \sigma S dW \tag{3.16}$$

Equation (3.16) possesses the explicit solution

$$S_t = S_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma W_t} \tag{3.17}$$

which affords us the ability to view the performance of the different solvers (the toolbox has a setting that forces it to write the W_t to disk and hence a plot of (3.17) can be made).

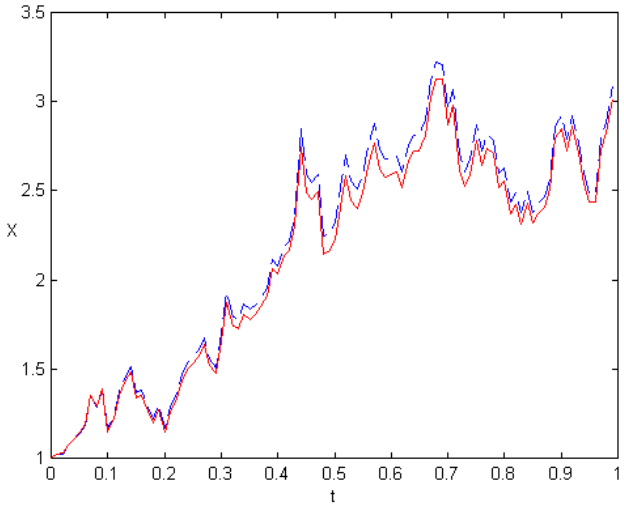

Figure 3.6: Equation (3.17) dashed, and the Euler-Maruyama (strong order 0.5) solution of equation (3.16) for a step size of 0.01 (and multiplier 1)

Figure 3.6 illustrates equation (3.17) (dashed) with the corresponding Euler-Maruyama solution of (3.16). Figure 3.7 is the same as Figure 3.6 but this time using the Runge-Kutta solver and produces a more accurate solution (in fact the solutions coincide perfectly) for the same step size. Finally, Figure 3.8 illustrates the ensemble average of 100 simulations with a different Wiener process used for each simulation.

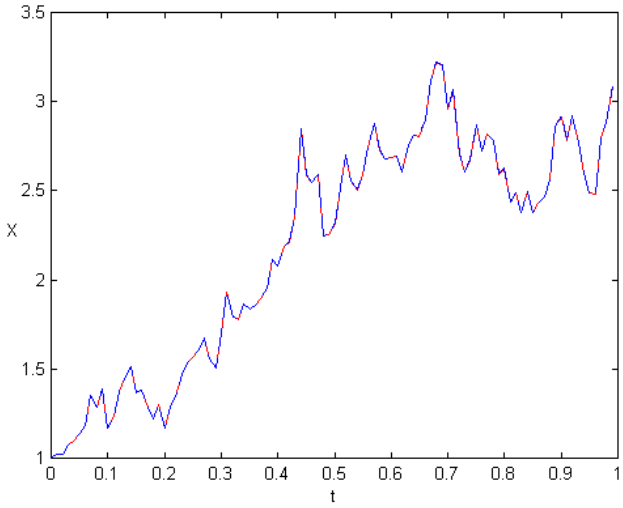

Figure 3.7: Equation (3.17) dashed, and the Runge-Kutta (strong order 1.5) solution of (3.16) for a step size of 0.01 (and multiplier 1)

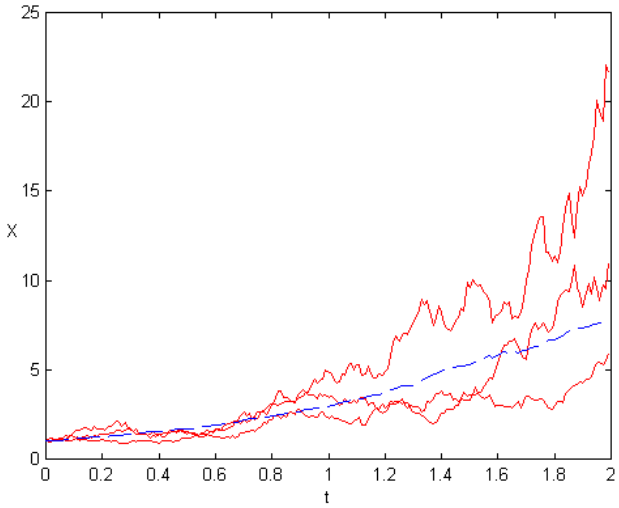

Figure 3.8: Three solutions of (3.16) using the Runge-Kutta solver with different Wiener process realisations and, superimposed (dashed), an average of 100 simulations.

3.5 Stochastic Limit Cycle Oscillators

In this section, we extend our method for synthesising prescribed limit cycle oscillators from Chapter II to the case of *stochastic* limit cycle oscillators. The (deterministic) system under consideration (reproduced here for convenience) is

$$\dot{x}_1 = \frac{\partial H}{\partial x_2} - \lambda H \frac{\partial H}{\partial x_1}$$
$$\dot{x}_2 = -\frac{\partial H}{\partial x_1} - \lambda H \frac{\partial H}{\partial x_2} \qquad \lambda > 0 \tag{3.18}$$

where we have shown previously that the limit cycle defined by

$$H(x_1, x_2) = 0$$

is asymptotically stable. We introduce Gaussian white noise by applying the following transformation

$$H(x_1, x_2) \to H(x_1, x_2) + \sigma\eta(t) \tag{3.19}$$

where σ is the intensity of the noise. Equation (3.19) might be interpreted as 'geometric additive noise'. Inserting (3.19) into (3.18) and replacing the white noise term by the derivative of the Wiener process, we formulate the vector stochastic differential system

$$dX = \begin{pmatrix} \dfrac{\partial H}{\partial X_2} - \lambda H \dfrac{\partial H}{\partial X_1} \\ -\dfrac{\partial H}{\partial X_1} - \lambda H \dfrac{\partial H}{\partial X_2} \end{pmatrix} dt + \sigma \begin{pmatrix} -\lambda \dfrac{\partial H}{\partial X_1} & 0 \\ 0 & -\lambda \dfrac{\partial H}{\partial X_2} \end{pmatrix} dW \tag{3.20}$$

where the column vectors are defined as

$$dX = \begin{pmatrix} dX_1 \\ dX_2 \end{pmatrix}, dW = \begin{pmatrix} dW_1 \\ dW_2 \end{pmatrix}$$

3.5.1 Stochastic Stability Analysis

In this section we investigate the stability of the limit cycle in (3.20) with a view to defining the behaviour. We consider a positive definite function, which is a nonlinear measure of the distance from the prescribed noise-free contour.

$$V = H^2 \tag{3.21}$$

We wish to show that the resultant behaviour of the oscillator in the neighbourhood of the prescribed contour forms an invariant set and thus trajectories may not leave this set with probability one. Writing the inequality

$$LV \le 0 \tag{3.22}$$

we may deduce from Theorem 3.2 that V evaluated along solution trajectories of (3.20) forms a super-martingale

$$E(V(t, X_t) \big| A_{t_0}) \le V(t_0, X_{t_0})$$

and the maximal martingale inequality

$$P\left(\sup_{t \ge t_0} V(t, X_t) \ge \varepsilon \right) \le \frac{1}{\varepsilon} V\left(t_0, X_{t_0}\right)$$

holds. Thus the expectation of the distance from the prescribed contour is non-increasing.

In particular we note this holds over the entire phase space since V is a nonlinear measure of distance and therefore

$$\lim_{x_1,x_2\to\infty} V(x_1,x_2)\to\infty$$

It remains to find under what conditions (3.22) is obeyed. For clarity, we separate (3.22) into two terms LV_1, LV_2 and thus $LV = LV_1 + LV_2$. The first term evaluates to the usual chain rule and hence the term used in proving asymptotic stability for the deterministic version (see Chapter II), namely

$$LV_1 = -2\lambda H^2\left[\left(\frac{\partial H}{\partial X_1}\right)^2+\left(\frac{\partial H}{\partial X_2}\right)^2\right]$$

which is negative semi-definite for $\lambda > 0.$ The second term evaluates to

$$LV_2 = \sigma^2\left\{\begin{array}{l}\underbrace{\lambda^2\left[\left(\frac{\partial H}{\partial X_1}\right)^4+\left(\frac{\partial H}{\partial X_2}\right)^4+2\left(\frac{\partial H}{\partial X_1}.\frac{\partial H}{\partial X_2}\right)^2\right]}_{positive\ \ definite}+ \\ \underbrace{\lambda^2 H\left[\left(\frac{\partial H}{\partial X_1}\right)^2\frac{\partial^2 H}{\partial X_1^2}+\left(\frac{\partial H}{\partial X_2}\right)^2\frac{\partial^2 H}{\partial X_2^2}+\frac{\partial H}{\partial X_1}.\frac{\partial H}{\partial X_2}.\frac{\partial^2 H}{\partial X_1\partial X_2}\right]}_{sign\ \ indefinite}\end{array}\right\}$$

and hence $LV_2 = O(\sigma^2)$. Thus we may state our main result with regard to stability of the prescribed stochastic limit cycle oscillator:

For sufficiently small σ, $LV \le 0$ and the distance from the prescribed contour H measured along the system trajectories is a super-martingale and thus an invariant set exists in the neighbourhood of the contour which we denote as a *stochastic limit cycle*.

In other words, the distance from the (noise free) contour measured along trajectories "tends to decrease" and trajectories in the neighbourhood of the contour must remain in the neighbourhood with probability one for all time.

3.5.2 Examples of Stochastic Limit Cycle Oscillators

We present two examples of stochastic limit cycle oscillators in this section.

Example 3.1: *Square Limit Cycle Oscillator*

Consider the square contour described by

$$H(x_1,x_2) = |x_1| + |x_2| - r^2 = 0$$

perturbed by additive noise $\sigma\eta(t)$ where σ is the intensity of the noise. Writing in the form of (3.20), the stochastic differential system is

$$dX = \begin{pmatrix} \operatorname{sgn}(X_2) - \lambda\left(|X_1| + |X_2| - r^2\right)\operatorname{sgn}(X_1) \\ -\operatorname{sgn}(X_1) - \lambda\left(|X_1| + |X_2| - r^2\right)\operatorname{sgn}(X_2) \end{pmatrix} dt + \sigma \begin{pmatrix} -\lambda \operatorname{sgn}(X_1) & 0 \\ 0 & -\lambda \operatorname{sgn}(X_2) \end{pmatrix} dW \qquad (3.23)$$

Figure 3.9 illustrates a numerical simulation of (3.23) with parameters as per the figure caption.

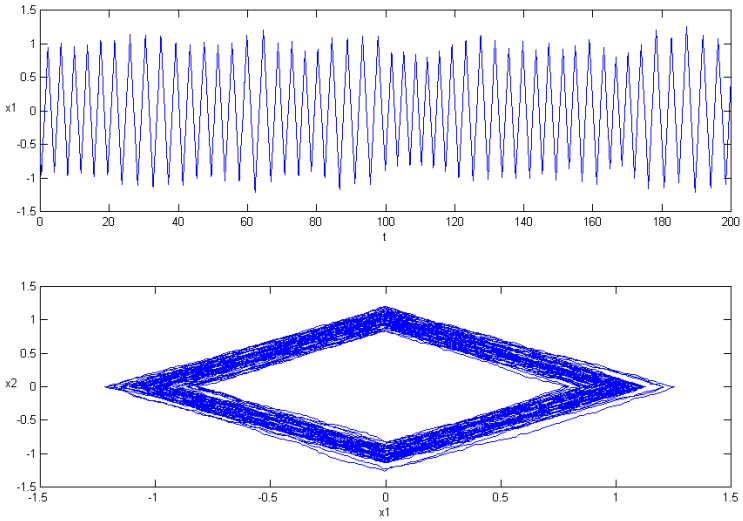

Figure 3.9: Stochastic Square Limit Cycle Oscillator
Simulation parameters: Runge-Kutta, *r=1, λ=0.08, σ=0.5,* $\Delta t = 0.01$, *T=200*

Example 3.2: *Steiner's Hypocycloid*

Consider Steiner's Hypocycloid contour given by

$$H(x_1, x_2) = \left(x_1^2 + x_2^2\right)^2 + 8rx_1\left(3x_2^2 - x_1^2\right) + 18r^2\left(x_1^2 + x_2^2\right) - 27r^4 = 0 \qquad (3.24)$$

perturbed by additive noise $\sigma\eta(t)$ where σ is the intensity of the noise. The stochastic limit cycle oscillator can be obtained by inserting (3.24) into (3.20) where the partial derivatives are given by

$$\frac{\partial H}{\partial x_1} = 4(x_1^2 + x_2^2)x_1 + 24rx_2^2 - 24rx_1^2 + 36r^2x_1$$

and

$$\frac{\partial H}{\partial x_2} = 4(x_1^2 + x_2^2)x_2 + 48rx_1x_2 + 36r^2x_2$$

Figure 3.10 illustrates a numerical simulation of (3.24) with parameters as per the figure caption.

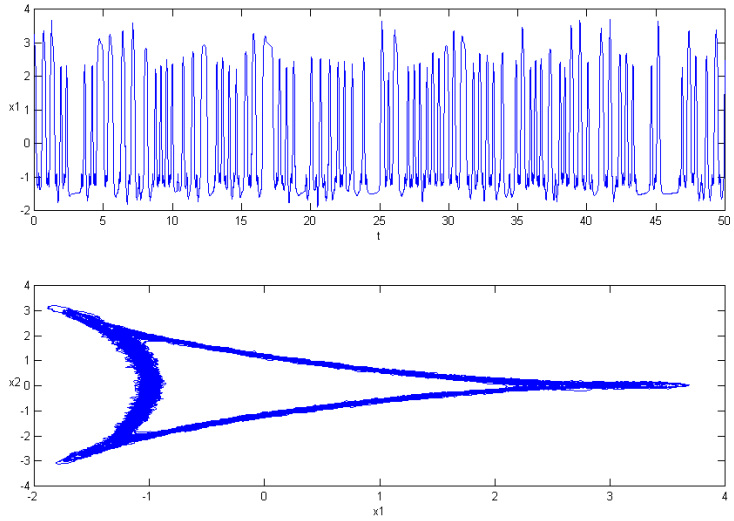

Figure 3.10: Stochastic Steiner's Hypocycloid Limit Cycle Oscillator. Simulation parameters: Runge-Kutta, $r=1$, $\lambda=0.1$, $\sigma=0.2$, $\Delta t = 0.0005$, $T=50$

3.6 Modelling the EEG Signal

In this section we will construct a mathematical model of the EEG signal based on the concept of stochastic limit cycle oscillators. This model might be termed an empirical model: we are attempting to find a simple mathematical model to describe an ordinary language model of the EEG i.e. "that the EEG is a manifestation of some equilibrium under noisy perturbation" where we interpret the equilibrium as limit cycle behavior. Limit cycles have been used as a model for electroencephalogram rhythms previously by DEWAN (1964) and derived from biophysical reasoning by WILSON and COWAN (1972) and FREEMAN (1975). Of course a limit cycle in the usual dynamical systems theory meaning does not produce a convincing model of the EEG since the real EEG signal does not exhibit regular oscillations and hence many of the models in the literature merely serve to indicate the *presence* of limit cycle behaviour. We create a model of the EEG comprising multiple *stochastic* limit cycle oscillators corresponding to the known spontaneous rhythms. The model is based on the construction above (3.20) with four independent oscillators to represent the theta, delta, alpha, and beta oscillations. Data from the model is visually compared to real EEG data by employing time series plots, 2-d phase plots, power spectral density plots, and amplitude histograms. In Chapter IV, we revisit the model using techniques from nonlinear time series analysis to enable more advanced comparisons with the actual EEG to be made.

3.6.1 Multiple Stochastic Limit Cycle Model of the EEG

Our model of the EEG is a stochastic one. Stochastic effects are ubiquitous in the brain; in fact neurons themselves require aperiodic activity simply to survive (FREEMAN, 1999). We impose some structure on the underlying stochastic process by regarding the EEG as the summated output of four stochastic limit cycle oscillators. The goal is to create an empirical model based on the noisy modulation of equilibria (limit cycles), which generates aperiodic solutions that simulate the statistics, spectra, and visually displayed patterns of the EEG. The proposed model is based on constrained noise, and is quite different to another type of common constraint called band-limiting. Clearly this model contrasts with other mechanisms such as deterministic chaos for generating aperiodic activity, e.g. as suggested by FREEMAN (1992) to model the EEG.

To construct the oscillators, we employ the simple topology of an ellipse

$$H(x_1, x_2) = \frac{x_1^2}{a^2} + \frac{x_2^2}{b^2} - 1 = 0 \tag{3.25}$$

Including the perturbation effects of additive noise $\sigma\eta(t)$ on the contour, the resultant vector stochastic differential system using (3.20) is

$$dX = \begin{pmatrix} \frac{2X_2}{b^2} - \lambda\left(\frac{X_1^2}{a^2} + \frac{X_2^2}{b^2} - 1\right)\frac{2X_1}{a^2} \\ -\frac{2X_1}{a^2} - \lambda\left(\frac{X_1^2}{a^2} + \frac{X_2^2}{b^2} - 1\right)\frac{2X_2}{b^2} \end{pmatrix} dt + \sigma \begin{pmatrix} -\lambda\frac{2X_1}{a^2} & 0 \\ 0 & -\lambda\frac{2X_2}{b^2} \end{pmatrix} dW \tag{3.26}$$

3.6.1.1 Time-Scaling Stochastic Differential Equations

To produce oscillators of different frequency we must time scale (3.26). However, since this is not an ODE, the usual time-scaling techniques are not appropriate and we must use a modified approach. Suppose $X(t)$ is the solution of the SDE

$$dX = a(t, X)dt + b(t, X)dW(t)$$

then by the Itô formula (Section 3.3.1.2) $Y(t) = f(X(t))$ also satisfies a SDE

$$dY = \tilde{a}(t, Y)dt + \tilde{b}(t, Y)dW(t)$$

For time-scaling, we are interested in

$$Z(\tau) = X(t(\tau))$$

where $t(\tau)$ is some time change. For small *h* we have

$$
\begin{aligned}
Z(\tau+h)-Z(\tau) &= X(t(\tau+h))-X(t(\tau)) \\
&\approx \mathrm{N}\left([t(\tau+h)-t(\tau)]a(t(\tau),X(t(\tau))),[t(\tau+h)-t(\tau)]b^2(t(\tau),X(t(\tau)))\right) \\
&\approx \mathrm{N}\left(h\tilde{a}(t(\tau),Z(\tau)),h\tilde{b}^2(t(\tau),Z(\tau))\right)
\end{aligned}
$$

where

$$
\begin{aligned}
\tilde{a}(\tau,Z) &= a(t(\tau),Z)t'(\tau) \\
\tilde{b}(\tau,Z) &= b(t(\tau),Z)\sqrt{t'(\tau)}
\end{aligned}
\tag{3.27}
$$

and $Z(\tau)$ is the solution of the SDE

$$dZ = \tilde{a}(t,Z)dt + \tilde{b}(t,Z)dW(t)$$

The time change $t(\tau)$ may also be random (see Theorem 8.5.7 in ØKSENDAL (1998) for a proof of the time change formula for Itô integrals). To scale (3.26) we introduce

$$t = \varepsilon\tau$$

and by using (3.27) we may restate (3.26) as (replacing τ with t)

$$
dX = \varepsilon\begin{pmatrix} \dfrac{2X_2}{b^2} - \lambda\left(\dfrac{X_1^2}{a^2}+\dfrac{X_2^2}{b^2}-1\right)\dfrac{2X_1}{a^2} \\ -\dfrac{2X_1}{a^2} - \lambda\left(\dfrac{X_1^2}{a^2}+\dfrac{X_2^2}{b^2}-1\right)\dfrac{2X_2}{b^2} \end{pmatrix}dt + \sigma\sqrt{\varepsilon}\begin{pmatrix} -\lambda\dfrac{2X_1}{a^2} & 0 \\ 0 & -\lambda\dfrac{2X_2}{b^2} \end{pmatrix}dW
\tag{3.28}
$$

3.6.1.2 Calibrating Oscillator Frequencies

Using the formulation (3.28), it is possible to configure the four oscillators corresponding to the delta, theta, alpha, and beta bands. Table 3.1 illustrates the frequency ranges, medians, and normalised medians for the four rhythms in the EEG. The gamma frequency band > 30 Hz is omitted since it is not usually regarded as a spontaneous rhythm - being induced by certain visual stimuli for example (NIEDERMEYER and LOPES DA SILVA, 1999).

TABLE 3.1: EEG RHYTHM FREQUENCY RANGE, MEDIAN, AND NORMALISED MEDIAN

Rhythm	Range	Median	Normalised Median
Delta	0.1 – 3.5 Hz	1.8 Hz	1.000
Theta	4 – 7.5 Hz	5.75 Hz	3.194
Alpha	8 – 13 Hz	10.5 Hz	5.833
Beta	14 – 30 Hz	22 Hz	12.222

An alternative method of scaling the frequencies of the oscillator comes from modifying the radius of the prescribed contour (3.25) and thus the amplitude of the oscillator. To see this (focusing on the deterministic case $\sigma = 0$ first and letting $a = b$ to form a circular contour), we may substitute the solution

$$X_1(t) = r\cos(\theta)$$
$$X_2(t) = r\sin(\theta)$$

into (3.26). Solving for r,θ allows us to restate the system in polar coordinates

$$\frac{dr}{dt} = -\lambda\left(\frac{r^2}{a^2} - 1\right)\frac{2r}{a^2} \tag{3.29a}$$

$$\omega = \frac{d\theta}{dt} = -\frac{2}{a^2} \tag{3.29b}$$

Equation (3.29a) has two equilibria. A stable equilibrium for $r = a$ and an unstable equilibrium at the origin $r = 0$ (as already noted in proof of Theorem 2.7, Chapter II). The second equation (3.29b), however, elucidates something interesting: the square of the amplitude is inversely proportional to angular frequency. When stated in terms of power (proportional to square of the amplitude) this translates to the *1/f power law*, a phenomenon ubiquitous in nature and observed by the EEG spectrum (BARLOW, 1993; WRIGHT and LILEY, 1996; FREEMAN *et al*, 2000c). This property of the oscillator facilitates a parsimonious mechanism for setting the oscillator frequency by altering the noise-free amplitudes of each component such that the frequency of the component's oscillator matches the normalised median value in Table 3.1. As a result of this selection procedure, the model's spectrum consisting of the four components naturally demonstrates a *1/f* power distribution. Now, referring back to (3.28) we may infer that to maintain the same noise intensity over different radii of the circle contour, we must scale the diffusion term by the radius. Thus, we finally arrive at (3.30) for the equation of a single oscillator.

$$dX = \begin{pmatrix} \frac{2X_2}{a^2} - \lambda\left(\frac{X_1^2}{a^2} + \frac{X_2^2}{a^2} - 1\right)\frac{2X_1}{a^2} \\ -\frac{2X_1}{a^2} - \lambda\left(\frac{X_1^2}{a^2} + \frac{X_2^2}{a^2} - 1\right)\frac{2X_2}{a^2} \end{pmatrix} dt + \sigma \begin{pmatrix} -\lambda\frac{2X_1}{a} & 0 \\ 0 & -\lambda\frac{2X_2}{a} \end{pmatrix} dW \tag{3.30}$$

3.6.1.3 Simulation Results

Simulations were carried out using SDESolver with a 1.5 strong order Runge-Kutta algorithm (Section 3.4.3.3). A step size of 0.001 was used with a multiplier of unity (Section 3.4.4.2). Figure 3.11 illustrates (in clockwise order starting from upper left) a time series, 2-d phase plot, amplitude distribution, and power spectral density for the model output. The radius of the circle contour prescribed by the oscillators was chosen

according to $a = 1/\sqrt{f}$ where f is the normalised median frequency for the corresponding frequency band as illustrated in Table 3.1.

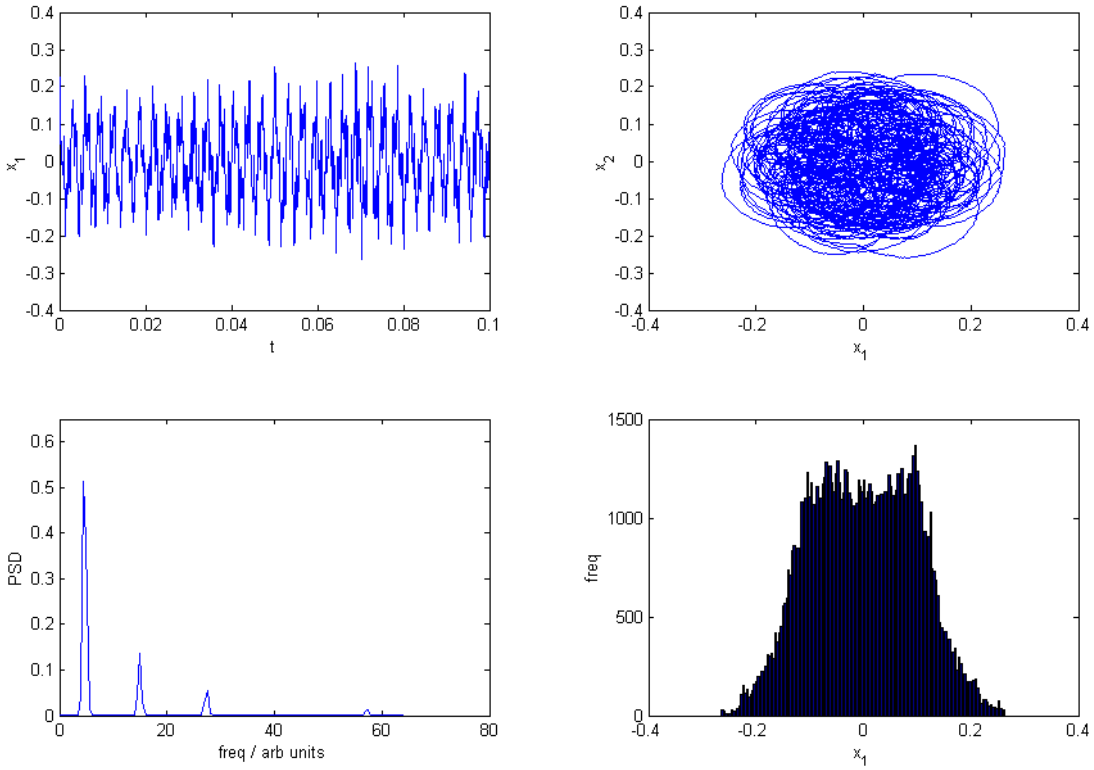

Figure 3.11: (In clockwise order) Time series, phase plot, amplitude histogram, PSD of the summated output of four stochastic limit cycle oscillators (3.30) corresponding to the delta, theta, alpha, and beta EEG components. Simulation details: Runge-Kutta, $\Delta t = 0.001$, $T = 100$, $\lambda = 0.01$, $\sigma = 4$. *1/f* power scaling used

In reality, the EEG power spectrum does not follow a strict *1/f* distribution, in particular it contains elevated spectral components at the alpha and beta frequencies (NIEDERMEYER, 1999; BARLOW, 1993; DUMERMUTH and MOLINARI, 1987). Accounting for this, Figure 3.12 illustrates the output where the power of both the alpha and beta components has been elevated by a factor of 2. For the purposes of comparison, EEG was acquired from the occipital region (Oz) with subjects seated in an upright position, eyes closed, and relaxed. Data was acquired at a sampling rate of 128 Hz and low-pass filtered with a cutoff frequency of 30 Hz. Figure 3.13 illustrates a comparison between actual EEG data and simulated model data. Finally, Figure 3.14 illustrates the time series, 2-d phase plot, amplitude distribution, and power spectral density of the actual EEG from Figure 3.13. The 2-D phase plot was achieved by time-delay embedding with a delay of 3 samples corresponding to 23.4 ms (ABARBANEL, 1996).

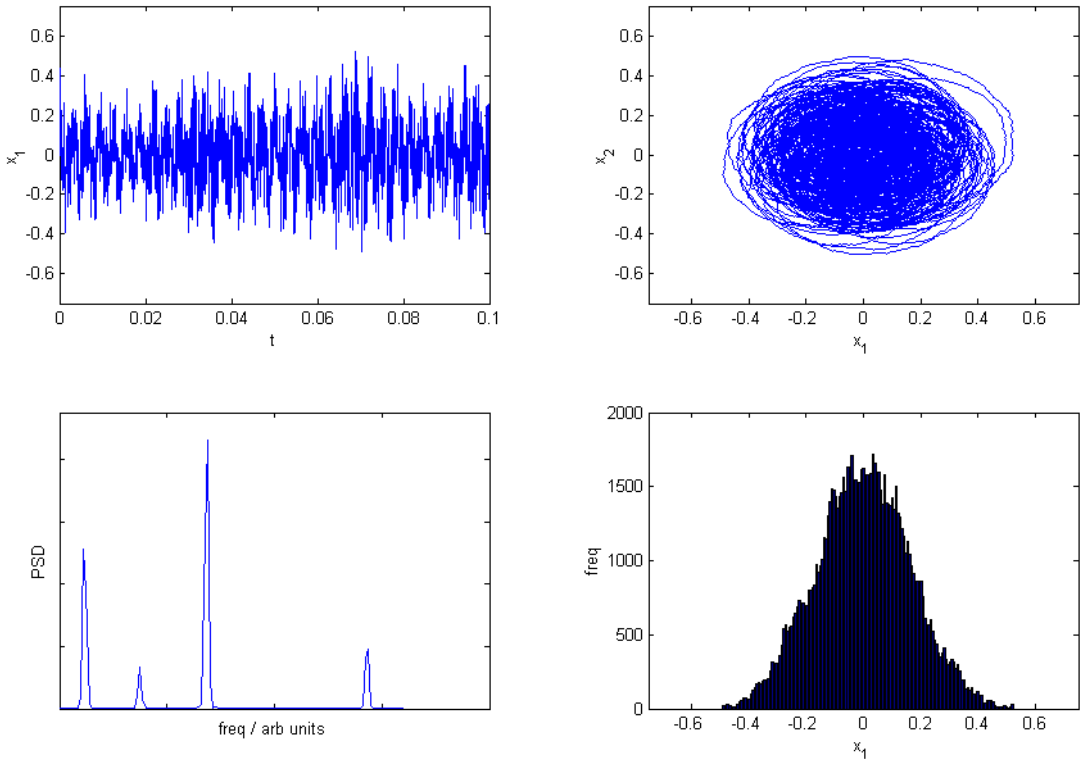

Figure 3.12: (In clockwise order) Time series, phase plot, amplitude histogram, PSD of the summated output of four stochastic limit cycle oscillators (3.30) corresponding to the delta, theta, alpha, and beta EEG components. Simulation details: Runge-Kutta, $\Delta t = 0.001$, $T = 100$, $\lambda = 0.01$, $\sigma = 4$. *1/f* power scaling used with x2 scaling of alpha and beta power

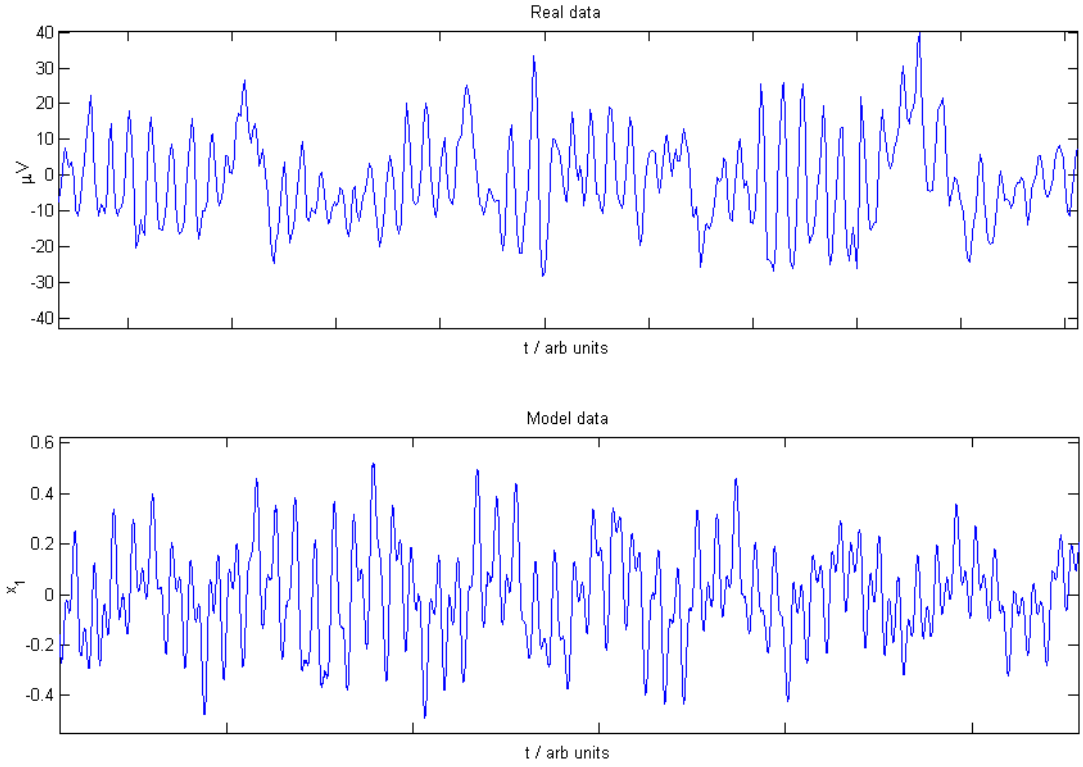

Figure 3.13: Comparison between actual EEG (top) and the simulated EEG (bottom) from the multiple stochastic limit cycle oscillator model. Simulation details: Runge-Kutta, $\Delta t = 0.001$, $T = 100$, $\lambda = 0.01$, $\sigma = 4$. *1/f* power scaling used with x2 scaling of alpha and beta power

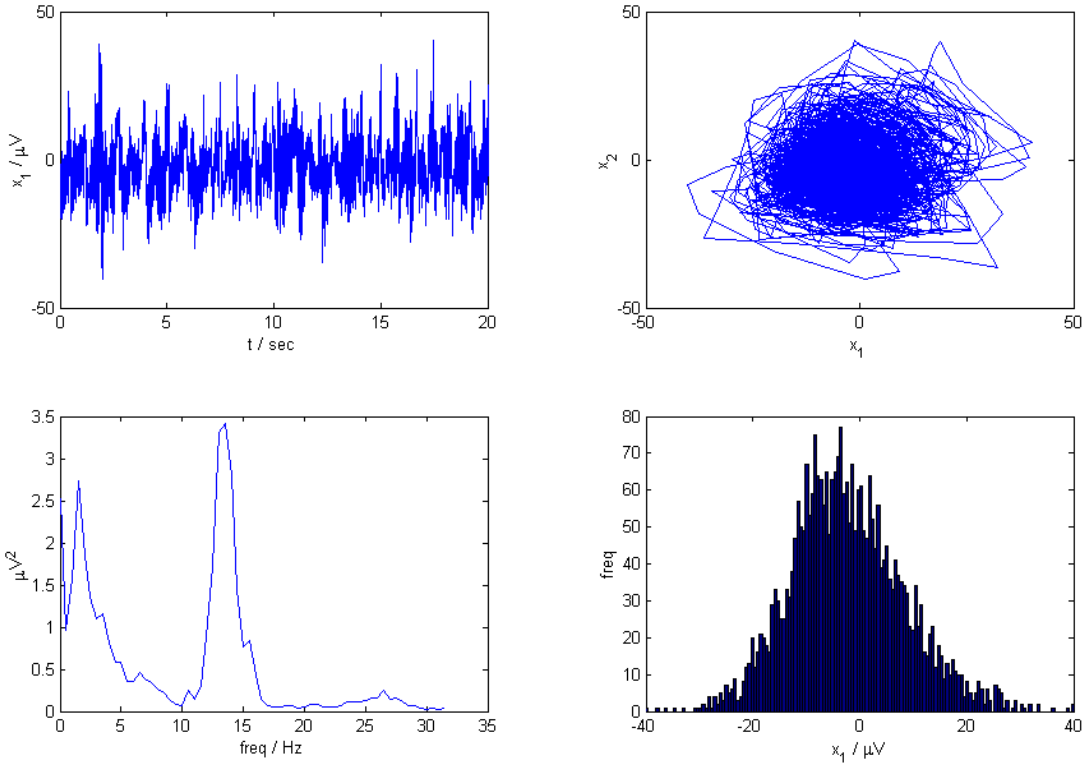

Figure 3.14: (In clockwise order) Time series, phase plot, amplitude histogram, and PSD of the actual EEG

3.6.1.4 Discussion

The visually displayed pattern of activity of the model agrees well with the actual EEG (in the awake, resting state from the occipital region) as is evident by Figure 3.13. In addition, the amplitude statistics appear in good agreement with the actual EEG, yielding a close approximation to a Gaussian distribution (the actual EEG is not always strictly Gaussian but approaches it under certain conditions (NIEDERMEYER and LOPES DA SILVA, 1999)). The 2-d phase plots of the model and actual EEG are also quite similar – an exact match would never be expected since the latter is a phase-space reconstruction and thus an invariant measure preserving mapping from the "true space". Nevertheless, a 2-d projection is a good tool for uncovering certain underlying structures (ABARBANEL, 1996).

The stochastic limit cycle model of the EEG is a natural extension to the models of WILSON and COWAN (1972) and DEWAN (1964) who employed deterministic limit cycle oscillators. These models clearly cannot produce the aperiodic activity and statistics produced by the current model. The outputs of the oscillators in our model are similar to that used by BARLOW (1993) where random modulation of extrema and slopes in the same proportion yield waves of constant frequency. The current model

differs from Barlow's being based on formal mathematics lending itself to analysis and extendable to any prescribed contour. Our model is in good agreement with Barlow who pointed out that random waves of constant frequency provide for an effective model the EEG with abundant alpha activity. That the EEG might be appropriately modelled as a stochastic limit cycle was reinforced by HERNÁNDEZ *et al* (1996). These authors showed that spike and wave activity from patients suffering from petit mal epilepsy could be accurately modelled by the output of a stochastic non-linear autoregressive model, which was shown to exhibit limit cycle behaviour when the noise component is removed. Finally, we note that we could have also included a white noise component in our model to more closely reflect measured EEG spectra. The white noise component is usually attributed mostly to measurement noise (DUMERMUTH and MOLINARI, 1987; ZETTERBERG, 1977).

3.7 Wilson-Cowan Stochastic Oscillator

The Wilson-Cowan model for densely interconnected excitatory and inhibitory populations of neurons was introduced in Section 1.3.1 (WILSON and COWAN, 1972). We now extend the model to account for noisy interactions from neighbouring populations by reformulating the original equations into a pair of stochastic differential equations. This resultant model builds on the previous section's purely mathematical approach for the EEG by providing a model with biophysical meaning. We perform a phase-plane analysis to consider the major phenomena elicited in the original work by Wilson and Cowan, namely the presence of multiple equilibria with hysteresis, damped sinusoidal impulse responses, and limit cycle behaviour, but this time taking stochastic effects into account.

The point of where to introduce stochastic effects into the original work by Wilson and Cowan is an interesting one; the equations are reproduced here for convenience

$$\begin{aligned}\tau_e \frac{dE}{dt} &= -E + (k_e - r_e E) S_e (c_1 E - c_2 I + P) \\ \tau_i \frac{dI}{dt} &= -I + (k_i - r_i I) S_i (c_3 E - c_4 I + Q)\end{aligned} \tag{3.31}$$

The original model employs two state variables (E, I) to represent a neural population consisting of two densely interconnected subpopulations of excitatory and inhibitory

neurons. The relevant variables are the proportion of cells of the subpopulations, which become active per unit time. This results in a deterministic model for the dynamics interpreted as a treatment of the mean values of the underlying statistical processes. Thus, introducing stochastic affects at the single neuron level does not make sense for this model. A more appropriate approach is to consider stochastic influences on the entire population resulting from neighbouring populations. For simplicity, we consider stochastic affects conducted by excitatory coupling only. The appropriate variable under consideration is *P*. We model stochastic effects by making the transformation

$$P \to P + \sigma\eta(t) \tag{3.32}$$

where η is white noise and σ is the noise intensity. The noise input is a variable of the activation function $S_e(x)$ which we take (as in the original work of Wilson and Cowan) to be the logistic sigmoid

$$S(x) = \frac{1}{1+\exp(-a(x-\theta)} - \frac{1}{1+\exp(a\theta)} \tag{3.33}$$

where *S(0) = 0* and the maximum slope of $a/4$ is attained at θ. Inserting (3.32) and (3.33) into (3.31) does not facilitate a SDE formulation. However, we may perform a Taylor expansion of $S_e(x)$ about $c_1E + c_2I + P$ to get

$$S_e(c_1E + c_2I + P + \sigma\eta) = S_e(c_1E + c_2I + P) + \sigma\eta S_e'(c_1E + c_2I + P) + h.o.t. \tag{3.34}$$

and for sufficiently small σ, we may neglect the higher order terms. Inserting (3.32) into (3.31), using (3.34), and replacing the noise term with the derivative of the Wiener process yields

$$\begin{aligned} \tau_e dE &= \left\{-E + (k_e - r_eE)S_e(c_1E - c_2I + P)\right\}dt + \left\{\sigma(k_e - r_eE)S_e'(c_1E - c_2I + P)\right\}dW \\ \tau_i dI &= \left\{-I + (k_i - r_iI)S_i(c_3E - c_4I + Q)\right\}dt \end{aligned} \tag{3.35}$$

where the derivative of (3.33) is given by

$$S'(x) = \frac{1}{\left[1+\exp(-a(x-\theta))\right]^2} a\exp(-a(x-\theta))$$

3.7.1 Hysteresis Behaviour

Wilson and Cowan suggested hysteresis constituting multiple stable equilibrium states as a basis for short-term memory. Without hysteresis, noisy perturbations could cause spurious transitions between stable equilibriums. Figure 1.6 illustrates a hysteresis loop formed by varying the excitatory input P. Our modification here to account for stochastic effects via (3.32) does not change this phenomenon – in fact hysteresis may be seen as a form of noise insensitiveness.

3.7.2 Limit Cycle Behaviour

Wilson and Cowan suggested limit cycle behaviour as a basis for EEG oscillations. However, their original model is unable to model the aperiodicity seen in the actual EEG. Figure 3.15 illustrates stochastic limit cycle behaviour for (3.35). Note that modelling stochastic effects exhibits random modulations of both the extrema and frequency of the resultant waveform E. That the frequency is also modulated is not surprising since Wilson and Cowan previously showed that the frequency of oscillation increases with P.

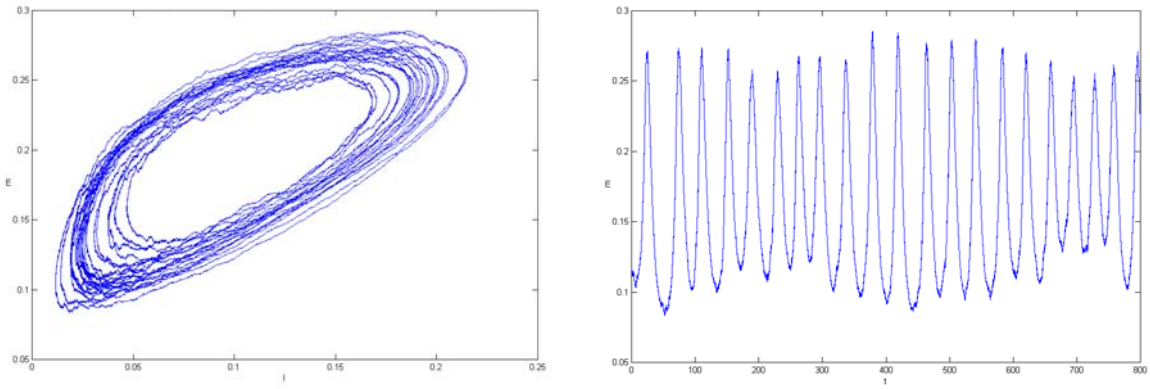

Figure 3.15: Limit Cycle behaviour produced by the stochastic Wilson-Cowan model (3.35). Parameters: $c_1 = 16$, $c_2 = 12$, $c_3 = 15$, $c_4 = 3$, $a_e = 1.3$, $a_i = 2$, $\theta_e = 4$, $\theta_i = 3.7$, $r_e = 1$, $r_i = 1$, $\tau = 8$, $P = 1.25$, $Q = 0$. Runge-Kutta solver: $\sigma = 0.015$, $\Delta t = 0.005$, $T = 800$

3.7.3 Damped Oscillatory Responses

For a certain range of parameters, the Wilson-Cowan model demonstrates a damped oscillatory response to a short impulse. An evoked potential may be interpreted as the difference between excitatory and inhibitory potentials measured in the neighbourhood of the recording electrode. Figure 3.16 illustrates four responses using a brief impulse to the excitatory input P with additive white noise (3.32). The result of modelling stochastic effects introduces randomness in the latency of the oscillatory response for different realisations of the Wiener process (i.e. experiment trials). Variation in

amplitude and latency across trials is well known in evoked and event-related potential recordings with many techniques for so-called latency corrected averaging having been suggested (e.g. GUPTA *et al*, 1996; YU *et al*, 1994). Interestingly, these techniques do not suggest any models to explain the latency, in fact they attempt to remove it in the averaging process, and it appears significant that the stochastic extension of the Wilson-Cowan model predicts the variation in latency.

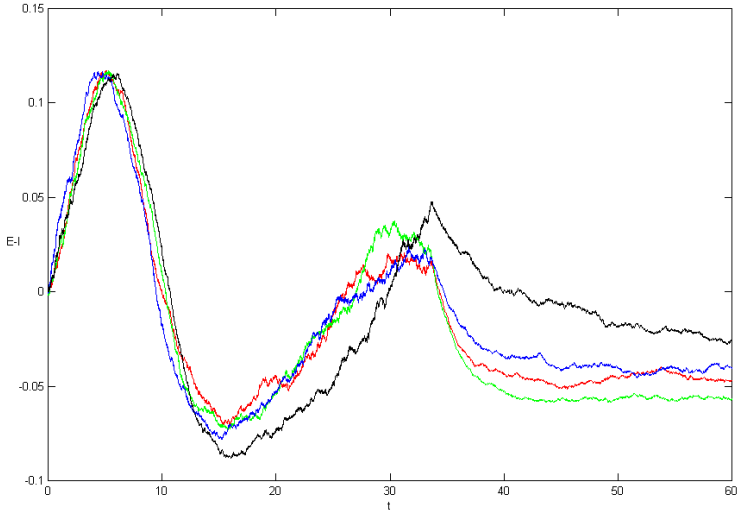

Figure 3.16: Damped oscillatory response to a brief, noisy stimulus produced by the stochastic Wilson-Cowan model (3.35). Parameters: $c_1 = 15$, $c_2 = 15$, $c_3 = 15$, $c_4 = 3$, $a_e = 1$, $a_i = 2$, $\theta_e = 2$, $\theta_i = 2.5$, $r_e = 1$, $r_i = 1$, $\tau = 10$, $Q = 0$. Runge-Kutta solver: $\sigma = 0.015$, $\Delta t = 0.005$, $T = 60$

3.8 Complex Behaviour From Simple Stochastic Oscillators

A major attraction of employing models of the EEG based on chaos is in the inherent ability of chaotic systems to produce complex aperiodic solutions from a small number of seemingly simple differential equations in addition to its suitability as a mechanism for information processing based on the existence of multiple attractors. The ability to produce complex aperiodic solutions from simple dynamical systems is not unique to chaotic dynamics and we demonstrate this here by employing a stochastic limit cycle oscillator possessing multiple homoclinic trajectories. Recall (from Chapter II) that a homoclinic trajectory is one which connects an equilibrium point with itself. For a 2-dimensional system the equilibrium point must be a saddle node and a good example of a simple oscillator with homoclinic trajectories is the Cassinian Ovals oscillator introduced in Example 2.4. By following the construction in Section 3.5 we can produce a stochastic limit cycle version by inserting

$$H(x_1, x_2) = \left[(x_1 - b)^2 + x_2^2\right]\left[(x_1 - c)^2 + x_2^2\right] - k \tag{3.36}$$

into (3.20). Figure 3.17 illustrates the time series and phase plots of a simulation with parameters as per the figure caption. As the trajectories approach the saddle node at the origin, they either continue around the outer contour or persist on the wing. The time series produced resembles that generated by the Lorenz attractor (although without the spiralling on each lobe), a popular demonstration model for chaos (THOMPSON AND STEWART, 1986) although here we are using a very different mechanism.

Chaotic activity has been proposed as a speculative mechanism for information processing in the brain by FREEMAN (1992). Sensitivity to initial conditions is suggested as a mechanism by which small microscopic inputs from sensory cells may elicit large-scale macroscopic responses across sensory cortex regions to facilitate Hebbian learning. In this model, a global attractor may exist for the sensory cortex developed over time by changes in synaptic strengths modifying the dynamics. In response to a particular stimulus, the system might jump to a corresponding wing of the attractor and this may be achieved with a minimal expenditure of energy.

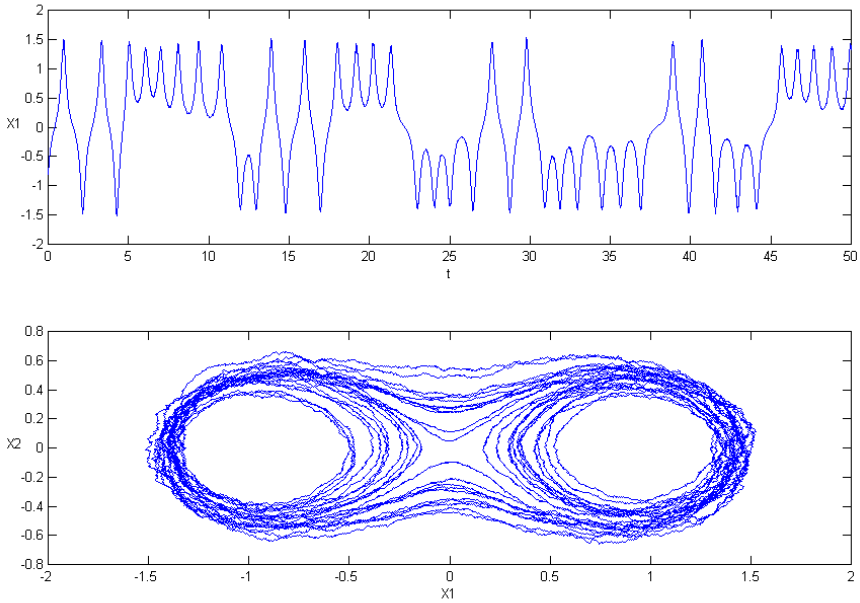

Figure 3.17: Cassinian Ovals stochastic limit cycle oscillator.
Runge-Kutta, $a=1, b=-1, k=1, \lambda=0.06, \sigma=0.5, \Delta t = 0.0005, T=50$

Freeman's model of information processing by chaotic mechanisms is also achievable by a stochastic limit cycle model similar to that illustrated in Figure 3.17. The only requirement on the system is that the attractor admits qualitatively different behaviour for small perturbations. This may be realised when the system is at a bifurcation point as is the case with the Cassinian Ovals oscillator for $k = 1$. In fact, a related argument

is made by HOPPENSTEADT and IZHIKEVICH (1997), whose modelling work is based on the concept that only those neurons participate nontrivially in the brain processes whose dynamics are at a bifurcation point, which they denote as the "Fundamental Theorem of Weakly Connected Neural Network Theory". The authors cite a number of mechanisms that might maintain neurons close to their bifurcation point – the most plausible might be the tonic background activity of neighbouring neurons. A natural extension of this theorem related to brain electrical activity might assume the existence of attractors close to a bifurcation; stochastic influences then admit aperiodic activity, and qualitatively different behaviour corresponding to information processing is also possible with the minimal expenditure of energy.

3.9 Summary

The construction for creating prescribed limit cycle oscillators (Method I) from the previous chapter was extended to the stochastic case by employing Itô calculus. The noise component was introduced in such a way that it actually interacts with the dynamics and might be termed *dynamical noise*. Itô calculus afforded us the ability to produce conditions for which the stochastic limit cycle behaviour may be described as an invariant set. A numerical toolbox (SDESolver) was created to solve SDEs and has proven useful throughout this work. This toolbox, being completely extensible and platform independent, should prove useful to other researchers.

The stochastic limit cycle oscillator was employed as a basis for creating a model of the EEG signal. Visually displayed patterns, spectra, and statistics of the model were shown to be in good agreement with real EEG data. The Wilson-Cowan model was also extended to the stochastic case and investigated via numerical techniques. The extended biophysically-based model was shown to produce stochastic limit cycle behaviour and variability in amplitude and latency of evoked potentials. Complex behaviour of a stochastic oscillator at a bifurcation was investigated with the motivation of demonstrating how the appealing features (for the purposes of modelling the EEG) of chaotic oscillators may also be approximated by stochastic oscillators.

In the next chapter, we employ nonlinear time series analysis in an effort to further validate the stochastic limit cycle oscillator model of the EEG against actual data.

CHAPTER IV

ANALYSIS IN RECONSTRUCTED PHASE SPACE

4.1 Introduction

This chapter is concerned with the examination of actual and model EEG data (the latter employing the stochastic limit cycle oscillator model from Chapter III) using *nonlinear time series analysis* (ABARBANEL, 1996; KANTZ and SCHREIBER, 1997). At this point, the reader might wish to review Section 1.4, which provides the motivation for the material presented here. Many of the studies in the literature regarding the nonlinearity of the EEG have been preoccupied with attempting to show that the aperiodic behaviour of the EEG might be ascribed to chaotic dynamics. In reality, a time series can be anything from purely random (independent identically distributed (i.i.d.) random variables) to strictly deterministic (KANTZ and SCHREIBER, 1997) and we shall study the EEG with this in mind.

The field of nonlinear time series analysis grew out of the pioneering work by TAKENS (1981). The core of the analysis focuses on reconstruction of a phase space from observed data known as *embedding*. Features of the system may be studied in reconstructed phase space by way of invariant measures such as attractor dimension, Lyapunov exponents, and entropies in addition to studying determinism by forecasting. The interested reader is referred to the excellent monographs of ABARBANEL (1996) and KANTZ and SCHREIBER (1997). A solid review of nonlinear dynamics is also recommended e.g. GUCKENHEIMER and HOLMES (1983).

Nonlinear time series analysis techniques are introduced by example with a model neuron exhibiting chaotic bursting behaviour. We subsequently investigate nonlinearity in the EEG by employing two principal techniques: correlation dimension (GRASSBERGER and PROCACCIA, 1983) and locally linear vs. globally linear (LL-GL) modelling (CASDAGLI, 1991). Results of the former have been reported in detail by many groups (e.g. BABLOYANTZ, 1985; BABLOYANTZ and DESTEXHE, 1986; PRITCHARD and DUKE, 1992; THEILER and RAPP, 1996), however the latter technique,

which forms a bridge between stochastic and deterministic methods, has only been briefly applied to the EEG. Here we apply both techniques to *qualitatively different* EEG data originating from the eyes-closed relaxed state, deep sleep, and petit mal absence seizures respectively. A statistical framework is employed facilitating comparison of the correlation dimension indication of nonlinearity with other published data and greatly facilitating our interpretation of the much less studied LL-GL technique. The significance of the LL-GL method is borne out of its direct practical consequence for EEG feature extraction, as we shall see later. In Section 4.4, we apply the same techniques to data obtained from the stochastic limit cycle oscillator model of the EEG. In particular, we are interested in how the nonlinearity present manifests itself in the results of phase space analysis techniques, which are commonly applied to EEG data in the literature. Finally, the chapter is concluded with a study of recent results regarding the applicability of nonlinear time series analysis (originally intended for deterministic systems) to stochastic data.

4.2 Nonlinear Time Series Analysis

It is instructive to introduce nonlinear time series analysis techniques by example on a model which we *know* possesses chaotic dynamics. A number of nonlinear time series analysis tools were reviewed before embarking on the work in this chapter; the TISEAN package (HEGGER *et al*, 1999) easily emerged as the superior toolbox and is used for the majority of the analyses that follow.

4.2.1 Hindmarsh-Rose Bursting Model Neuron

In this section we present a modified version of the FitzHugh-Nagumo model, itself a qualitative abstraction of the Hodgkin-Huxley model presented in Chapter I (HINDMARSH and ROSE, 1982; HINDMARSH and ROSE, 1984). The model equations are

$$\frac{dx}{dt} = y - x^3 + 3x^2 + I$$
$$\frac{dy}{dt} = 1 - 5x^2 - y \qquad (4.1)$$

where x represents the membrane potential, y is the recovery variable, and I is an externally applied current. Bistability is introduced by virtue of the fact the $\dot{y} = 0$ nullcline intersects the $\dot{x} = 0$ nullcline in three places to form three critical points: a stable node $x = -0.5(1+\sqrt{5})$, a saddle node $x = -1$, and an unstable node $x = +0.5(-1+\sqrt{5})$,

which is surrounded by a stable limit cycle as illustrated in Figure 4.1. Note that the stable manifold of the saddle point forms a separatrix and thus trajectories will approach the limit cycle for a sufficiently large perturbation from the resting state.

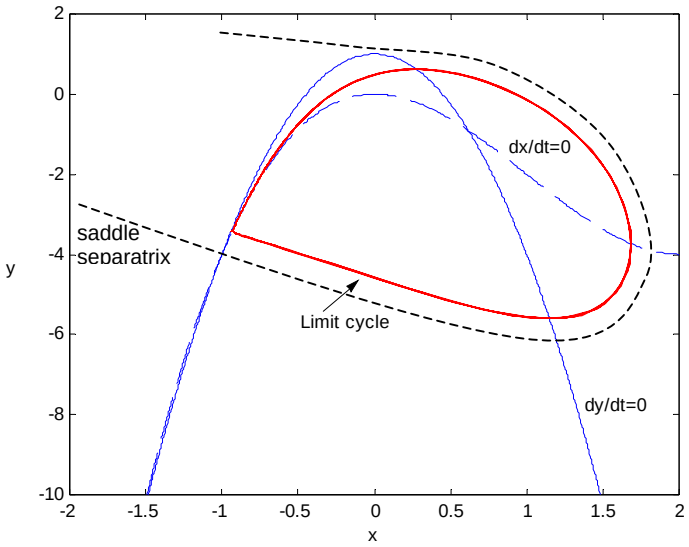

Figure 4.1: Phase plane for (4.1) with $I = 0$

To generate bursting, a slow variable is added (corresponding to the effect of an additional slow inward current) so that the membrane potential variable can be moved in and out of the bistable regime. The resulting equations are given by

$$\frac{dx}{dt} = y - x^3 + 3x^2 + I - z$$
$$\frac{dy}{dt} = 1 - 5x^2 - y$$
$$\frac{dz}{dt} = r\left[s(x - c_x) - z\right] \qquad (4.2)$$

The system exhibits chaos for parameter values: I = 3.281, c_x = -1.6, s = 4.0, r = 0.0021 (ABARBANEL *et al*, 1996). Figure 4.2 illustrates a time series plot of the x variable (representing membrane potential) and a 2-d phase plot projection (x,y) exhibiting chaotic dynamics.

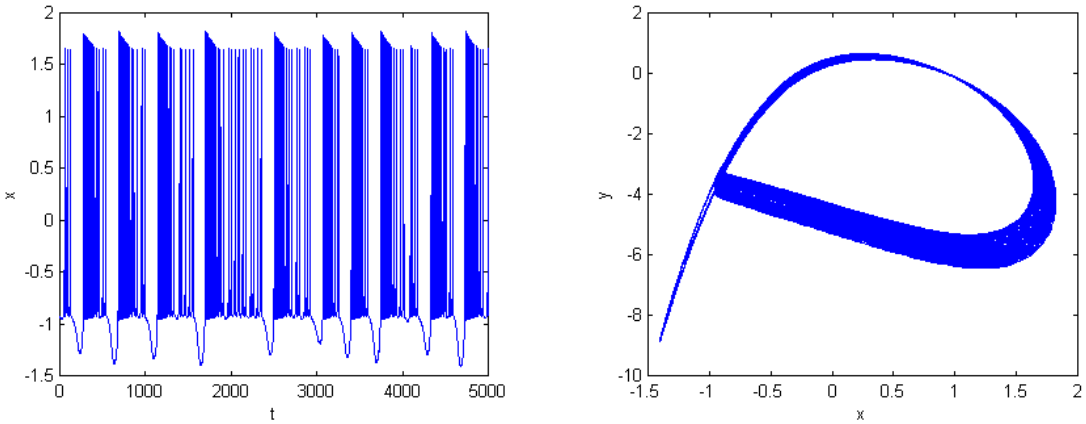

Figure 4.2: Time series x and phase plot (x,y) for (4.2)

Simulations were carried out via a variable step 4th order Runge-Kutta algorithm and the output was down-sampled to a step-size of 0.1. While isolated neurons can show chaotic signatures e.g. ABARBANEL *et al* (1996), it is worth noting that neurons embedded in the cortex typically generate pulse trains that are mostly random, exhibiting a Poisson distribution (FREEMAN, 1999). Nevertheless, the Hindmarsh-Rose model (4.2) serves its purpose here by providing a reference chaotic system to introduce nonlinear time series analysis techniques. There are alternative chaotic dynamical models we could have used here instead (e.g. Lorentz, Rössler etc) for which the techniques in the next section are very well behaved. The Hindmarsh-Rose model presented a challenge however, particularly for calculating a useful time delay for embedding and for accurately estimating the correlation dimension; the example should help alert the reader that nonlinear time series analysis is still a new field and requires careful application to achieve meaningful results.

4.2.2 Reconstruction of Phase Space

Embedding refers to a technique by TAKENS (1981) for transforming a scalar measured time series into a vector value time series. The sequence of vectors is interpreted as a trajectory of the dynamical system producing a *phase space reconstruction*. The reconstructed phase space is related to the unknown phase space of the underlying dynamical system by some smooth coordinate transform (not readily available in practice), assuming the measurement function was smooth. As a result, invariant measures such as Lyapunov exponents, attractor dimension, and entropies are the same for the reconstructed phase space and actual phase space. The *embedding theorem* may be summarised thus:

Theorem 4.1 (TAKENS, 1981; SAUER *et al*, 1991) Given are a dynamical system $\dot{x} = f(x)$ in a phase space $\Gamma \subset \mathbb{R}^d$, a measurement function $h : \mathbb{R}^d \to \mathbb{R}$, and a sampling interval Δt. Let the trajectory *x(t)* be confined to an *f*-invariant set $A \subset \Gamma$, with a box-counting dimension D_f. Denote the scalar measurements obtained through the sampling by $s_n := h(x(t = n\Delta t))$. Consider the delay embedding space spanned by delay vectors $S_n = (s_n, s_{n-\tau}, s_{n-2\tau}, ..., s_{n-(m-1)\tau})$. If $m > 2D_f$, then there exists a unique smooth map from A into the delay embedding space, which is invertible and has nonzero

derivative on the image of A in $\mathbb{R}^m$. A is then said to be immersed in $\mathbb{R}^m$. This holds for generic *h*, generic *f*, almost all Δt, and every $\tau \in N$.

The reconstructed phase space is formally equivalent to the actual phase space, i.e. for the purposes of evaluating invariant measures the reconstructed phase space is the same as actual phase space. From a heuristic point of view, the justification for using a time delay embedding can be seen by observing that information in higher order derivatives differs by simply adding a new delayed version of the signal e.g.

$$\frac{dx}{dt} \approx \frac{x(t+\Delta t) - x(t)}{\Delta t}$$

$$\frac{d^2x}{dt^2} \approx \frac{x(t+2\Delta t) - 2x(t+\Delta t) + x(t)}{\Delta t^2}$$

It is indeed possible to create an embedding by using a derivative coordinate embedding (as used in Section 2.3.2). Derivative coordinates are nothing but linear transformations of the standard delay coordinates (KANTZ and SCHREIBER, 1997). However, delay coordinates suffer from the problem that they accentuate noise and hence we do not use them here.

Theorem 4.1 does not suggest an optimal delay τ to use – in fact in theory, for an infinite amount of infinitely accurate data, any delay will work. In practice, this of course is not the case and we look more closely at this problem next. Theorem 4.1 also does not specify the ideal dimension *m* to use (the box-counting dimension can only be estimated after embedding – see Section 4.2.3). We review a method for estimating dimension in Section 4.2.2.2. The reader should keep in mind that visual interpretation of 2-d and 3-d reconstructed phase space is probably one of the most powerful tools in nonlinear time series analysis – the human eye is an exceptionally well-trained pattern recogniser! Visual analysis is recommended both as a first step and as a 'sanity check' for choosing delays and dimensions for reconstructing the phase space.

4.2.2.1 Finding the Optimal Delay

If we choose a time interval τ (to construct the lagged coordinates) too short the system will not have evolved significantly. On the other hand if the interval is too large, two consecutive points will be statistically independent of each other by virtue of the exponential divergence of neighbouring trajectories. A good prescription in practice (ABARBANEL, 1996) is to take the abscissa coordinate value at the first

minimum in average mutual information versus time delay between pairs of samples. Given a source $A=\{a_i\}$, and a related measure $B=\{b_j\}$, the average mutual information can be written as

$$I_{AB} = H_A - H_{A/B}$$

where

$$H_A = \sum_i P(a_i) \log_2 \left(\frac{1}{P(a_i)} \right)$$

is the entropy of source A and

$$H_{A/B} = \sum_{i,j} P(a_i, b_j) \log_2 \left(\frac{1}{P(a_i / b_j)} \right)$$

is the conditional entropy – the uncertainty remaining about A after observing B. In terms of our problem domain, the average mutual information can be written as

$$I(\tau) = \sum_{s_n, s_{n+\tau}} P(s_n, s_{n+\tau}) \log_2 \left[\frac{P(s_n, s_{n+\tau})}{P(s_n) P(s_{n+\tau})} \right] \tag{4.3}$$

where τ is an integer. Note, for two unrelated observations $P(s_n, s_{n+\tau}) = P(s_n) P(s_{n+\tau})$ the mutual information is zero. For two trivially identical observations the mutual information just reflects the entropy. Equation (4.3) may be thought of as a kind of nonlinear autocorrelation function. Figure 4.3 illustrates a plot of (4.3) for the Hindmarsh-Rose model (4.2).

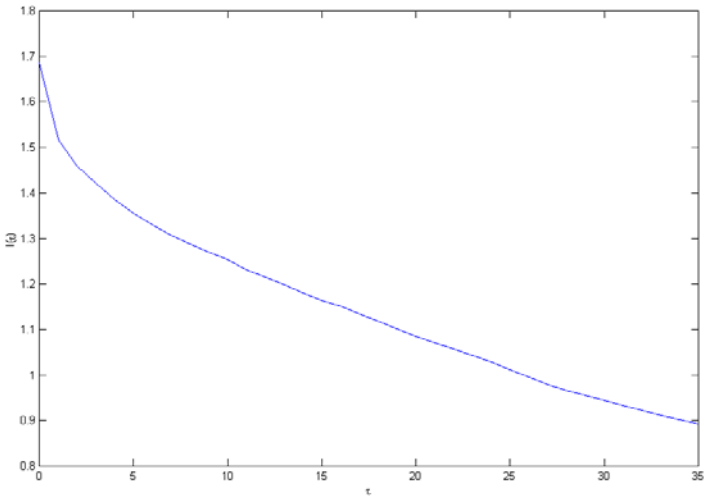

Figure 4.3: Average mutual information plot for x in (4.2)

In this case there is no immediately discernable minimum, as occurs for the classic Lorentz or Rössler attractors for instance. A difference from these other models might have been expected since it is chaotic *bursting* behaviour that (4.2) exhibits. We must

resort to estimation of a sensible time delay by inspecting the reconstructed phase space - too small a delay will yield vectors clustered in a diagonal, too large a delay will yield an erroneous spread of the vectors with the structure being destroyed by even small noise contributions (KANTZ and SCHREIBER, 1997). Phase plots in 3-d suggest a delay of 30 samples as a reasonable compromise.

4.2.2.2 Finding the Optimal Dimension

Theorem 4.1 states that any embedding dimension $m > 2D_f$ is valid for the phase space reconstruction. In many circumstances a smaller dimension can be found. However, too small a dimension will result in erroneous overlaps of the attractor. The technique employed here for embedding dimension estimation is called *false nearest neighbours* (ABARBANEL, 1996) and is based upon the concept that a pair of points constituting false neighbours are close together as a result of projection on a dimension that is too small to properly unfold the attractor. Projecting to larger embedding dimensions will result in the elimination of the false neighbours. We call the minimum dimension where all false neighbours are eliminated the *embedding dimension* d_E. Denoting a vector in reconstructed phase space as

$$S_n = (s_n, s_{n-\tau}, ..., s_{n-(m-1)\tau})$$

we may write its nearest neighbour in dimension m as

$$\hat{S}_k = (s_k, s_{k-\tau}, ..., s_{k-(m-1)\tau})$$

The square of the Euclidean distance between nearest neighbours in d is

$$R_d^2[n] = \sum_{m=1}^{d} \left(\hat{S}_{n-(m-1)\tau} - \hat{S}_{k-(m-1)\tau} \right)^2$$

The difference in distance between nearest neighbours in dimension $d+1$ compared with the distance in dimension d can be written as

$$\left(\frac{R_{d+1}^2[n] - R_d^2[n]}{R_d^2[n]} \right)^{1/2} = \frac{\left| \hat{S}_{n-\tau d} - \hat{S}_{k-\tau d} \right|}{R_d[n]}$$

When this value is greater than a certain threshold, we consider the nearest neighbours to be false neighbours. In practice, a factor of ~15 works well and results are quite insensitive to variations in this (ABARBANEL, 1996). Figure 4.4 illustrates the fraction of false nearest neighbours as a function of dimension suggesting an embedding dimension of $d_E = 3$, which agrees with the actual dimension of (4.2). The false nearest

neighbour method degrades slowly and gracefully under the influence of noise and thus is quite suited for application to real data (ABARBANEL, 1996).

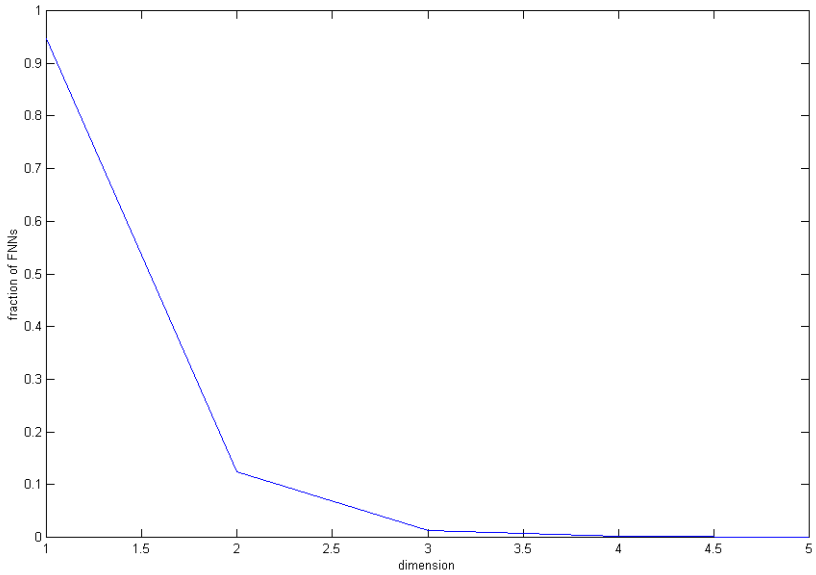

Figure 4.4: Fraction of FNNs for x in (4.2)

4.2.3 Invariant Measures

After a successful embedding, according to Theorem 4.1, we end up with a finite sample of points on an invariant set. We wish to characterise this invariant set by measures that are independent of the details of the measurement process and the reconstruction of the state space. We call such estimates *invariant measures*, where in particular we assume ergodicity (values are independent for almost any choice of initial conditions) so that quantities may be defined as averages over the *natural measure* of the phase space (KANTZ and SCHREIBER, 1997). By averages over the natural measure in phase space we mean averages centred around the data points. The common invariant measures used in nonlinear time series analysis include dimensions, Lyapunov exponents, and entropies. Not all invariant measures can be obtained robustly from data. In this section, influenced by the EEG literature, our main focus is on the correlation dimension. We briefly touch on Lyapunov exponents, a simple concept that helps build intuition about chaotic dynamics.

4.2.3.1 Correlation Dimension

Dissipative attractors for chaotic systems possess a complicated, *strange* geometry and are *fractals* in the sense of MANDELBROT (1985). Non-integer dimensions exhibit *self-*

similarity on all length scales. The classic mathematical example of self-similarity is the Cantor set. A chaotic attractor, on the other hand, only shows self-similarity locally and possesses position dependent scaling factors (KANTZ and SCHREIBER, 1997). Since a strange attractor exhibits a fractal dimension it is thus an interesting property of the chaotic system to measure. Consider a point set located in $\mathbb{R}^m$. If we cover the set with a grid of boxes of length ε and denote the number of boxes which contain at least one point by $M(\varepsilon)$, then for a self-similar set we obtain

$$M(\varepsilon) \propto \varepsilon^{-D_f}, \quad \varepsilon \to 0$$

where D_f is called the *box counting dimension*. There are other ways to define and generalise dimensions e.g. based on the Renyi entropies (ABARBANEL, 1996; KANTZ and SCHREIBER, 1997). In practice, the most robust and efficient estimate of dimension is the *correlation dimension* proposed by GRASSBERGER and PROCACCIA (1983). We first define the correlation sum for a collection of points x_n in some vector space as the fraction of all possible pairs of points which are closer than a given distance ε

$$C(\varepsilon, N) = \frac{2}{N(N-1)} \sum_{i=1}^{N} \sum_{j=i+1}^{N} \Theta(\varepsilon - \|x_i - x_j\|)$$

where Θ is the Heaviside step function, $\Theta(x) = 0$ if $x \le 0$ and $\Theta(x) = 1$ if $x > 0$. Taking $N \to \infty$ and for small ε we observe scaling according to $C(\varepsilon) \propto \varepsilon^D$

We define the correlation dimension *D* by

$$d(N, \varepsilon) = \frac{\partial \ln C(\varepsilon, N)}{\partial \ln \varepsilon}$$

$$D = \lim_{\varepsilon \to 0} \lim_{N \to \infty} d(N, \varepsilon)$$

This definition coincides with the usual notion of dimension when applied to non-fractal objects. In practice, there is often a limited amount of data available, a lower bound on the size of ε due to finiteness of samples, and an upper bound on the size ε due to the finite size of the attractor. As a consequence, we have to select a region (if any) where $C(\varepsilon) \propto \varepsilon^D$ by visual inspection. This precludes the use of robust automatic methods for dimension estimation and due to subjectivity, is often a source of errors (KANTZ and SCHREIBER, 1997).

The correlation sum represents the probability that a pair of randomly chosen points on the reconstructed attractor are less than a certain distance apart, where we assume that the distance between pairs of points is due solely to geometry. In reality,

dynamically correlated points will also be close in space. To avoid the effects of this temporal correlation, we introduce the Theiler window n_{min} (THEILER, 1986) to suppress trivial pairs of points. In practice, we estimate the correlation sum over a range of embedding dimensions (the correct dimension D can only be seen for $m > D$) and thus (including the Theiler window) we now restate the correlation sum as

$$C(m,\varepsilon) = \frac{2}{(N-n_{\min})(N-n_{\min}-1)} \sum_{i=1}^{N} \sum_{j=i+n_{\min}}^{N} \Theta\left(\varepsilon - \left\|S_i - S_j\right\|\right)$$

A natural choice for n_{min} might be the time when the autocorrelation function goes to zero. However, this only measures linear correlations. Happily, the problem of selecting a safe value for n_{min} has been solved by PROVENZALE *et al* (1992) using a *space-time separation plot*. The technique involves plotting contour maps of the fraction of points closer than a distance ε at a given time separation Δn i.e. we plot constant curves of

$$P\left(\left\|S_{n+\Delta n} - S_n\right\| < \varepsilon\right)$$

as a function of Δn. Figure 4.5 illustrates a space-time separation plot for x in (4.2). The contour lines may be interpreted in how far one must go to find a given fraction of pairs for a particular separation in time Δn. The plot suggests a Theiler window of 120.

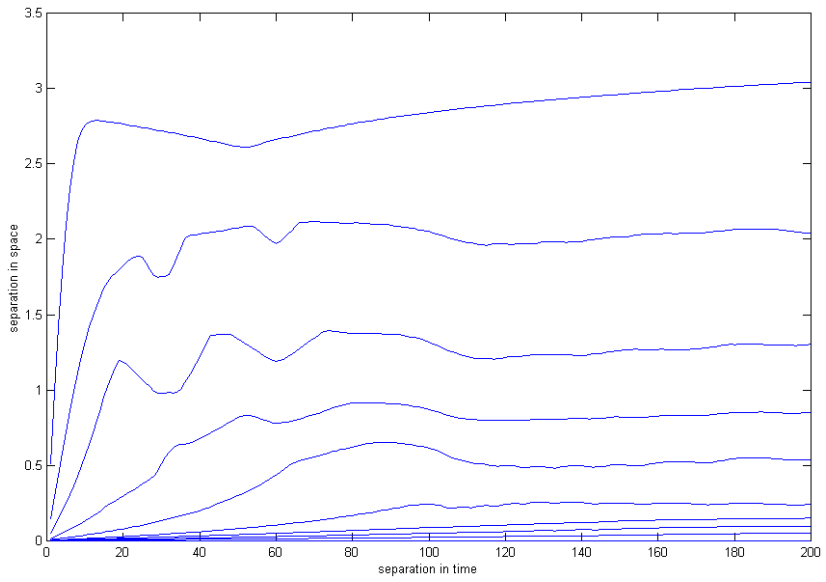

Figure 4.5: Space-time plot for x in (4.2). Fractions of pairs 0.1, 0.2, etc (bottom to top) for a given space and time resolution. Note saturation for $\Delta n \approx 120$

Figure 4.6 illustrates the correlation sum (left) and corresponding dimension values (right) for the chaotic Hindmarsh-Rose model (4.2). We can locate a scaling region that yields an estimate of $D = 1.98 \pm 0.1$.

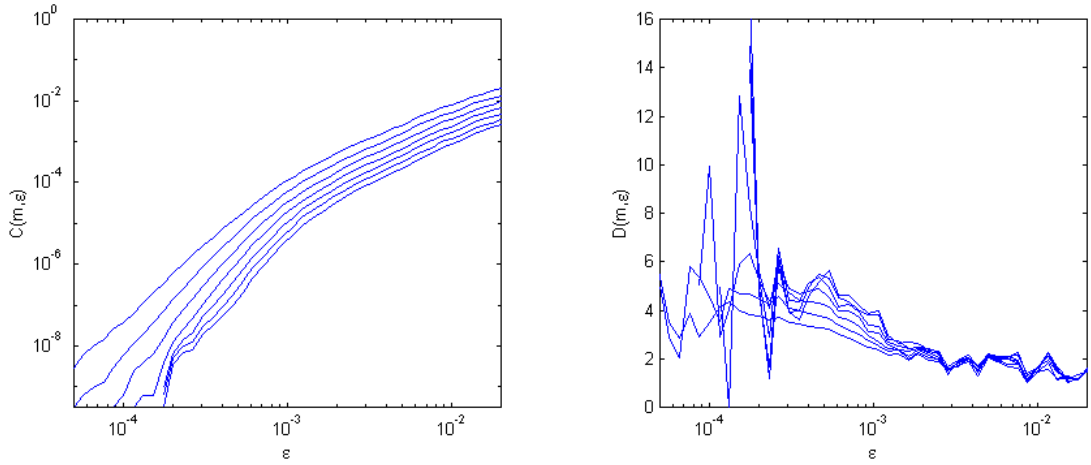

Figure 4.6: Correlation sum (left) and local slopes (right) for dimensions m = 3,…,9 (left: top to bottom; right: bottom to top) for x in (4.2)

For a typical plot of $D(m,\varepsilon)$ vs. ε, one may differentiate four regions. For large ε effects of the finite size of the attractor are seen. On a smaller scale the scaling region should exhibit the correct value for D. On smaller scales again, noise dominates (particularly for large dimensions) and finally on the smallest scale there are large fluctuations due to the finiteness of the samples.

4.2.3.2 Lyapunov Exponents

The hallmark of chaos is *sensitivity to initial conditions*, that is, exponential divergence of neighbouring trajectories. The average exponent of the divergence is known as the Lyapunov exponent. Let S_{n_1} and S_{n_2} denote two infinitesimally close points in state space such that $\left\|S_{n_1} - S_{n_2}\right\| = \delta_0$. The distance $\delta_{\Delta n}$ at a time in the future is

$$\left\|S_{n_1+\Delta n} - S_{n_2+\Delta n}\right\| = \delta_{\Delta n}$$

The maximal Lyapunov exponent λ may be defined as

$$\lim_{\delta_0 \to 0} \lim_{\Delta n \to \infty} \delta_{\Delta n} = \delta_0 e^{\lambda \Delta n}$$

The two limits are required to avoid saturation effects due to distance. A positive exponent is observed for chaos (and imposes a limit on the prediction horizon – see next section), a negative exponent is observed for a stable fixed point, and an exponent of zero is observed for flows constrained to a volume of phase space. There are in fact as many Lyapunov exponents as there are dimensions in the system. The full spectrum of Lyapunov exponents may be obtained by determining the product of the Jacobians at every point of the trajectory and obtaining the logarithm of the eigenvalues (actually

this is somewhat more technical than described here, requiring the formation of the Oseledec matrix – see ABARBANEL (1996) for more details).

In this chapter, we are not concerned with estimating Lyapunov exponents from data – the interested reader is referred to the methods by WOLF *et al* (1985), ECKMANN *et al* (1986), and KANTZ (1994).

4.2.4 Prediction

For a deterministic system, knowledge of the current state and its dynamics (equations of motion) are all that is required to predict the next state with complete certainty. For a chaotic system, the same holds but our prediction horizon is adversely affected by the sensitivity to initial conditions. In this section we discuss how to construct models of nonlinear dynamical systems. A successful model may be used for prediction, control, or used to create an arbitrary amount of artificial data to study (the bootstrapping method). Our interest in prediction stems from the fact that it is a very robust indication of determinism (KANTZ and SCHREIBER, 1997) and we employ it in the next two sections.

FARMER and SIDOROWICH (1987) were the first to demonstrate how to extract equations of motion from the delay embedding space. A deterministic dataset sampled at discrete times is described by

$$\hat{s}_{n+1} = G(S_n) \tag{4.4}$$

Denote by $\Gamma_n \subset \Re$ a small neighbourhood of S_n. To find the predicted value of S_n the *zeroth order approximation* employs the average of the future neighbours in reconstructed phase space i.e.

$$\hat{s}_{n+1} = \left\langle S_{j+1} \right\rangle_{\Gamma_n}$$

We can extend this to produce a *local linear model* (valid in a small neighbourhood of the reference point)

$$\hat{s}_{n+1} = a_n + b_n S_n$$

where the subscripts on a and b emphasise that we create a local linear model for each vector n in reconstructed phase space. We select a and b such that they minimise

$$e_n^2 = \sum_{S_{j \in \Gamma_n}} (\hat{s}_{n+1} - a_n + b_n S_j)$$

The collection of all these local linear models constitutes the global nonlinear dynamics and may be interpreted as linearisations of the global dynamics at various points in the phase space.

Another method of modelling involves choosing an appropriate functional form for *G* in (4.4) such that it is capable of modelling the *whole* attractor. A popular method is to consider decompositions onto basis functions and perform an optimisation of the weighted superposition of the basis functions during the fitting procedure (KANTZ and SCHREIBER, 1997). Choices for *G* include polynomials and radial basis functions, with neural networks providing for a possible mechanism for optimisation. This procedure can be considered parametric (as opposed to the previous non-parametric techniques) and could be of utility in feature extraction. More explicitly, signals for which nonlinear time series techniques are well suited could be fitted with an appropriate functional form for *G* and its free parameters used as features for pattern recognition applications.

4.2.4.1 Locally Linear vs. Globally Linear Modelling

Although a time series can be anything from purely random (independent identically distributed (i.i.d.) random variables) to strictly deterministic, there only exist established methods to optimally exploit *either* linear correlations or nonlinear determinism (KANTZ and SCHREIBER, 1997). Our hypothesised stochastic limit cycle oscillator model of the EEG, where noise contributions actually interact with the dynamics, produces a time series that is at neither of the two extremes. In this section we present a recent technique by CASDAGLI (1991) that attempts to bridge the gap between stochastic and deterministic approaches and facilitates an exploratory approach to the investigation of the dynamics underlying a time series. In the following sections we build on the preliminary investigations of Casdagli with application to the EEG and apply the method to the stochastic limit cycle oscillator model of the EEG from Chapter III.

Consider as before, a dataset sampled at discrete times described by (4.4). We wish to construct local linear approximations to *G* by using a variable number *k* of neighbours.

A small value of *k* corresponds to a deterministic approach to modelling (as described in the previous section) and a maximum value of *k* corresponds to fitting a stochastic linear autoregressive model to the data. *Intermediate values of k correspond to fitting non-linear stochastic models*. To apply the procedure, we select a test point $\tilde{s}_i$ and its associated delay vector $\tilde{S}_i$ for a Δn-forecasting test. The *fitting set* is defined as all vectors that do not include the component $\tilde{s}_i$ (i.e. we wish to perform an out-of-sample test). The *k* nearest neighbours $S_{j(1)},...,S_{j(k)}$ of $\tilde{S}_i$ from the fitting set are obtained and the following autoregressive (AR) model fitted

$$s_{j(l)+\Delta n} = \alpha_0 + \sum_{n=1}^{m} \alpha_n s_{j(l)-(n-1)\tau}, \quad l=1,...,k \tag{4.5}$$

where α_i are the AR coefficients, *m* is the embedding dimension, and τ is the delay time (for more information on AR models, see Chapter VI). We use (4.5) to generate a Δn forecast $\hat{s}_{i+\Delta n}(k)$ for the test point $\tilde{s}_i$ and write the error as

$$e_i(k) = \left|\hat{s}_{i+\Delta n}(k) - s_{i+\Delta n}\right|$$

The normalised error over the entire test set is

$$E(k) = \frac{1}{\sigma}\left[\sum_i e_i^2(k)\right]^{\frac{1}{2}} \tag{4.6}$$

where σ is the standard deviation of the time series. We refer to graphs of *E(k)* versus *k* as LL-GL plots.

Figure 4.7 illustrates (superimposed on the same plot for ease of comparison) the LL-GL plot for the Hindmarsh-Rose model and for the AR(2) process given by

$$x_{n+1} = 1.985x_n - 0.995x_{n-1} + \xi_n$$

The chaotic model demonstrates a clear deterministic signature with accurate predictions for small neighbourhoods - in fact the local linear (globally nonlinear) model yields two orders of prediction accuracy above the globally linear model. In contrast, the AR(2) model shows a diminishing error for increasing neighbourhood size. These examples demonstrate modelling at two ends of the scale. For a high dimensional attractor, and for limited data, there will be insufficient data to approximate (4.4) and we will be forced to approximate it with a nonlinear stochastic model (intermediate *k*) (CASDAGLI, 1991).

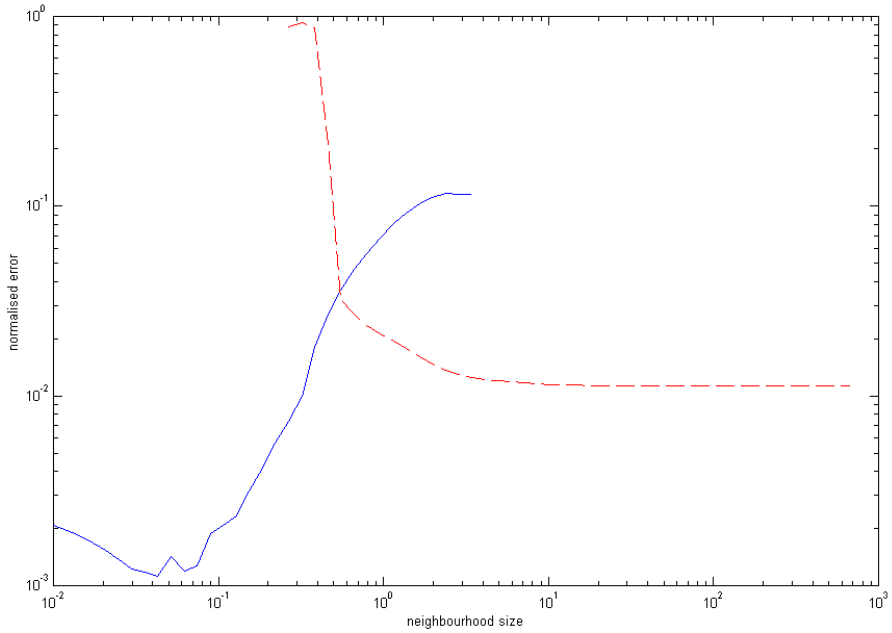

Figure 4.7: LL-GL plot for x in (4.2) (solid) and for the AR(2) process (dashed). See text for details

In general, an AR model cannot be recommended in practical applications just because the error $E(k)$ is minimised for large k. However, application to the EEG is a special case since the AR construction, being an all-pole model, reproduces the characteristics of the EEG exceptionally well and *is* suitable as a model for feature extraction and spectral estimation (PARDEY *et al*, 1996). Thus, the LL-GL allows us to reaffirm, with the added knowledge nonlinear time series analysis brings to bear, that the AR model is in fact optimum for these practical tasks.

4.3 Testing for Nonlinearity in the EEG

One might easily be forgiven for assuming that chaotic dynamics could underlie the EEG signal: the microscopic neuronal processes are themselves nonlinear, and chaotic solutions exhibit aperiodic activity with broadband spectra somewhat similar to that of the EEG. Indeed, many systems are composed of a huge number of internal microscopic degrees of freedom, but nevertheless produce signals which are found to be low dimensional (KANTZ and SCHREIBER, 1997). Nonlinear determinism is a necessary but not sufficient condition for chaos and thus detecting nonlinear determinism in the EEG is a prerequisite to discussions on chaotic dynamics. In this section we build on the preliminary investigations of CASDAGLI (1991) by applying the LL-GL technique to qualitatively different EEG data: eyes-closed relaxed state,

deep sleep, and petit mal absence seizures. In addition, we calculate the correlation dimension and compare with surrogate data using a robust statistical framework (SCHREIBER and SCHMITZ, 2000), which affords us some ability to relate our LL-GL results with the literature. Our approach here is best described as *exploratory* – we wish to better understand the characteristics of the dynamics exhibited by the EEG. In Section 4.4 we will apply the same techniques to the stochastic limit cycle oscillator model for comparison.

4.3.1 Surrogate Data and Statistical Comparison

Investigating supposed nonlinearities in data by absolute values of invariant measures quickly leads to conclusions open for interpretation and difficult to relate to other studies (this is particularly evident in the early studies of nonlinear analysis of the EEG). What is required is a proper statistical framework to distinguish features of a series with a specified level of significance. For a detailed review on testing for nonlinearity with surrogate time series see SCHREIBER and SCHMITZ (2000). Suppose we have some measure λ obtained from a time series, which we want to probe for nonlinearity. We require a *null hypothesis* to test against. A convenient null hypothesis is that the data was produced from a Gaussian linear stochastic process. The null hypothesis can be formulated by stating that all structure to be found in the time series is fully described by the mean, variance, and auto-covariance. Indeed, Gaussian linear stochastic processes are fully described by their power spectrum (by the Wiener-Khintchine relation), which does not contain any information pertaining to the direction of time. It is possible to obtain a randomised series with the same mean, variance, auto-covariance, and power spectrum by first taking the Fourier transform, randomising the phases, and subsequently taking the inverse Fourier transform. To test the null hypothesis, that is to *accept* or *reject* it, we must specify a level of significance. A significance level of $p < 0.05$ means that there is a 5% chance we reject the null hypothesis even though it is true – the test is then said to be valid at the 95% significance level. The use of parametric statistical tests e.g. those assuming a normal distribution for λ is discouraged since many nonlinearity measures do not follow a normal distribution (SCHREIBER and SCHMITZ, 2000). Instead, a rank-order test is recommended (THEILER *et al*, 1992). We select a probability α of false rejection, corresponding to a significance level of $(1-\alpha)\times 100\%$. For a *one-sided* test (e.g. the 'smallest value'), we generate $M = 1/\alpha - 1$ surrogate series. Therefore, including the

series in question, we have a total of $1/\alpha$ series. The probability that one of these series has the smallest value is exactly α, as initially prescribed. For a *two-sided* test we need $M = 2/\alpha - 1$ surrogate series. Assuming a significance of 95%, this translates to requiring 19 surrogate series for a one-sided test and 39 surrogate series for a two-sided test.

4.3.2 Results

The datasets in this section were investigated using the average mutual information and false nearest neighbours techniques as an initial guide to selecting delays and dimensions. Rather than be bound by these initial prescriptions, we follow the recommendation of HEGGER *et al* (1999) that what constitutes an optimal choice is largely dependent on the application (in this case prediction and dimension calculation). Filtering was performed using a 4th order FIR approach since IIR filters may introduce new dynamics and could confound results (ABARBANEL *et al*, 1993). Since endpoint mismatch results in artefacts in the Fourier transformation when creating surrogate data, subsets of data series were chosen such that the endpoints were matched in value and first derivative (SCHREIBER and SCHMITZ, 2000). For analysing correlation dimension, a one-sided test at the 95% level was performed (the null hypothesis may be rejected when the dimension of the original series is smaller than all the surrogates). Rather than extract a value for the dimension, we just plot the correlation sum. This way we avoid the dubious search and evaluation of a scaling region – a difficult task even for mathematically generated data (see Section 4.2.3.1).

In the preceding sections, we have avoided the discussion of nonstationarity. Needless to say, the techniques presented here assume stationarity. In the analyses that follow, we have endeavoured to maintain stationarity of the data under study. EEG datasets are limited to a maximum of 10 seconds and selected for constancy of spectra (e.g. uniform alpha, delta, or spike-wave activity).

4.3.2.1 EEG Data I (Eyes-Closed Relaxed State)

EEG data was recorded from normal subjects seated upright in a relaxed, eyes-closed state. A single channel from the occipital area (Oz) was sampled at a rate of 100 Hz and filtered digitally 0.1 Hz to 30 Hz. Artefact free segments consisting of 10 seconds (N = 1000) of data were selected for analysis. Figure 4.8 (left) illustrates the LL-GL

plot over a range of dimensions. The dynamics in this case show little evidence of nonlinearity, being optimally described by an autoregressive model. The global model shows a 300% improvement in forecasting accuracy compared with nonlinear techniques. Indeed, in Figure 4.8 (right) the correlation sum of the original series (solid) is compared with 19 surrogates to reveal no statistically significant ($p < 0.05$) low dimensionality.

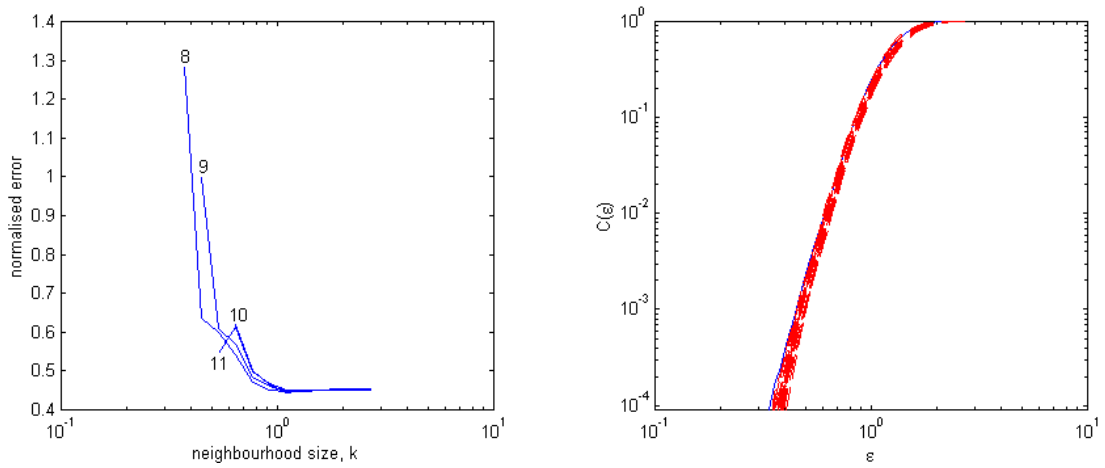

Figure 4.8: Left: LL-GL plot for eye-closed EEG, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 11. Right: correlation sum (solid) at m = 11 and for 19 surrogates (dashed). A time delay of τ = 6 was used for embedding. Theiler window n_m = 15

Figure 4.9 illustrates similar results for another EEG dataset.

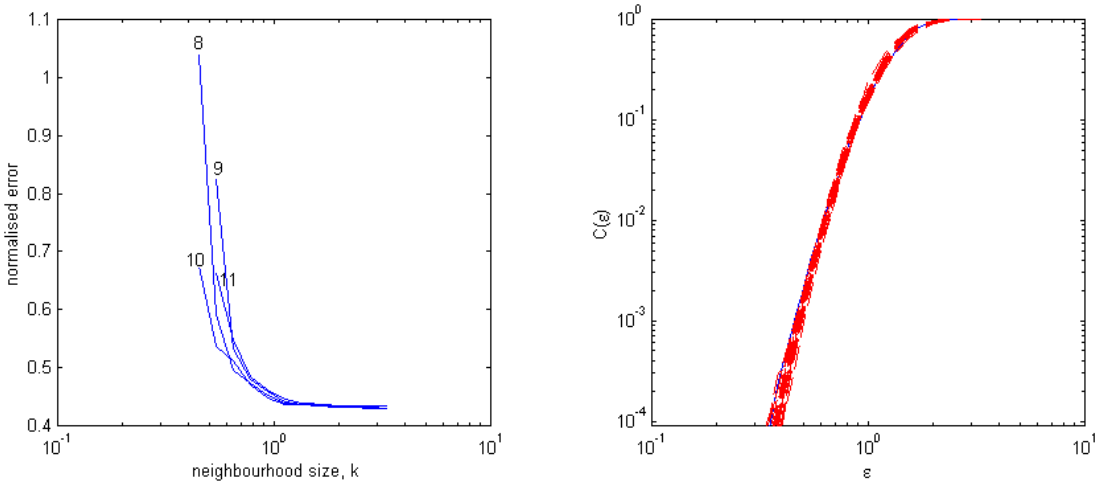

Figure 4.9: Left: LL-GL plot for eye-closed EEG, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 11. Right: correlation sum (solid) at m = 11 and for 19 surrogates (dashed). A time delay of τ = 6 was used for embedding. Theiler window n_m = 15

Our results here are not surprising, and in good agreement with CASDAGLI (1991), who applied the LL-GL technique in a preliminary example on the EEG. These findings are also in agreement with THEILER and RAPP (1996) who investigated the dimension of a large number of similar datasets, concluding that the null hypothesis could only be rejected in 106 of 1320 series tested ($p < 0.025$). However, this is more than the expected 33 so not all of the datasets can be regarded as linear Gaussian noise. Referring to Figure 4.8 (right), the distribution of the surrogates does appear to be biased towards increased slopes implying it is conceivable that the null hypothesis might be rejected for similar datasets. However, the LL-GL is less forgiving, displaying the signature of a linear stochastic process. The datasets for Figures 4.8 and 4.9 display abundant alpha activity and hence are consistent with STAM *et al* (1999) who concluded that 98.75% of epochs studied containing alpha activity could not be distinguished from filtered noise (a nonlinear prediction statistic was used in this case).

4.3.2.2 EEG Data II (Stage 4 Sleep)

For this section, EEG data was recorded during sleep studies. A professionally trained electroencephalographer distinguished the stages. Two datasets marked as stage 4 sleep (abundant delta activity occurring during deep sleep), each ~8 seconds long (N = 4096) were analysed. The data was recorded from the C3 channel, sampled at a frequency of 500 Hz, and filtered from 0.1 to 30 Hz. Figure 4.10 (left) illustrates the LL-GL plot over a range of dimensions.

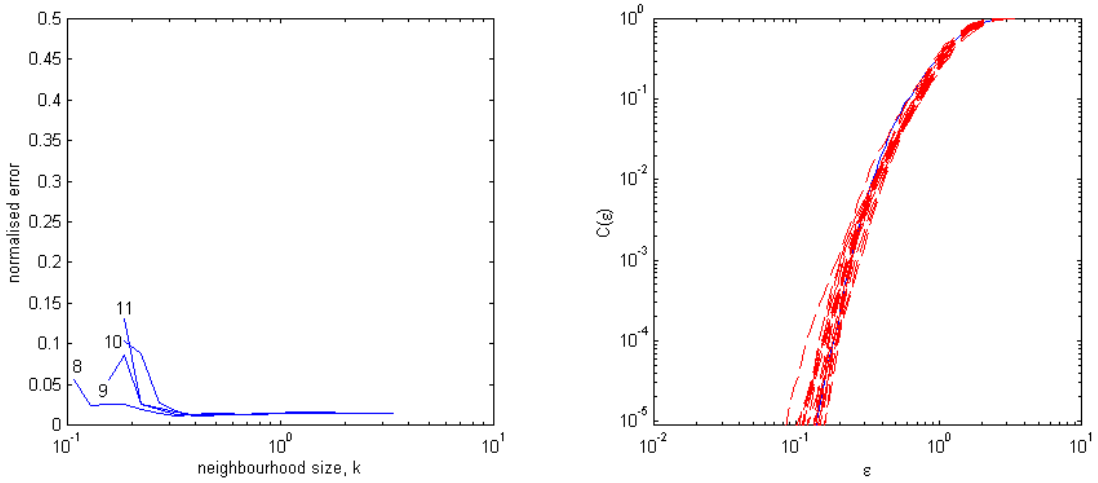

Figure 4.10: Left: LL-GL plot for stage 4 sleep, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 8. Right: correlation sum (solid) at m = 8 and for 19 surrogates (dashed). A time delay of τ = 30 was used for embedding. Theiler window n_m = 200

It is slightly more difficult to draw solid conclusions here - one may argue that a small improvement in prediction accuracy is seen for a smaller neighbourhood size. The improvement is very slight and the data is interpreted here as demonstrating stochastic behaviour as opposed to high dimensional dynamics. Referring to Figure 4.10 (right) the null hypothesis cannot be rejected in this case. Figure 4.11 tells a similar story for a dataset from another subject.

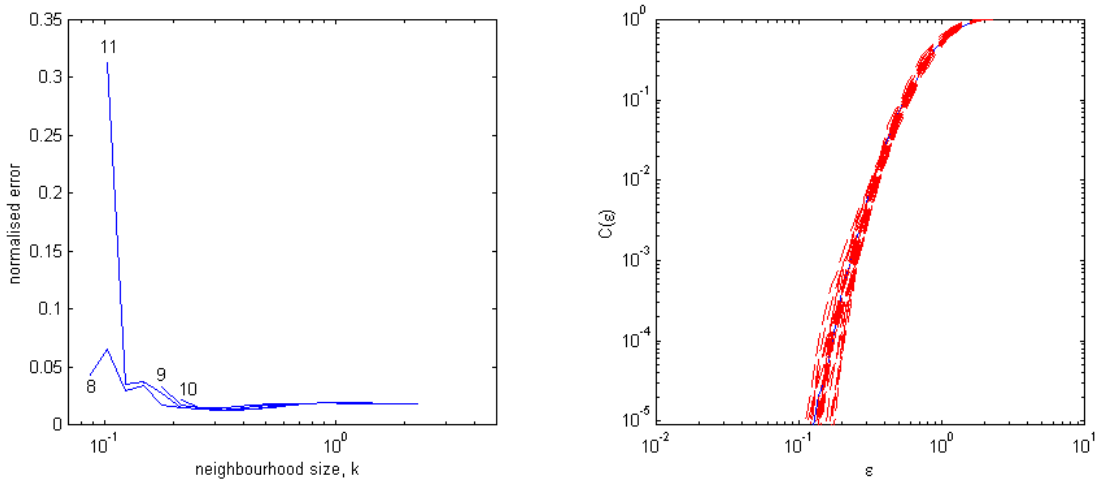

Figure 4.11: Left: LL-GL plot for stage 4 sleep, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 9. Right: correlation sum (solid) at m = 9 and for 19 surrogates (dashed). A time delay of τ = 30 was used for embedding. Theiler window n_m = 200

4.3.2.3 EEG Data III (Petit Mal Absence Seizure)

The petit mal absence seizure can occur mainly in children older than 4 years, with a declining incidence throughout adolescence. Usually, an attack consists of a sudden lapse of consciousness and impairment of mental functions with typical durations between 5 and 20 seconds (NIEDERMEYER and LOPES DA SILVA, 1999). In this section we analyse ~8 seconds of ictal (seizure) activity sampled at 500 Hz (N = 4096 samples) from the frontal regions (F8 and F3 respectively). Data was filtered from 0.16 Hz to 30 Hz. The LL-GL plot in Figure 4.12 (left) exhibits a clear deterministic signature. In particular, the locally linear (globally nonlinear) model exhibits a ~200% improvement in prediction over a globally linear AR model. The correlation sum is illustrated in Figure 4.12 (right) along with 19 surrogates. The original time series exhibits a larger correlation integral than all of the surrogates and we reject the null hypothesis ($p < 0.05$). Figure 4.11 illustrates similar results for the same ictal seizure from channel F3.

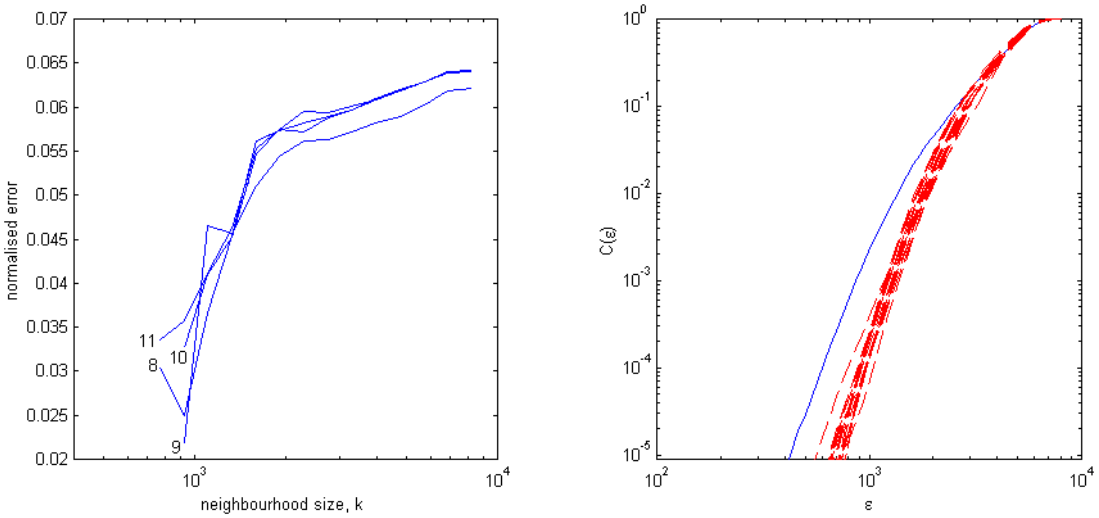

Figure 4.12: Left: LL-GL plot for seizure data, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 9. Right: correlation sum (solid) at m = 9 and for 19 surrogates (dashed). A time delay of τ = 30 was used for embedding. Theiler window n_m = 300

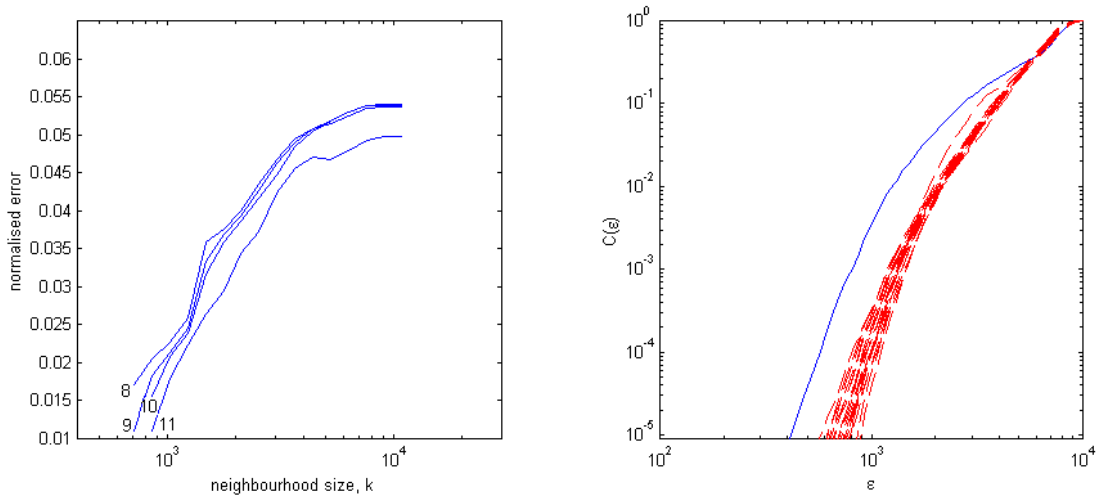

Figure 4.13: Left: LL-GL plot for seizure data, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 9. Right: correlation sum (solid) at m = 10 and for 19 surrogates (dashed). A time delay of τ = 30 was used for embedding. Theiler window n_m = 300

The ictal EEG of the petit mal absence is characterised by general synchronous 3 Hz spike-wave discharge that when observed visually and in reconstructed phase space, suggests deterministic dynamics. Our results for dimension estimation are consistent with PIJN *et al* (1991) who were able to differentiate rat epileptic datasets from their phase-randomised surrogates. Similarly, VAN PUTTEN and STAM (2001) concluded that the interictal EEG activity seen in West syndrome could not be described accurately by a Gaussian linear stochastic process. CASDAGLI *et al* (1997) found statistically significant nonlinearities in invasive EEG recordings from patients with temporal lope epilepsy. Our application of the LL-GL reinforces these findings – in fact the

prediction error characteristics in Figures 4.12 and 4.13 are quite similar to a mathematical model of chaos in Figure 4.7. However, based on the data we may only conclude significant evidence for nonlinearity. Indeed HERNÁNDEZ *et al* (1996) concluded that petit mal data is likely to be a form of stochastic disturbed limit cycle behaviour rather than chaos. In Section 4.6 we show that nonlinearity demonstrated in the nonlinear time series techniques does not necessarily imply *dynamical* nonlinearity.

4.4 Analysis of a Stochastic Limit Cycle Model of the EEG

Armed with the experience of applying nonlinear time series analysis techniques to both mathematically generated data and actual EEG data, we perform the same analysis on the stochastic limit cycle model of the EEG introduced in Chapter III. The data used to produce Figures 3.12 and 3.13 are employed again here (N = 4096 samples). Previously, we compared traditional statistics of the model with EEG from the eyes-closed relaxed state (same data that produced Figures 4.8 and 4.9 in this chapter) including power spectra and amplitude histograms. Here, we are interested in how the nonlinearity of the model manifests itself under the microscope of nonlinear time series analysis. Methods of time series analysis are optimised for *either* stochastic or deterministic data (KANTZ and SCHREIBER, 1997). They do not take into account that noise can actually interact with the dynamics, as is clearly the case for the stochastic limit cycle oscillator model of the EEG. The LL-GL approach goes some way to bridging the gap between deterministic and stochastic analysis and hence explains our motivation for employing it in this chapter. Figure 4.14 (left) illustrates the LL-GL plot over a range of dimensions. The results here are remarkably similar to those obtained from actual EEG data in Figures 4.8 – 4.11, with the data being optimally described by an autoregressive model instead of a nonlinear model. The correlation sum is illustrated in Figure 4.14 (right) for the original series (solid) and 19 surrogates. Similar to the results for actual EEG data from the eyes-closed and deep sleep states, we cannot reject the null hypothesis at the 95% level. Notably, however, the stochastic limit cycle oscillator model in its current form does not account for the nonlinearity evident in pathological EEG data (as illustrated in Figures 4.12 and 4.13). While this suggests more research is required to investigate under what conditions the model exhibits increased nonlinearity in its simulated data, we show in Section 4.6

that it is possible that the introduction of a simple static nonlinearity (as might occur in a static measurement function) can account for the appearance of nonlinear dynamics.

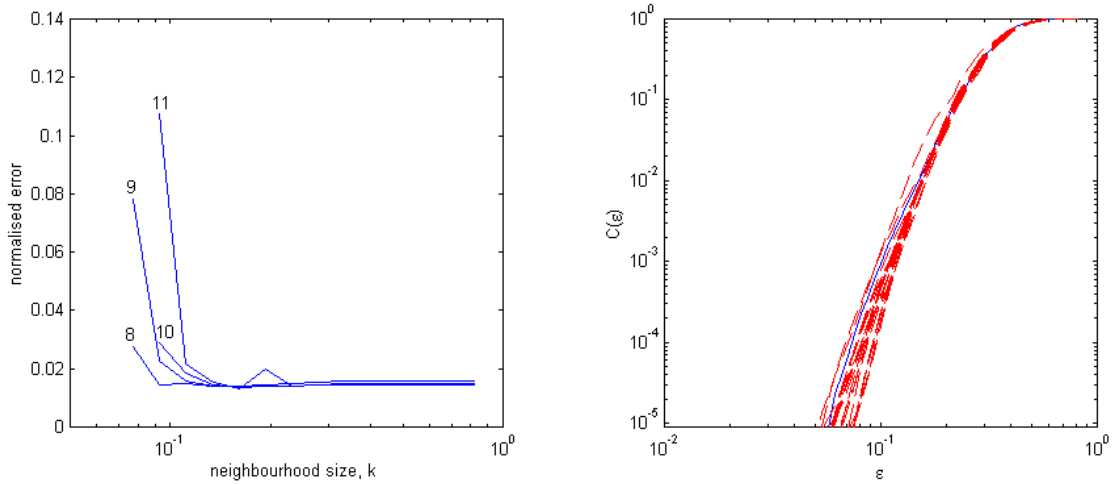

Figure 4.14: Left: LL-GL plot for the stochastic limit cycle model of the EEG, dimensions m = 8, 9, 10, 11. Minimum error is achieved for embedding dimension 10. Right: correlation sum (solid) at m = 10 and for 19 surrogates (dashed). A time delay of τ = 27 was used for embedding (after downsampling by 4). Theiler window n_m = 100

4.5 Implications for Modelling the EEG in Practice

It is very interesting that despite the fact that nonlinear dynamics underpin the stochastic limit cycle oscillator model, no evidence for nonlinearity is revealed by the nonlinear time series techniques (Figure 4.14). This situation is reflected in actual (non-pathological) EEG data also, where the mechanisms producing the EEG data are almost certainly nonlinear. SCHREIBER (1999) makes the point: "…whether the observed irregularity is due to an intrinsic instability, the large number of neurons, or noise, may not be decidable based on time series data alone". Choosing a technique for practical applications therefore, should be justified by the data itself. With the possible exception of pathological EEG data, this chapter suggests that nonlinear time series analysis might *not* appear to provide any benefit over traditional techniques as little evidence of nonlinearity is present in the data itself. Indeed, BILNOWSKA and MALINOWSKI (1991) concluded that the goodness of prediction via the method of SUGHARA and MAY (1990) (a variation of the zeroth order approximation technique – see Section 4.2.4) gave no benefit over the traditional AR approach. The application of the LL-GL technique in this chapter confirms these authors' findings. Based on these conclusions, Chapter VI builds on the AR model for feature extraction for practical applications involving direct brain interfacing using the EEG.

Given its lack of displayed nonlinear characteristics, one might be forgiven for concluding that nonlinear time series techniques are not applicable to the EEG. Furthermore, much of the research in the last decade or two with regard to application of these techniques to the EEG might appear to have been in vain. We have shown that pathological data differs significantly from non-pathological data by the increased presence of nonlinearity thus suggesting applicability of nonlinear time series techniques in practice for this kind of data. Recently, LEHNERTZ and ELGER (1998) interpreted the correlation dimension as an 'operational definition' and showed its ability to predict epileptic seizures several minutes prior to the event. In the final section of this chapter we show that nonlinear time series analysis may be applicable to stochastic data subjected to a static, nonlinear measurement function.

4.6 Stochastic Dynamics and Nonlinear Time Series Analysis

The LL-GL and correlation dimension techniques applied to petit mal absence seizure data suggest strong nonlinearity in the data. In this section, we show that time series techniques may correctly identify nonlinearity in the data but that it might be *non-dynamical* nonlinearity that is being detected. KANTZ (2001) recently demonstrated an AR(2) process with a static nonlinearity was well suited to nonlinear time series prediction schemes. In this section we provide a similar presentation but using an alternative formulation – via stochastic differential equations, which are the basis of our model of the EEG signal. In addition, we briefly review some recent results, which facilitate one to extract stochastic differential equations from time series data. These results could be of particular interest to modelling of the EEG in the future.

Consider the stochastic differential system

$$dX_1 = -aX_1 dt + bdW \tag{4.7a}$$

$$dX_2 = X_1 dt \tag{4.7b}$$

where $a > 0$. The solution to (4.7a) is called the Ornstein-Uhlenbeck process (ØKSENDAL, 1998). We can show that the discrete time approximation to (4.7b) is an AR(2) process. Consider a time discretisation with $0 = \tau_0 < \tau_1 < \ldots < \tau_n < \ldots < \tau_N = T$ and $\Delta t = T/N$. We may write the Euler-Maruyama approximation for (4.7a) as

$$X_{1_{n+1}} = X_{1_n} - aX_{1_n}\Delta t + b\left(W_{\tau_{n+1}} - W_{\tau_n}\right)$$

Similarly, for (4.7b) we write

$$\frac{X_{2_n} - X_{2_{n-1}}}{\Delta t} = X_{1_n}$$

Inserting this into the previous equation and replacing the Wiener process by its white noise equivalent, yields the AR(2) process

$$X_{2_{n+1}} = (2 - a\Delta t)X_{2_n} - (1 - a\Delta t)X_{2_{n-1}} + \xi, \qquad \xi \in N(0, b^2\Delta t^3)$$

Consider a static, nonlinear measurement function

$$Y(t) = \text{sgn}(X_2(t))\sqrt{|X_2(t)|} \tag{4.8}$$

Figure 4.15 illustrates the effect of (4.8) on Euler-Maruyama (see Chapter III) solution of (4.7a) and (4.7b).

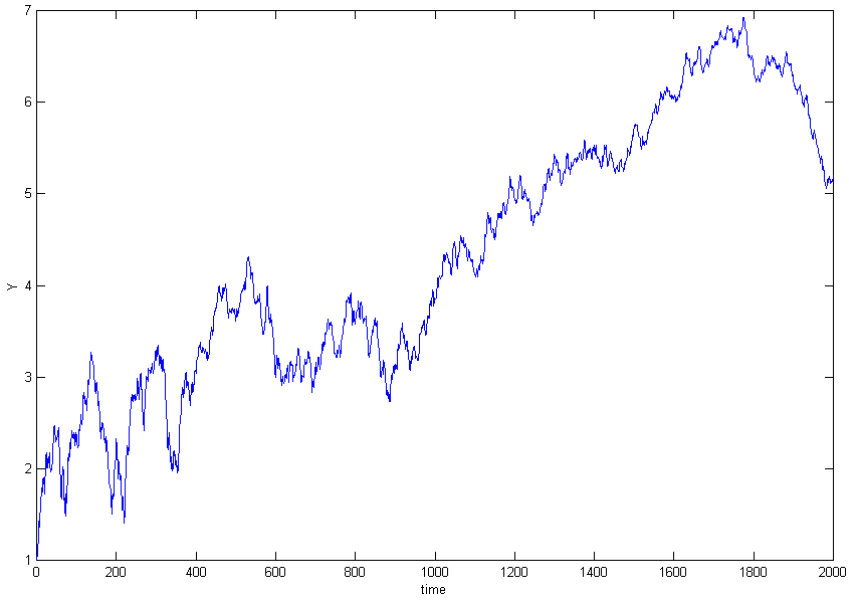

Figure 4.15: Euler-Maruyama solution of (4.7a) and (4.7b) where $a = 2$, $b = 1$ subject to (4.8). Simulation parameters: $\Delta t = 0.01$, $T = 2000$. Scaled to unit sampling

Figure 4.16 illustrates the LL-GL plot for (4.8). The plot reveals that a simple stochastic process (here the integral of the Ornstein-Uhlenbeck process) measured through a static nonlinear function is better modelled via a locally linear (globally nonlinear) scheme than an autoregressive process (actually, if the static nonlinearity was known one could invert it and subsequently apply an AR model). KANTZ (2001) puts forward an interesting explanation for this. Consider the generalisation of a dynamical system to a Markov process where the transition rules are defined by a set of transition probabilities. A Markov process is one where its future transition is dependent only on its current state (or more generally, for a Markov model of order *m*, the transition probability depends on the last *m* states). A deterministic dynamic

system may be considered a limit of a Markov process where the transition probability is δ-peaked. In the stochastic setting, for a dimension m reconstruction, the locally constant prediction strategy $\hat{s}_{n+1} = \langle S_{j+1} \rangle_{\Gamma_n}$ (see Section 4.2.4) may be considered to be the mean of the conditional probability $p(s_{j+1} \mid s_j, s_{j-1}, \ldots, s_{j-m+1})$ integrated over the neighbourhood Γ_n. For certain probability distributions, the prediction error might be minimised by the mean. Adding a static nonlinear measurement function to the stochastic limit cycle oscillator suggests one way of accounting for increased nonlinearity seen in pathological data.

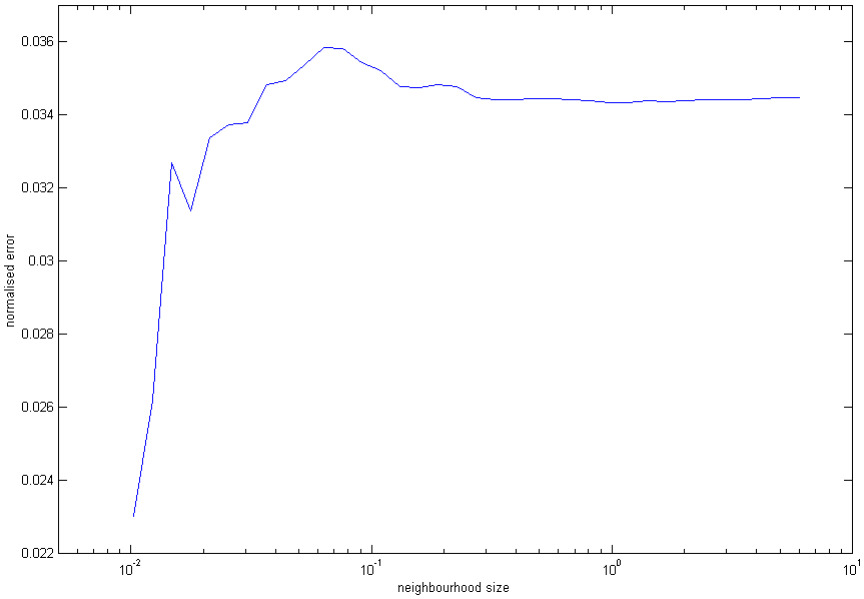

Figure 4.16: LL-GL plot for a stochastic process subject to a static nonlinearity (4.8)

We conclude this chapter by mentioning the recent results of SIEGERT *et al* (1998) and GRADISEK *et al* (2000) for determining the drift and diffusion terms from data with an underlying stochastic differential equation. Consider the (for simplicity 1-dimensional) stochastic differential equation

$$dX = a(t, X)dt + b(t, X)dW(t)$$

The family of solutions are called an *Itô diffusion* (subject to the sensible constraints of Theorem 3.1) and are homogenous (in time) Markov processes. For sufficiently smooth a and b, the transition probabilities of such a Markov processes have a density $p = p(s, x; t, y)$ satisfying the Fokker-Planck equation

$$\frac{\partial p}{\partial t} + \frac{\partial}{\partial y}(ap) + \frac{1}{2}\frac{\partial}{\partial y^2}(b^2 p) = 0$$

The probability density p may be estimated by constructing a histogram from the data itself. Moreover, the drift (a) and diffusion (b) terms may be estimated from their statistical definitions

$$a(t,x) = \lim_{\tau \to 0} \frac{1}{\tau} E\left(X(t+\tau) - X(t) \mid X(t) = x \right)$$

$$b^2(t,x) = \lim_{\tau \to 0} \frac{1}{\tau} E\left(\left(X(t+\tau) - X(t) \right)^2 \mid X(t) = x \right)$$

The conditional expectations may be evaluated by numerically integrating over the transition probability density p. Thus, it is possible to estimate the drift and diffusion contributions and hence the stochastic differential equation underlying the dynamics. This may be seen as the stochastic counterpart to modelling the global deterministic dynamics discussed in Section 4.2.4. The motivation here is the same though, namely to create a concise model of the underlying dynamics for use in practical applications (for example feature extraction in pattern recognition applications).

4.7 Summary

The primary goal of this chapter was to provide further support for the stochastic limit cycle model of the EEG (see Chapter III) by invoking techniques from nonlinear time series analysis. A secondary goal was to compare and contrast *qualitatively different* EEG data from a dynamical perspective. Nonlinear time series prediction, a robust indication of determinism, was employed to compare real and model data. Data generated from the model produces an identical dynamical signature to that of non-pathological EEG data. Furthermore, for both model and non-pathological real data, and using the correlation sum, the null hypothesis that the data was produced from a Gaussian linear stochastic process could not be rejected at the 95% level. In contrast, epochs of pathological data (arising from a petit mal seizure) were shown to be significantly nonlinear and this suggests further work to investigate how the stochastic limit cycle oscillator model might produce increased nonlinearity in its data - a tentative suggestion of a static nonlinear measurement function was put forward as one possibility. Another possibility might be to weakly couple the oscillators or narrow the state transition probability distribution (equivalent to reducing the noise intensity).

In the next chapter, we study more complex dynamical systems involving coupled oscillators for modelling evoked potentials – an important feature of the EEG with application to direct brain interfacing.

CHAPTER V
EVOKED POTENTIALS MODELLING WITH COUPLED OSCILLATORS

5.1 Introduction

In this chapter we explore an important phenomenon that occurs ubiquitously in populations of biological oscillators, that of synchronisation. Mutual synchronisation can occur in populations of cells ranging, for example, from cardiac pacemaker cells (JALIFE, 1984) to circadian pacemaker cells (WINFREE, 1967; WINFREE, 1980) to neurons in the hippocampus (TRAUB *et al*, 1989). Our interest in synchronisation arises as a possible mechanism underlying evoked potentials and event-related potentials in the EEG. Evoked potentials are the electric responses of the nervous system to motor or sensory stimulation. They may be easily elicited by visual, auditory, or somatosensory stimuli (MISULIS, 1994). Event-related potentials are similar to evoked potentials in response magnitude and morphology except they are usually defined as endogenous potentials, resulting from cognitive or initiative processes. Both evoked and event-related potentials have found application in direct brain interfacing – see Chapter VI for more details.

The generators of evoked and event-related potentials are typically regarded as the spatial and temporal summation of Excitatory Postsynaptic Potentials (EPSP) and Inhibitory Postsynaptic Potentials (IPSP), with the possible contribution of action potentials from sub-cortical generators (MISULIS, 1994). However, very little is known about the temporal dynamics responsible for the generation of these potentials. This might be due at least in part by the fact that there is no generally accepted model of the EEG. Typically, models interpret the EP as an impulse response of an under damped oscillatory system, where the oscillatory system is intended to be a model of the ongoing EEG. WILSON and COWAN (1972) suggest parameters for their model of tightly coupled excitatory and inhibitory neurons to admit marginally stable damped response behaviour as a basis for evoked potentials. PROTOPAPAS *et al* (1999) model an evoked potential as an impulse or "shock" applied to a large-scale network of realistic neurons. JANSEN *et al* (1993) describe an evoked potential as a response to a

volley of activity impinging on a lumped model of the cortex by LOPES DA SILVA *et al* (1976) with an additional excitatory feedback loop, and later employing coupled lumped models (JANSEN and RIT, 1995). In the latter, the authors suggest a possible role for synchronicity in addition to damped sinusoidal response behaviour as the primary mechanisms underlying evoked potentials. Our approach here involves explicitly focusing on possible mechanisms underlying the generation of evoked potentials and may be collectively termed under the general heading of *phase reordering*.

That synchronisation might underpin the generation of evoked potentials appears quite plausible. Synchronisation has been suggested as a mechanism to solve the so-called 'binding problem' in the visual cortex – that is the problem of decoding a distributed representation of multiple objects in the visual field (MILNER, 1974; MALSBURG, 1981). Neuronal coding for the same object might be bound together by synchronous firing and differentiated from other neural assemblies representing different objects by the absence of synchrony between them. This concept and the existence of synchrony in other parts of the cortex are supported by recent experimental evidence (ECKHORN, 1994; KÖNIG and ENGEL, 1995; SINGER and GRAY, 1995).

The key concept we wish to explore in this chapter is that *afferent (sensory) stimulation naturally triggers synchronisation*. We shall see that *phase reordering* including both *phase resetting* and *synchronisation* play important roles. We approach the process of modelling synchronisation on two scales. The first we call a microscopic scale, describing interactions between biologically realistic models of neurons. The second, a more abstract approach, is focused at the macroscopic level and deals with neural populations as the basic unit (BURKE and DE PAOR, 2002). We conclude the chapter by introducing the novel concept of stochastic synchronisation.

5.1.1 A Survey of Models of Neural Synchronisation

It is useful to review some of the methods used in the literature to model neural synchronisation to help put the work contained in this chapter into context. Study of neural synchronisation invariably leads to a description of coupled oscillators. Coupled oscillators are difficult to analyse, often leading to large numbers of intractable equations. This area of applied mathematics is still quite new and there are no standard

techniques available. A major simplification can be obtained using *phase models* (ERMENTROUT, 1981; ERMENTROUT and KOPELL, 1984; ERMENTROUT and KOPELL, 1991), which reduces differential systems with asymptotically stable limit cycle behaviour to a one-dimensional system parameterised only by phase. Small perturbations can also be studied in the resultant phase equations by averaging (GUCKENHEIMER and HOLMES, 1983), as can stability of different types of coupling. A convenient representation of coupling terms can also be obtained in terms of phase response functions (KOPELL and ERMENTROUT, 1986), although these usually have to be determined numerically (see Section 5.2.1 for more details).

To produce more realistic models of neural synchronisation, MIROLLO and STROGATZ (1990) have used the so-called integrate-and-fire model for neural spiking. These authors showed that globally excitatory interaction results in synchronization of zero phase lag. However, it is not true that excitatory coupling automatically leads to synchronisation. HANSEL *et al* (1995) differentiate 'Type I' and 'Type II' phase response functions. The former is positive over the whole cycle whereas Type II response curves are negative at the beginning and positive at the end of the cycle. This suggests that excitatory coupling for Type I phase response functions cannot lead to stable synchronisation. On the other hand, fast enough excitation for Type II phase response curves can lead to synchronisation, as seen in MIROLLO and STROGATZ (1990). Interestingly, more comprehensive models of neurons such as the Hodgkin-Huxley neuron (although not a general rule – see HANSEL *et al* (1995) for examples) possess a Type II phase response curve thus suggesting that excitatory coupling might result in synchronisation (we will make use of this in Section 5.2). ERMENTROUT (1996) made a connection between the two types of phase response function and the transition to oscillations. For Type I functions, oscillations can be sustained at arbitrarily low frequencies whereas for Type II phase response functions, oscillations are usually triggered through a Hopf bifurcation (GUCKENHEIMER and HOLMES, 1983).

As one models larger networks, more complex behaviour can be seen. STURM and KÖNIG (2001) describe how both static noise (such as inhomogeneous distribution of parameters) and stochastic noise (small random fluctuations of the dynamic variables) can lead to the emergence of irregular, partially synchronised, and alternating states of synchrony in large networks. Higher levels of noise encourage clustering of synchrony

into smaller and smaller clusters, transcending into complete asynchrony. We will show in Section 5.5 that another phenomenon is also possible in coupled oscillators, that of synchronisation of irregular, stochastic behaviour between neighbouring oscillators.

5.2 Microscopic Synchronisation: Networks of Neurons

Computational neuroscience is concerned with understanding the information content of neural signals by modelling the nervous system at many different structural scales, including the biophysical, the circuit, and the systems levels (KOCH and SEGEV, 1998). In this section, we use techniques from computational neuroscience, by employing biological realistic compartment models of neurons in a small neural circuit to explore synchronisation. Before introducing the techniques and the model itself, it is useful to build intuition by first considering phase models. The latter can help illuminate some common phenomena observable in more complex systems of coupled oscillators in a mathematical tractable way.

5.2.1 Phase Models

Consider a biological oscillator modelled as a structurally stable dynamical system exhibiting an asymptotically stable limit cycle where the system is described by

$$\frac{dx}{dt} = F(x) \tag{5.1}$$

and $x \in R^n$. Assuming the oscillator is on its steady state limit cycle at all times, we may make a dramatic simplification by describing the oscillator by a single state variable called the phase $\theta(t)$ .We rescale θ and use modular arithmetic so that it takes on values from 0 to 2π radians and parameterise θ so that it flows *uniformly* around the limit cycle. In essence, we have replaced the limit cycle by a circular limit cycle that flows at a constant speed. Strictly speaking, the new limit cycle is topologically equivalent to the original but not conjugate (see Chapter II). This yields a simple equation for (5.1)

$$\frac{d\theta}{dt} = \omega \tag{5.2}$$

where ω is the natural frequency of the oscillator and the solution of the equation is given by

$$\theta(t) = \omega t + \theta(0) \pmod{2\pi}$$

The solution is constrained to always lie on the interval $0 \le \theta < 2\pi$. This simplification of using phase equations discards information pertaining to the original equations in exchange for making analysis of coupled oscillators mathematically tractable. For example, studying four coupled Hodgkin-Huxley model neurons would require sixteen nonlinear differential equations for the neuron somas alone if we were to use the original equations. Indeed, there are many situations where phase models are quite suited, for example when no model for the biological oscillator in question exists or when we wish to model an oscillator that generates an event at a certain phase.

Consider N coupled biological oscillators given by their phase equation (5.2). The resultant coupled system exists on an N-dimensional torus T^N. We may now write

$$\frac{d\theta_i}{dt} = \omega_i + \alpha h_i(\theta_1, \theta_2, ..., \theta_N) \tag{5.3}$$

where h_i represents the coupling effect of all other oscillators on oscillator i and α is the strength of coupling between oscillators. The coupling function h_i, which depends only on relative phases, is chosen such that it is 2π-periodic in each of its arguments. This is necessary so that the flow is uniquely defined at each point on the torus T^N. We may further describe the coupling function as being the sum of contributions between pairs of oscillators, i.e.

$$h_i(\theta_1, \theta_2, ..., \theta_N) = \sum_{j=1}^{N} h_{ij}(\theta_i, \theta_j) \tag{5.4}$$

where j represents the presynaptic oscillator and thus h_{ij} describes the effect of oscillator j on oscillator i.

When $|\omega_i - \omega_j| = O(\alpha)$ it can be shown (ERMENTROUT and KOPELL, 1984) via averaging theory that a coordinate change allows one to write

$$\alpha h_{ij}(\theta_i, \theta_j) = \alpha H_{ij}(\theta_j - \theta_i) + O(\alpha^2)$$

and thus for weak coupling $|\alpha| \ll 1$ we consider the functions in (5.4) to depend only on phase difference. This is known as *diffusive* or *electrotonic* coupling. For coupling far from the weak condition, results of this theory can still hold (KOPELL, 1988). Note that if $H_{ij}(0) = 0$ then in-phase synchronisation is permitted. ERMENTROUT and KOPELL (1991) have shown that for strong coupling the results of the averaging technique, namely employing diffusive coupling, can still be applicable under certain conditions.

In particular, employing diffusive coupling was shown to be appropriate if the interaction effects (e.g. impulses imparted) are distributed around the cycle of the oscillator. An example where this might occur of relevance to this thesis can be given by considering a neural population coupled to a neighbouring neural population. Each population is composed of subunit oscillators (neurons). We assume strong coupling between subunit neighbours such that they are synchronised but not in-phase (this can be accomplished by assuming additional inhibitory coupling). Thus the pulse-coupled interaction of a neural population on its neighbour is distributed about the oscillator's cycle. In addition, if one of the subunits is phase-shifted by a stimulus, it is assumed that the other subunits follow almost immediately. We use this result for selecting the form of the coupling terms in Section 5.3.

An alternative type of coupling is called *synaptic* or *chemical.* Consider an oscillator subject to periodic forcing. If θ_m is the phase just before the m^{th} stimulus, we may write

$$\theta_{m+1} = \theta_m + \omega T + R(\theta_m)$$

where ω is the natural frequency of the oscillator and T is the period of the forcing. R is called the *phase response curve* and gives the phase shift of the oscillator in response to a stimulus given at a particular phase (WINFREE, 1980). The phase response curve is obtained numerically by perturbing the oscillator and measuring the phase shift after the system relaxes back to its limit cycle. We may write the previous difference equation as a differential equation

$$\frac{d\theta}{dt} = \omega + \delta(t \bmod T) R(\theta)$$

ERMENTROUT and KOPELL (1990) suggest it is reasonable (since real stimuli are not instantaneous) to replace the Dirac delta function with a distributed function *P(t)* that is sufficiently narrow to mimic the effects of coupling two oscillators by their phase response curves. Thus, the synaptic coupling term in (5.3) may be written as $h_{ij}(\theta_i, \theta_j) = P(\theta_j) R(\theta_i)$.

5.2.1.1 Analysis of a Chain of Coupled Oscillators

We present a simple example of a chain of coupled oscillators by RAND *et al* (1988). Here we employ diffusive coupling $h_{ij}(\theta_i, \theta_j) = H_{ij}(\theta_i - \theta_j)$. Since H is a periodic function, we may expand it in its Fourier series, and keeping only the odd component

of the first term in the expansion, yields $H_{ij}(\theta_i,\theta_j) = a_{ij}\sin(\theta_j - \theta_i)$. Consider a chain of four oscillators with nearest neighbour coupling given by

$$\begin{aligned}
\frac{d\theta_1}{dt} &= \omega_1 + a\sin(\theta_2 - \theta_1) \\
\frac{d\theta_2}{dt} &= \omega_2 + a\sin(\theta_3 - \theta_2) + a\sin(\theta_1 - \theta_2) \\
\frac{d\theta_3}{dt} &= \omega_3 + a\sin(\theta_4 - \theta_3) + a\sin(\theta_2 - \theta_3) \\
\frac{d\theta_4}{dt} &= \omega_4 + a\sin(\theta_3 - \theta_4)
\end{aligned} \tag{5.5}$$

Now, by introducing the change of variables

$\varphi_i = \theta_i - \theta_{i+1}$,

we may write the system in matrix form as

$$\frac{d\varphi}{dt} = \Omega + AS$$

where

$$\varphi = \begin{bmatrix} \varphi_1 \\ \varphi_2 \\ \varphi_3 \end{bmatrix}, S = \begin{bmatrix} \sin\varphi_1 \\ \sin\varphi_2 \\ \sin\varphi_3 \end{bmatrix}, \Omega = \begin{bmatrix} \omega_1 - \omega_2 \\ \omega_2 - \omega_3 \\ \omega_3 - \omega_4 \end{bmatrix}$$

and

$$A = a\begin{bmatrix} -2 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -2 \end{bmatrix}$$

For phase locked behaviour we let $\frac{d\varphi}{dt} = 0$ and thus

$$S = -A^{-1}\Omega \tag{5.6}$$

No solution exists for (5.6) if any of the components of $A^{-1}\Omega$ have an absolute value greater than 1. Now, assuming a constant frequency gradient across the chain, that is

$\omega_i - \omega_{i+1} = c$,

we may write

$$\Omega = c\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

and hence (5.6) becomes

$$\begin{bmatrix} \sin\varphi_1 \\ \sin\varphi_2 \\ \sin\varphi_3 \end{bmatrix} = \frac{c}{2a}\begin{bmatrix} 3 \\ 4 \\ 3 \end{bmatrix}$$

The hardest condition to meet (for a solution to exist) is for φ_2 and results in the following condition for 1:1 phase-locked motion

$$\left|\frac{c}{a}\right| \leq \frac{1}{2} \tag{5.7}$$

Equation (5.7) tells us that for phase-locked motion to occur, there is a tradeoff between coupling strength (a) and frequency difference (c). This is quite an intuitive result, as oscillators closer in frequency require 'less effort' or coupling to bring them into synchronisation. An interesting aspect of diffusive coupling is that synchronisation in a chain invariably occurs with a phase lag between neighbour oscillators. Figure 5.1 illustrates the phase lags for different ratios of c/a .

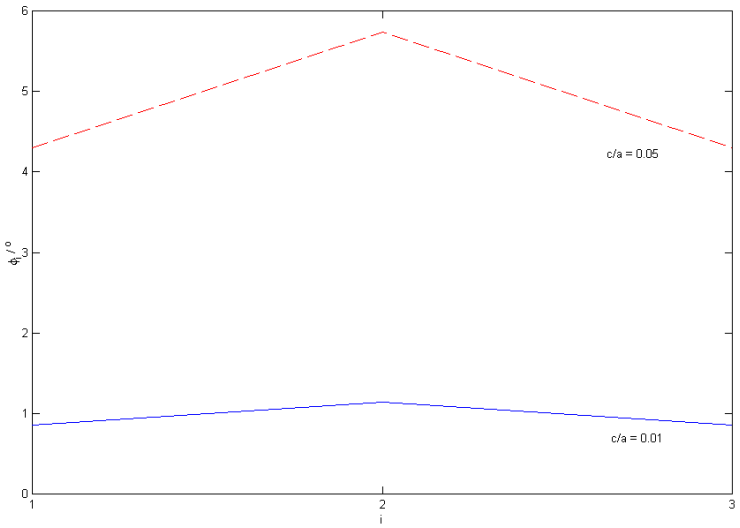

Figure 5.1: Plot of φ_i against i for different values of c/a

Although the analysis here is very much oversimplified, this behaviour is characteristic of more complex systems of coupled oscillators. In order to maintain synchrony, coupling between oscillators must be sufficiently strong; weak coupling can only synchronise oscillators that are very close in natural frequency (ERMENTROUT and KOPELL, 1991). Indeed, this observation is the basis of the model for synchronisation in Section 5.2.3.

5.2.2 Biologically Realistic Neurons

Compartmental modelling is a powerful technique for constructing biologically realistic descriptions of neurons (KOCH and SEGEV, 1998; BOWER and BEEMAN, 1998). The central concept behind compartmental modelling is that small sections of the neuron can be treated as isopotential elements so that the continuous structure of a

neuron may be approximated by a collection of connected compartments. In particular, nonuniformity in physical properties such as diameter, resistances, and capacitance occur between compartments rather than within them thus facilitating different types of compartments (e.g. representing a soma, axon, or dendrite) to be coupled. Figure 5.2 illustrates a pyramidal neuron with dendrites, soma, and axon, with a possible compartmentalisation. The dendrites form a tree-like structure and receive synaptic inputs from other neurons. Synaptically activated ion channels in the dendrites create postsynaptic potentials that we assume to be passively propagated to the pyramid shaped soma. The axon hillock is located near the base of the soma and contains a high concentration of voltage-activated channels that can create action potentials. The action potentials propagate along the axon.

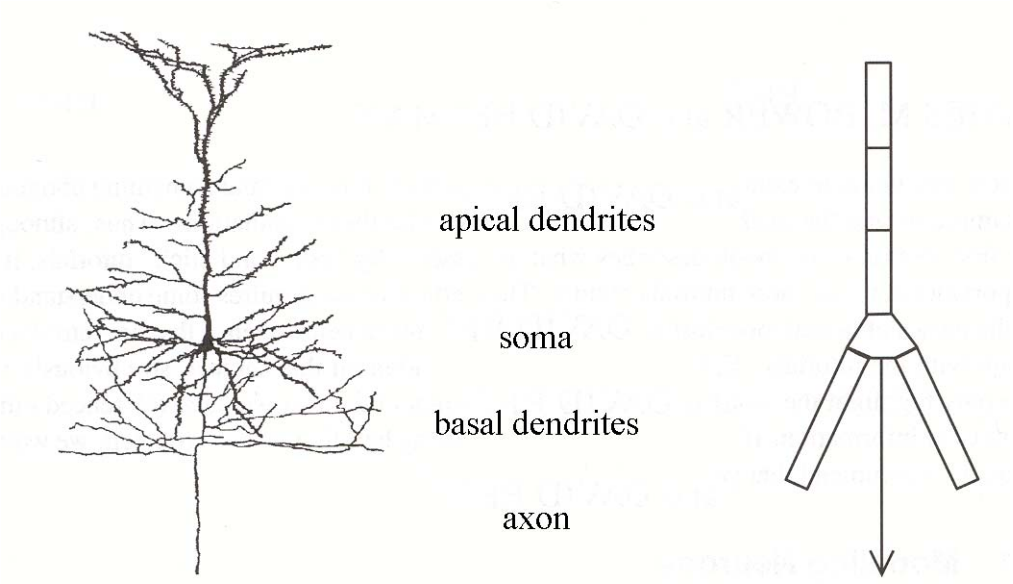

Figure 5.2: Illustration of a pyramidal neuron (left) with a suggested compartmentalisation (right) (BOWER and BEEMAN, 1998)

Each compartment in Figure 5.2 is modelled with differential equations corresponding to an equivalent electrical circuit model; the entire neuron is modelled as a coupled system of differential equations, which itself may be coupled to other neurons. The number of compartments employed depends on what is being investigated. Studying a single neuron might suggest using a large number of compartments for detailed analysis, however modelling a large network might require (for computational efficiency) a much smaller number of compartments per neuron. There are also powerful techniques for reducing the number of compartments, for example replacing dendritic trees with equivalent cylinders is valid under certain conditions (the main one being that the sum of the daughter branches diameters raised to the power of 3/2

must equal the parent branch diameter raised to the power of 3/2 (KOCH and SEGEV, 1998)). Finally, we should point out that much of the effort of practical computational neuroscience modelling is concerned with extracting sensible physical values for the compartments. Typically, intracellular injection of tracing substances and visualisation techniques are used to isolate and dimension individual neurons. Estimating electrical parameters is known as the 'inverse estimation' problem; the basic criterion used is to ensure that the model neuron exhibits the same steady-state input resistance as the actual cell and similar behaviour during transient perturbations (KOCH and SEGEV, 1998).

5.2.2.1 Membrane Models

Figure 5.3 illustrates the equivalent circuit model for a patch of membrane.

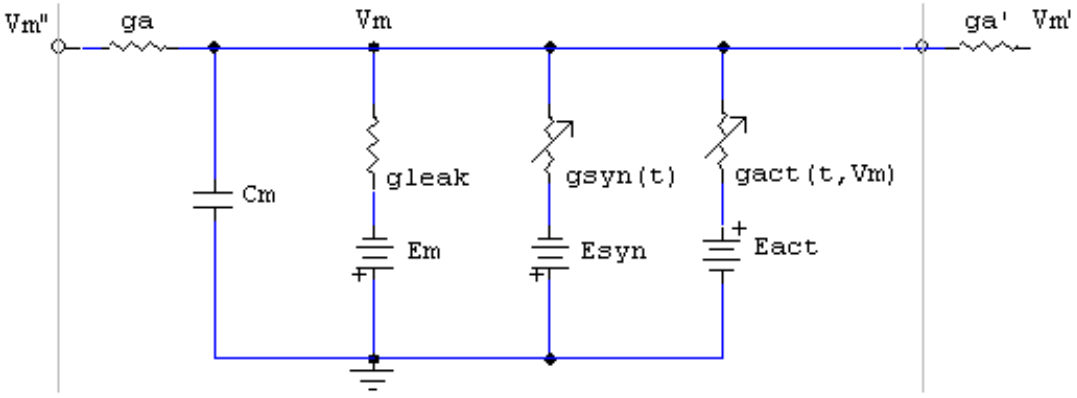

Figure 5.3: Equivalent electrical circuit for an isopotential patch of membrane

There are usually three types of conductance branch considered when modelling patches of membrane (although not all necessarily present): *passive*, *synaptic*, and *active*. The passive branch is the simplest. A number of passive ionic channels are lumped together to form a leakage current (primarily made up of a chloride current). This is modelled by a small reversal potential E_m and the conductance g_{leak} (the reversal potential is that value which results in no current flowing for the ion species in question and corresponds to the Nernst potential). In practice, we also lump the resting potential (usually between –40 and –100 mV) in the E_m term also. Ohm's law describes the current through the passive branch

$$I_{leak} = g_{leak}(V_m - E_m)$$

The membrane acts as a capacitor and hence is represented by C_m. Flow into and out of the compartment occurs passively through the axial conductances g_a and g_a'.

Synaptic channels are typically activated by external chemical agents and are concentrated on the dendrites and soma. We will describe the time-dependent form of g_{syn} in the next section. Its contribution to the current is given by

$$I_{syn} = g_{syn}(t)(V_m - E_{syn})$$

where E_{syn} is the reversal potential for the ion species involved. Voltage dependent channels are represented by a voltage- and time-dependent conductance $g_{act}(t, V)$ and a reversal potential E_{act}. The contribution to the total current is given by

$$I_{act} = g_{act}(t, V_m)(V_m - E_{act})$$

We have already come across these kinds of channels in the Hodgkin-Huxley model described in Chapter I and will use them here in what follows. Finally, a single differential equation for a generalised compartment may be written as

$$C_m \frac{dV_m}{dt} = (E_m - V_m) g_{leak} + (E_{syn} - V_m) g_{syn}(t) + \sum_k (E_{act_k} - V_m) g_{act_k}(t, V_m) + (V_m' - V_m) g_a' + (V_m'' - V_m) g_a + I_{inject}$$

where I_{inject} corresponds to an externally applied current.

5.2.2.2 Synapses and Axons

Synaptic channels change their conductance when the appropriate chemical stimulus (neurotransmitter) binds to the receptor associated with these channels. Rather than deal with the complexity of the reaction kinetics, this is usually simply modelled by a voltage-independent but time-dependent function known as the alpha function (KOCH and SEGEV, 1998). The functional form is given by

$$g_{syn}(t) = g_{\max} \frac{t}{t_p} e^{(1 - t/t_p)} \tag{5.8}$$

The function increases rapidly to a maximum of g_{max} at $t = t_p$ and subsequently decreases more slowly to zero. A more general function involving two exponentials is also possible. In practice, a postsynaptic potential is computed by convolving a train of impinging pulses (action potentials) with (5.8). While an axon may be modelled as a series of compartments, a common approach just involves modelling the axon as a simple delay line for the propagation of action potentials (BOWER and BEEMAN, 1998).

5.2.3 Modelling Synchronisation in a Small Network of Neurons

Our goal here is to demonstrate how afferent (sensory) stimulation naturally triggers synchronisation in a small network of biologically realistic neurons with excitatory

coupling. The phenomenon shown here, elucidated with the help of a microscopic model of synchronisation, might be considered as the basis for evoked potentials which occur as a result of sensory stimulation on a much larger scale in the brain. We use the GENESIS simulation package for modelling in this section (BOWER and BEEMAN, 1998). The scripts that define the model can be found in Appendix A4 and have been tested on GENESIS 2.2.

5.2.3.1 Model Description

The network consists of a circular chain of reciprocally, nearest neighbour connected excitatory neurons. We justify the use of excitatory coupling only on the basis of simplicity and the fact that the majority of the cortex (70-80%) consists of pyramidal cells, which themselves are excitatory (DEFELIPE and FARINAS, 1992). Figure 5.4 illustrates the connections between oscillators employed in the model network.

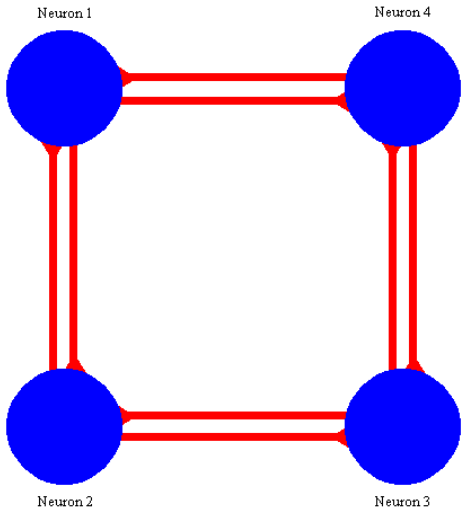

Figure 5.4: Connection topology for a small circuit of neurons.
Flared endings on connections represent excitatory synapses

Each individual model neuron is composed of two compartments: a dendrite and a soma. The dendrite compartment consists of passive membrane with excitatory synapses employing the sodium reversal potential. The soma compartment consists of passive membrane with Hodgkin-Huxley type potassium and sodium ionic channels (see Chapter I). The axon is modelled as a simple delay line. We employ the same specific resistance and capacitances used in HODGKIN and HUXLEY (1952). Specific units enable one to specify parameters that are independent of cell dimensions. The *specific membrane resistance* R_M has units $\Omega.m^2$. The *specific membrane capacitance*

C_M has units F/m^2. The *specific axial resistance* R_A of a cylindrical compartment is proportional to its length and inversely proportional to its cross-sectional area and has units $\Omega.m$. For a compartment of length l and diameter d we have

$$R_m = \frac{R_M}{\pi l d}, \quad C_m = \pi l d C_M, \quad R_a = \frac{4 l R_A}{\pi d^2}$$

Table 5.1 summarises the simulation parameters.

TABLE 5.1: SIMULATION PARAMETERS

Parameter	***Description***	***Value***	***Units***
R_M	Specific membrane resistance	0.3333	$\Omega.m^2$
C_M	Specific membrane capacitance	0.01	F/m^2
R_A	Specific axial resistance	0.3	$\Omega.m$
E_M	Resting potential + leakage reversal potential	-0.07 + 0.0106	V
E_{Na}	Sodium reversal potential	0.045	V
E_K	Potassium reversal potential	-0.082	V
l_{soma}	Length of soma (cylinder)	30×10^{-6}	m
d_{soma}	Diameter of soma (cylinder)	30×10^{-6}	m
$l_{dendrite}$	Length of dendrite (cylinder)	100×10^{-6}	m
$d_{dendrite}$	Diameter of dendrite (cylinder)	2×10^{-6}	m
E_{syn}	Reversal potential for synaptic ionic species	0.045	V
g_{max}	Maximum conductance of synapses	5×10^{-10}	siemens
t_p	Time to peak for postsynaptic potential	0.003	s
α	Synaptic weight	0.03	-
Δ_{axon}	Axon propagation delay	0.02	s
$I_{inject1}$	Current injection to neuron 1	0.000222×10^{-6}	V
$I_{inject2}$	Current injection to neuron 2	0.00023×10^{-6}	V
$I_{inject3}$	Current injection to neuron 3	0.00024×10^{-6}	V
$I_{inject4}$	Current injection to neuron 4	0.00025×10^{-6}	V

Before explaining the basic principle underlying synchronisation in the model, we need to study how the period of an oscillating neuron varies with applied current. Figure 5.5 illustrates a plot of the reciprocal of the period versus applied somatic current for a single, uncoupled oscillator with parameter values described in Table 5.1.

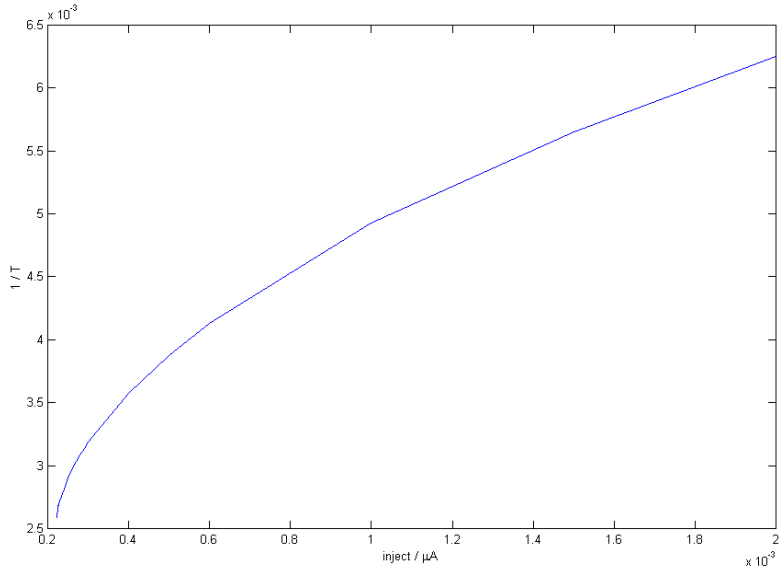

Figure 5.5: Plot of frequency of oscillation versus injected current

As mentioned earlier, ERMENTROUT (1996) showed that Type II phase responses occur in models that obtain rhythmicity via a Hopf bifurcation (GUCKENHEIMER and HOLMES, 1983). The Hodgkin-Huxley is one such model. At a certain critical value for injected current, the asymptotically stable equilibrium resting potential becomes unstable and is surrounded by limit cycle oscillation as illustrated by Figure 5.5.

As the injected current is increased two effects are noticeable. The first is an increase in frequency and the second is a *saturation* of the frequency. Since most neurons most of the time are just below threshold (FREEMAN, 1979), it is the lower part of Figure 5.5 that is of interest. Indeed, (as already touched on in Chapter III) HOPPENSTEADT and IZHIKEVICH (1997) only consider neurons whose dynamics are at a bifurcation point to participate nontrivially in the brain processes. We select the injected current values for the four oscillators (Table 5.1) so they are incommensurate and take on values on the near vertical leg of the frequency-current curve. The effect of a sudden, afferent increase in stimulation impinging on all the neurons increases the frequency of oscillation but also brings the frequencies *closer* together. Recall from the discussion in Section 5.2.1.1 that there is generally a tradeoff between coupling strength and frequency gradient. By first approximation, we can regard the decrease in frequency gradient occurring as a result of afferent stimulation as equivalent to an increase in coupling strength thus encouraging synchronisation.

5.2.3.2 Simulation Results

To demonstrate the effect of afferent stimulation, we apply a step increase in injected current of 0.001 μA to each neuron. Figure 5.6 illustrates the membrane potential of the four neurons. The step occurs at $t = 0.25$. Figure 5.7 illustrates an alternative visualisation using a raster plot where each dot corresponds to a firing event. The vertical lines correspond to the firing event of the first oscillator. Before the step increase in current, it is clear that the neurons are oscillating independently and drifting in phase relative to each other. After the step increase in current, the oscillators move closely into phase, increase in frequency, and maintain synchronisation for the duration of the applied current.

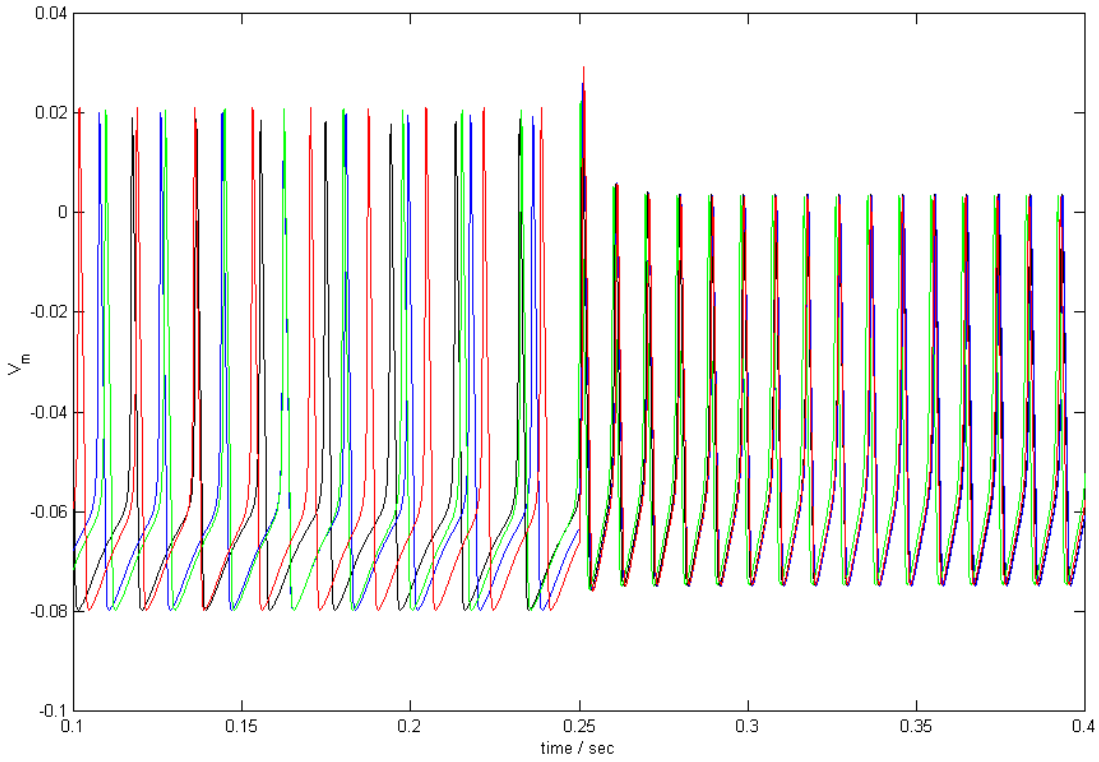

Figure 5.6: Plot of the membrane potential for four neurons. A step increase in injected current occurs at $t = 0.25$

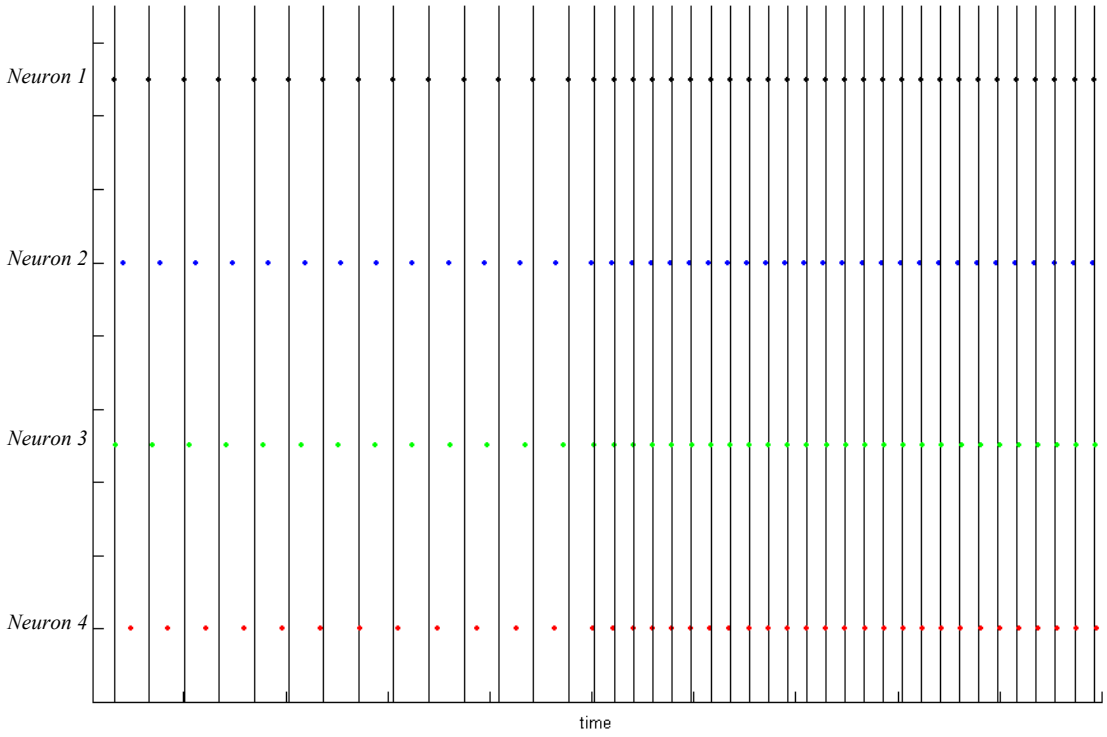

Figure 5.7: Raster plot of firing events for four neurons. A step increase in injected current results in synchronisation and increase frequency in the latter part of the plot. Vertical lines refer to timing of Neuron 1 to aid comparison

5.2.3.3 Discussion

We have demonstrated a biologically plausible mechanism for synchronisation as a result of afferent stimulation. The basic principle is that frequency saturation in response to increased synaptic activity encourages synchronisation by bringing the oscillator frequencies closer together. Different frequency of oscillation in neighbouring neurons is to be expected as a result of the natural variation in synapse density, cell dimensions, and electrophysiological properties. Increased synaptic activity (assumed to occur in response to stimulation) was simulated as a step increase in current to the soma and is equivalent to a step increase of presynaptic activity.

The demonstration of this stimulus-triggered mechanism for synchronisation appears to be new in the literature, yet intuitively simple. The phenomenon is suggested primarily as a possible mechanism underlying evoked potentials whereby a transient synchronisation of neural activity (in response to a visual, auditory, or somatosensory stimulus) produces an electrical response detectable in the background EEG. A closer look at Figures 5.6 and 5.7 reveals phase reordering comprising of two possible mechanisms. Firstly, the decrease in frequency gradient that occurs naturally serves to increase the strength of interaction between neighbouring oscillators and secondly, the step increase works to phase reset the oscillators and bring them into phase more quickly (we discuss phase resetting in more detail later in this chapter). The model and analysis here suggests a simple reason why synchronisation occurs ubiquitously in populations of neurons.

5.3 Macroscopic Synchronisation: Neural Populations

A key step in neuroscience modelling is the choice of scale. Since cortical columns of neurons share activity and are often regarded as the primary unit for behavioural processing (FREEMAN, 2000a), a more appropriate oscillator unit for modelling evoked potentials might be the neural population. To that end, we construct a new model of evoked potentials based on the transient synchronisation of coupled oscillators representing neural populations. We present the fundamental mechanism of the model next, namely that synchronisation of neural populations might occur due to dynamic changes in the coupling strength resulting from sensory stimulation with the coupling level also being influenced by cortical arousal level. Lyapunov stability theory, as with

previous chapters, provides a framework for constructing the model. The model is explored via numerical experiments.

5.3.1 A Mechanism for Synchronisation

It is well established that pools of neurons in local neighbourhoods tend to share activity thus leading to the concept of neural populations (see Chapter I for a review). Freeman's work (FREEMAN, 2000a; Chapter I) on the dynamical interactions of neural populations indicates the roles of two main state variables namely the dendritic current density over the population (w) and axonal pulse density over the population (p). The w-p conversions occur at the trigger zones of the neurons and the p-w conversions occur at the synapses within the population; the product of dw/dp and dp/dw dictates the forward gain of a population. The slope dp/dw is dependent on 1) sensory excitation and 2) arousal level, and is the main determinant in the forward gain of the population (FREEMAN, 1979). Most neurons most of the time are just below threshold and thus an increase in afferent stimulation from the resting level p_0 results in an effective increase in dp/dw - see Figure 5.8. In addition to being input-dependent, the dp/dw characteristic is also state-dependent, being influenced by arousal level. Notably, deep anaesthesia has the effect of diminishing the slope to zero (FREEMAN, 1979).

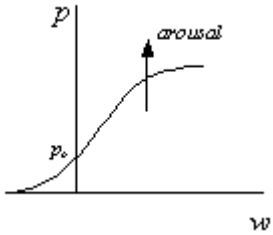

Figure 5.8: Wave-pulse (w-p) relationship. Note that an increase in w from afferent stimulation or an increase in arousal results in an effective increase in slope dp/dw

The basic unit of the proposed model is the neural population and is represented by a limit cycle oscillator. This approach has also been previously employed in WILSON and COWAN (1972) as an interpretation of oscillating populations and also in DEWAN (1964) to represent the periodic component of the EEG. The forward gain of the i^{th} population to the j^{th} population is given by the product of dw_j/dp_i and dp_j/dw_j and we interpret this as the coupling strength between oscillators representing neural populations. An evoked potential might thus be considered as being generated from a transient synchronization of the coupled oscillators resulting from an increase in coupling strength as a consequence of a brief afferent stimulus.

5.3.2 Modelling Synchronisation

We use the prescribed oscillator Method I (2.7) to create a simple limit cycle oscillator, reproduced here for convenience

$$\dot{x}_1 = \frac{\partial H}{\partial x_2} - \lambda H \frac{\partial H}{\partial x_1}$$

$$\dot{x}_2 = -\frac{\partial H}{\partial x_1} - \lambda H \frac{\partial H}{\partial x_2}, \quad \lambda > 0$$

For simplicity, a circular contour given by

$$H(x_1, x_2) = x_1^2 + x_2^2 - r^2 \tag{5.9}$$

is employed. The resultant oscillator produces an asymptotically stable circular limit cycle that is traversed with uniform speed. The oscillator may be regarded as topologically equivalent to an equivalent oscillator representing the aggregated field activity of a neural population. In other words, there exists a mapping *h* from the original state space to the system given by (2.7) and (5.9). To couple the oscillators, we introduce a term proportional to the phase difference between neighbouring oscillators. When the oscillators are 1:1 phase synchronized, this term reduces to zero. This is diffusive coupling and is justified on the basis of simplicity and the results of ERMENTROUT and KOPELL (1991) - the interaction effects for the equivalent neural population oscillator on its neighbours are distributed around the cycle of oscillator, and hence can be represented by diffusive coupling (see the discussion in Section 5.2.1 for more details).

To show that synchronisation can occur, we consider a pair of oscillators (x_1,x_2) and (x_3,x_4) where for simplicity we have assumed $\lambda = 1$. The two oscillators and the coupling constants α and β are given in (5.10) where we have introduced a term ω_i to allow for frequency variation in the oscillators (this term is added by time scaling and noting that the term $H.\partial H/\partial x_i$ reduces to zero on the limit cycle).

$$\begin{aligned}
\dot{x}_i &= \omega_i x_{i+1} - (x_i^2 + x_{i+1}^2 - 1)x_i - \alpha(x_i - x_{i-2}) - \alpha(x_i - x_{i+2}) \\
\dot{x}_{i+1} &= -\omega_i x_i - (x_i^2 + x_{i+1}^2 - 1)x_{i+1} - \beta(x_{i+1} - x_{i-1}) - \beta(x_{i+1} - x_{i+3})
\end{aligned} \tag{5.10}$$

Introducing the coordinate system

$$p = x_1 - x_3, \; q = x_2 - x_4$$

then for two oscillators and using (5.10) yields

$$\begin{aligned}
\dot{p} &= -2\alpha p - r^2 p + \omega_1 x_2 - \omega_2 x_4 - x_1^3 - x_1 x_2^2 + x_3 x_4^2 + x_3^3 \\
\dot{q} &= -2\alpha q - r^2 q - \omega_1 x_1 + \omega_2 x_3 - x_2^3 - x_1^2 x_2 + x_3^2 x_4 + x_4^3
\end{aligned} \tag{5.11}$$

where 1:1 synchronisation corresponds to the origin. We can investigate stability of the synchronised state by suggesting a candidate Lyapunov function

$$V = p^2 + q^2$$

Consider its time derivative, initially for identical frequency oscillators, $\omega_1 = \omega_2$,

$$\begin{aligned} 0.5\dot{V} = & -p^2\left(2\alpha - r^2 - x_1x_3 - x_2x_4\right) \\ & -q^2\left(2\beta - r^2 - x_1x_3 - x_2x_4\right) \\ & -\left(x_1^4 + x_2^4 + x_3^4 + x_4^4 + 2x_1^2x_2^2 + 2x_3^2x_4^2 + 2(x_1x_3 - x_2x_4)^2\right) \end{aligned} \qquad (5.12)$$

Clearly, the third bracketed term in (5.12) is negative definite. The first two terms are also negative definite for sufficiently large values of the coupling α, β subject to the condition that x_i is bounded. Thus, if (5.12) is negative definite, the origin of the (p, q) phase plane is asymptotically stable, and the two oscillators synchronise. It remains to show the bounded condition and this may be achieved by suggesting a new Lyapunov function for (5.12) for the case of two oscillators

$$V^* = x_1^2 + x_2^2 + x_3^2 + x_4^2$$

Its time derivative is given by

$$\begin{aligned} 0.5\frac{dV^*}{dt} = & -x_1^2\left(x_1^2 + x_2^2 - r^2\right) - x_2^2\left(x_1^2 + x_2^2 - r^2\right) \\ & -x_3^2\left(x_3^2 + x_4^2 - r^2\right) - x_4^2\left(x_3^2 + x_4^2 - r^2\right) \\ & -\alpha\left(x_1 - x_3\right)^2 - \beta\left(x_2 - x_4\right)^2 \end{aligned}$$

and is negative far from the origin and therefore all distant trajectories approach the 2-torus at the origin and hence are bounded.

Now, if the two oscillators under consideration have differing frequencies $\omega_1 \neq \omega_2$, a new term is added to (5.12)

$$(\omega_1 - \omega_2)(x_2x_3 - x_1x_4)$$

This is a sign-indefinite term that violates Lyapunov's stability criterion for small x_i and results in limit cycle behaviour for (5.12) with a radius that decreases with increasing coupling and closer frequencies (converging to a stable equilibrium point for identical oscillators). This translates to a slight time lag in synchronization between neighbouring oscillators and does not pose any problems for our purposes as we may make the lag arbitrarily small by increasing the coupling strength. To model evoked potentials, we employ a chain of N coupled oscillators given by (5.10) where $i = 1 \ldots 2N-1$, $r^2 = 1$, $\lambda = 1$. As a first approximation, we consider the brain and surrounding

structures as an infinite homogenous conductor, the contribution to the field potential in volts for each current source I_i in amperes is

$$\Phi(\vec{r},t) = \frac{1}{4\pi\sigma}\sum_{i=1}^{N}\frac{I_i(t)}{R_i}$$

where σ is the brain tissue conductivity and R_i is the distance of the i^{th} current source from the field point (PLONSEY and BARR, 1998). In our model, the output of each oscillator in the chain is taken to represent the generated extracellular current of a neural population. Considering the chain as uniform, all the coefficients R_i are equal, $R_i = R$. Ignoring scaling constants the EEG can be modelled as a simple summation of the output of each oscillator. Finally, to simulate the unsynchronised background EEG, the set $\{w_i\}$ is chosen from a random Gaussian distribution N(1, 0.33).

5.3.3 Exploring the Model I

To explore the model, a custom application EPSync written in Java was created. Based on a fourth order, fixed step Runge-Kutta solver (PRESS *et al*, 1992), the application employs a graphical interface to configure the simulation parameters and (more uniquely) to specify the coupling time series to use. Figure 5.9 illustrates the class diagram. The source code can be found in Appendix A3 (CD-ROM).

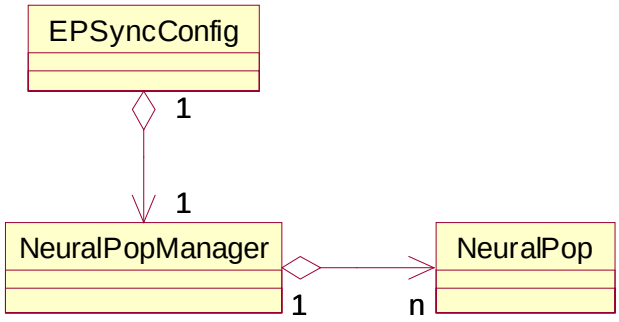

Figure 5.9: UML class diagram for EPSync

Each oscillator is modelled via the NeuralPop class, which interacts with other neighbour oscillators in an object-oriented fashion. The public method updateState() is called repetitively by the NeuralPopManager class to advance the solution. The EPSyncConfig class wraps the GUI interface and instantiates the NeuralPopManager. Different oscillators can be used in place of NeuralPop and different network topologies may be specified within the BuildTopology() method in NeuralPopManager.

Figure 5.10 illustrates the effect of a pulse in coupling strength on a pair of incommensurate coupled oscillators (ω_1 = 1.0 and ω_2 = 1.6). Following the increase in coupling strength, the oscillators become immediately synchronized and subsequently drift apart again when the coupling returns to zero. The oscillators synchronise at the average frequency of the uncoupled oscillators - a phenomenon that has been shown to occur for a chain of coupled phase oscillators (RAND *et al*, 1988).

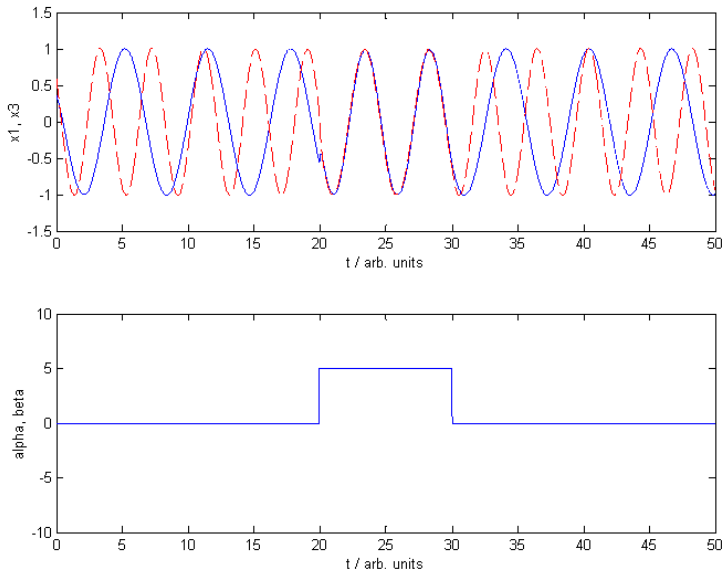

Figure 5.10: A pair of oscillator undergoing transient synchronisation. Coupling (α,β) time series is shown at the bottom

To simulate an evoked potential, 50 oscillators with Gaussian distributed frequencies are subjected to a transient increase in coupling strength as a result of some sensory stimulus. Figure 5.11 illustrates the summated output of the first state variable of each oscillator, subject to a given coupling time series. The simulation models an evoked potential as the oscillators are briefly brought into phase and subsequently go out of phase due to their differing frequencies. During the interval of synchronicity, and triggered by sensory stimulation, the complexity of the system briefly reduces to simple limit cycle behaviour. That the complexity of the dynamics can simplify in our model under coupling is consistent with BAŞAR (1983). Başar argued that neural populations in the brain might be regarded as a large number of coupled oscillators and suggested that if *n* oscillators are left uncoupled, the resulting attractor will be an *n*-dimensional torus corresponding to *n* independent frequencies. If the oscillators are coupled, the dimension will be reduced, possibly resulting in simple limit cycle behaviour.

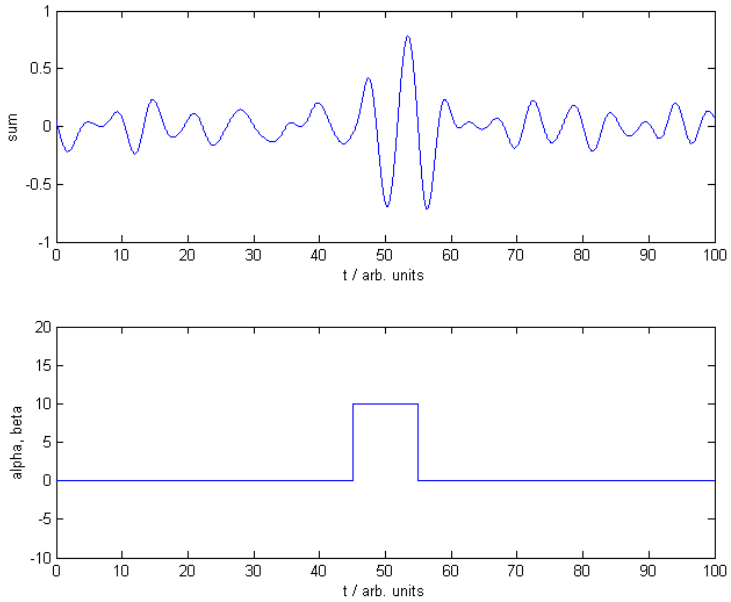

Figure 5.11: Summated output of 50 oscillators with Gaussian distributed frequencies N(1, 0.33) undergoing transient

5.3.4 Phase-locking the Response to the Stimulus

The model proposed thus far exhibits one aspect of phase reordering, that of phase synchronisation. However, no phase-locking to the trigger is displayed. This is clearly demonstrated in real evoked potentials where there is a small variance in the latency of the evoked response with respect to the trigger (hence the origin of names such as P100 and N145, which refer to positive and negative evoked response peaks that occur at 100 ms and 145 ms after a stimulus respectively). Indeed, averaging over many trials is the principal technique used in clinical practice for recovering the evoked potential from the background (much larger) EEG signal (MISULIS, 1994). To model phase-locking to the trigger, it is necessary to include a mechanism for *phase resetting* the oscillators. We proceed by first studying phase resetting for the oscillators given by (2.7) and (5.9) but with a change of coordinates. Our approach here is based on a modified version of MURRAY (1990). We will show that by adding a stimulus component corresponding to a volley of afferent activity impinging on each oscillator, the evoked response is phase-locked to the stimulus.

We transform (2.7) with (5.9) to polar coordinates by specifying the transform

$x_1 = r\cos\theta$

$x_2 = r\sin\theta$

where $r = \sqrt{\left(x_1^2 + x_2^2\right)}$. This yields

$$\frac{dr}{dt} = -\lambda r\left(r^2 - 1\right)$$
$$\frac{d\theta}{dt} = -\omega \tag{5.13}$$

The steady-state solution of (5.13) is given by

$$r(t) = 1,\ \ \theta(t) = -\omega t + \theta(t_0)$$

We wish to study the change in phase of (5.13) subject to a stimulus. A convenient graphical tool is the phase transition curve (PTC), which gives the new phase versus the old phase subject to a stimulus of a given magnitude applied at the old phase (WINFREE, 1980). Figure 5.12 illustrates (5.13) subject to an instantaneous (without loss of generality) horizontal stimulus of magnitude *I*. The effect of the stimulus is to displace the trajectory horizontally to the left from *P* to *Q* resulting in a new phase angle ϕ. Noting that

$$m\cos\phi = \cos\theta - I$$
$$m\sin\phi = \sin\theta$$

we can write the new phase angle in terms of the old angle

$$\phi = \tan^{-1}\left(\frac{\sin\theta}{\cos\theta - I}\right) \tag{5.14}$$

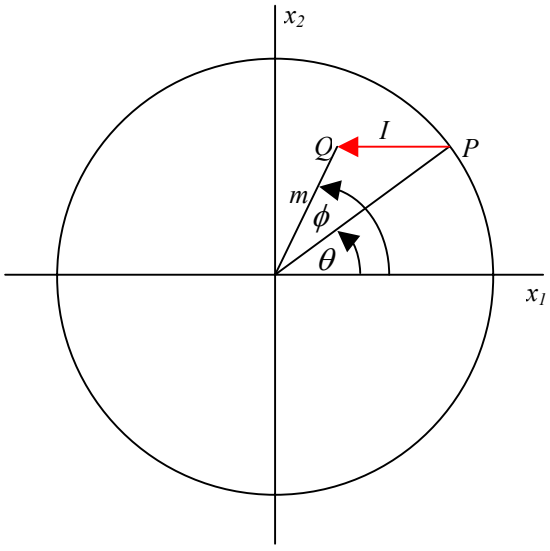

Figure 5.12: Impulse applied to (5.14) with $|I| < 1$

Differentiating (5.14) with respect to θ yields

$$\left(1 + \tan^2\phi\right)\frac{d\phi}{d\theta} = \frac{1 - I\cos\theta}{(\cos\theta - I)^2} \tag{5.15}$$

Now, for $|I| < 1$, we can see from (5.15) that $d\phi/d\theta > 0$ always and thus the new phase is always delayed with respect to the old phase. Figure 5.13 illustrates a PTC for

I = 0.5. The average slope is equal to 1 and hence is called a Type 1 phase resetting curve (WINFREE, 1980). Note that any value of ϕ is attainable for a range of values of θ.

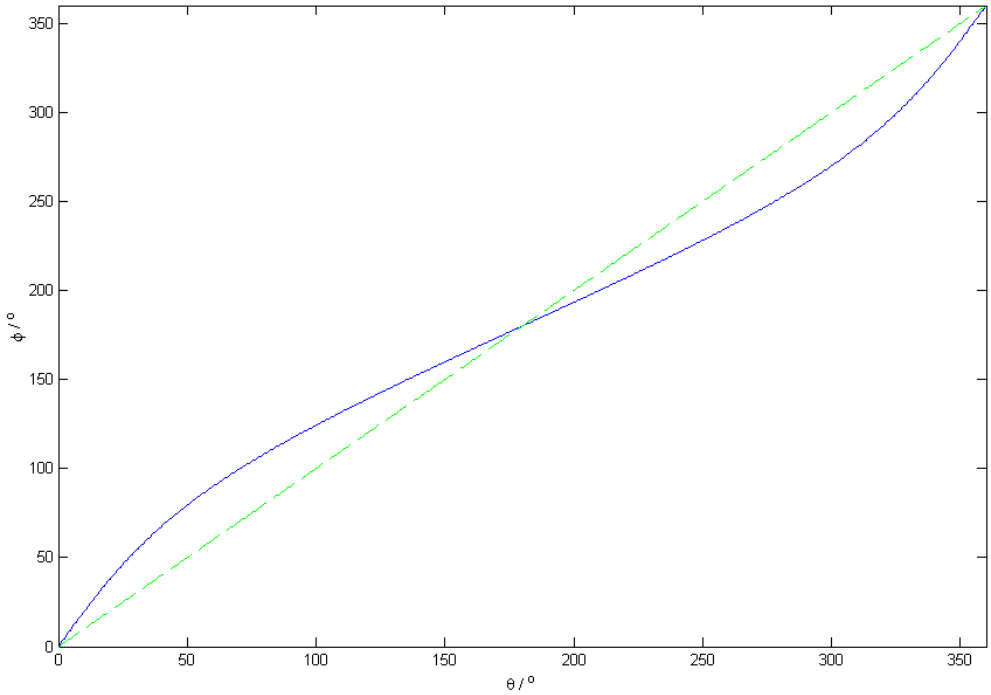

Figure 5.13: Type I phase resetting (I = 0.5). Dashed line corresponds to $\phi = \theta$

Now consider the case for $|I| > 1$. According to (5.15) $d\phi/d\theta$ can now take on negative values. Figure 5.14 illustrates a phase plot.

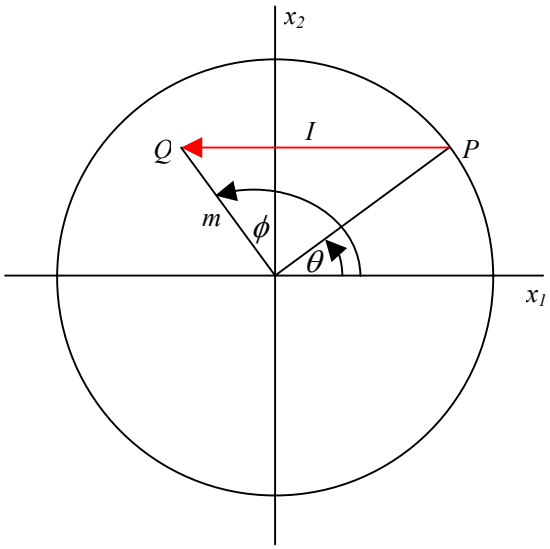

Figure 5.14: Impulse applied to (5.14) with $|I| > 1$

The effect of the stimulus on the oscillator is to always move the state into the left hand plane. This is reflected in the PTC where the new phase is clearly constrained to

a subset of values as illustrated in Figure 5.15 for I = 1.5. This is known as Type 0 phase resetting since the average slope of the curve is 0 (WINFREE, 1980).

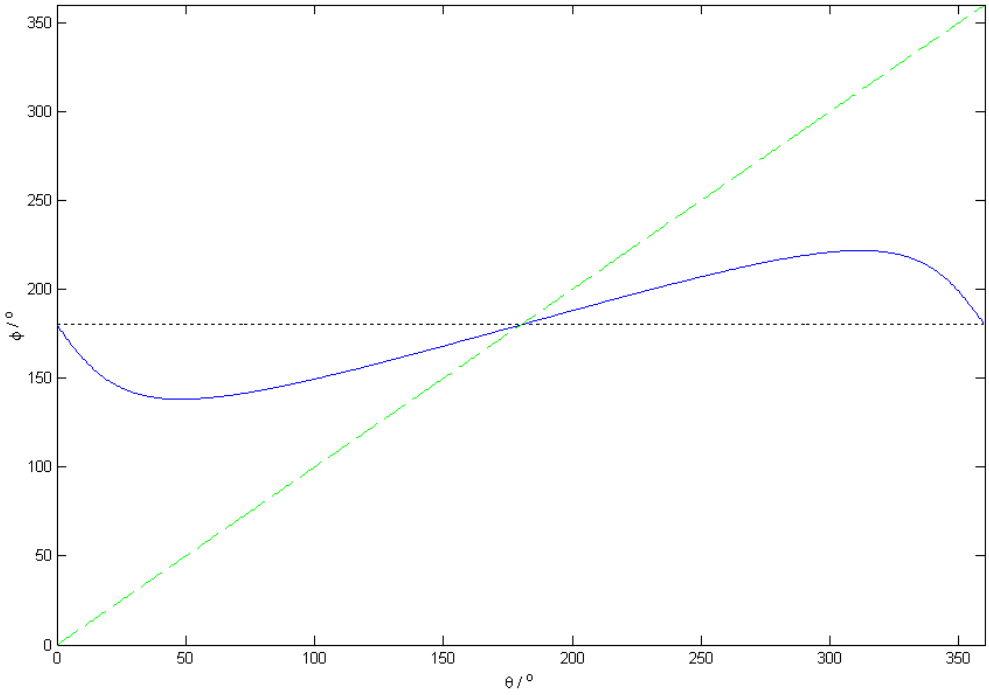

Figure 5.15: Type 0 phase resetting (I = 1.5). Dashed line corresponds to $\phi = \theta$

We can make use of Type 0 phase for a sufficiently large stimulus ($|I| > 1$) to achieve phase-locking to the trigger. Subjecting a large number of uncoupled oscillators with (initially) uniformly distributed random phases to a stimulus ($|I| > 1$) results in a distribution of new phases about a mean value as illustrated by the dotted horizontal line in Figure 5.15. By including a subsequent phase synchronisation effect, the oscillator phases are 'pulled together' causing the standard deviation of the new phases about the mean to reduce to zero. The end result is that the synchronised oscillators are phase locked to the stimulus trigger. In other words, we have now introduced a phase reordering effect comprised of both phase resetting and phase synchronisation. The additional phase resetting mechanism also serves to bring the oscillators into synchronisation faster than with just the synchronisation effect alone. We extend the oscillators (5.10) used in the macroscopic model of evoked potentials by adding a brief stimulus to the first state variable corresponding to a volley of afferent activity. Since real stimuli are not instantaneous, we model the stimulus pulse as a distribution given by

$$I(t) = a \exp\left(-\left(\frac{t-\tau}{b} \right)^2 \right) \tag{5.16}$$

Figure 5.16 illustrates a plot of equation (5.16).

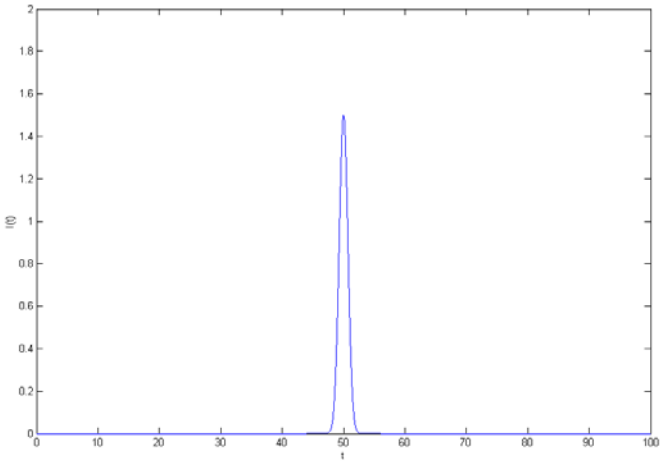

Figure 5.16: Plot of (5.16) for $a = 1.5$, $b = 1$, $\tau = 50$

5.3.5 Exploring the Model II

5.3.5.1 Ensemble Averaging

We repeat the numerical experiment of subjecting 50 oscillators with Gaussian distributed frequencies to a transient increase in coupling strength but this time with the simultaneous addition of a phase resetting impulse to each oscillator and ensemble averaging over 10 responses with random initial conditions for each oscillator. Figure 5.17 illustrates the summated output of the first state variable of each oscillator (top) with the coupling strength waveform (middle) and stimulus imparted (bottom).

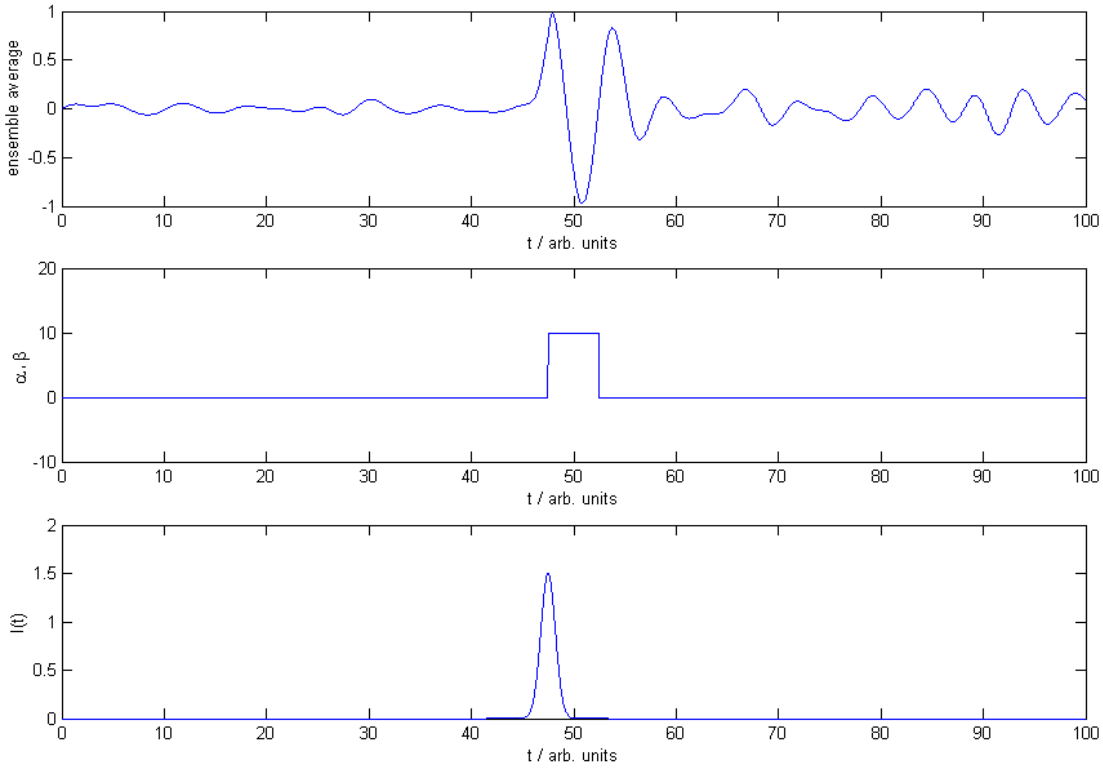

Figure 5.17: Ensemble average over 10 evoked responses (top) subject to a change in coupling strength (middle) and stimulus (bottom)

From Figure 5.17 two obvious effects as a result of the introduction of the stimulus can be observed. The first is the oscillators synchronise faster than before (compare with Figure 5.11) and we have been able to reduce the width of the coupling pulse required to synchronise the oscillators. Secondly, the evoked response is now preserved during averaging, i.e. the response is phase-locked to the trigger.

5.3.5.2 Arousal Level Effects

In Section 5.3.1 we described how the strength of coupling is also dependent on the level of arousal. Application of anaesthetics can reduce the level of arousal and changes in auditory and visual evoked potentials have been used in methods for detecting depth of anaesthesia (NOGAWA, 1991). In general, lowering of the arousal level is seen to increase the latency and often reduce the amplitude of auditory and visual evoked potentials for many anaesthetics (THORNTON 1991; NOGAWA 1991). We can model this phenomenon by reducing the coupling strength to correspond to the low arousal occurring during anaesthesia. The stimulus received by all oscillators is also expected to be reduced significantly as the strength of the synapses is dramatically decreased by the application of anaesthetics (FREEMAN, 1979). We note from the PTC for a weak stimulus $|I|<1$, the effect is to delay some oscillators and advance others by a small amount. On average the slope is 1, and measured across a large number of initially unsynchronised oscillators, the net change in phases is zero.

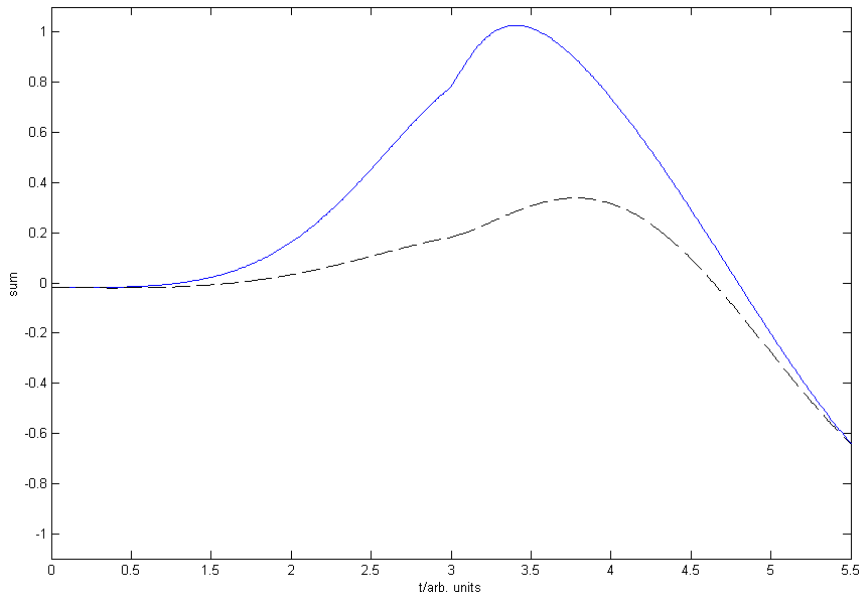

Figure 5.18: Latency increase and amplitude decrease for lower coupling and stimulus (single-trial, 50 oscillators). Coupling and stimulus as per Figure 5.17 (solid) and as a result of a decrease by a factor of 5 (dashed)

The proposed model suggests that the lowering in arousal (coupling and stimulus strength) will result in latency in the evoked response peak by virtue of the fact that it will take longer for the oscillators to arrive at a synchronized state (the phase resetting component has no net effect). In addition, since the oscillators will not synchronize as tightly, a lower amplitude response will result. Figure 5.18 illustrates the onset of synchronisation for the original stimulus and coupling strength, and for a decrease in stimulus and coupling by a factor of 5.

5.3.5.3 Desynchronisation Waveform

REILLY *et al* (1996) demonstrated that the result of resetting all the phases of n uncoupled sinusoidal oscillators resulted in the summated output of

$$y(t) = n \exp(\frac{-\sigma^2 t^2}{2}) \sin \mu t \tag{5.17}$$

where n is the number of oscillators, σ is the standard deviation and μ is the mean of a Gaussian distribution of frequencies. It turns out that this equation is a scaled version of the Morlet wavelet. The output of the first state variable of the oscillators described by (2.7) and (5.9) can be described by a simple sinusoidal function and thus the summated desynchronisation of a large number of oscillators should also approximate the Morlet function. Figure 5.19 illustrates a repeat of Figure 5.17 where the time axis starts at the instant the coupling falls to zero and thus the oscillators are drifting apart at this point. Superimposed (dashed) is a scaled version of (5.17).

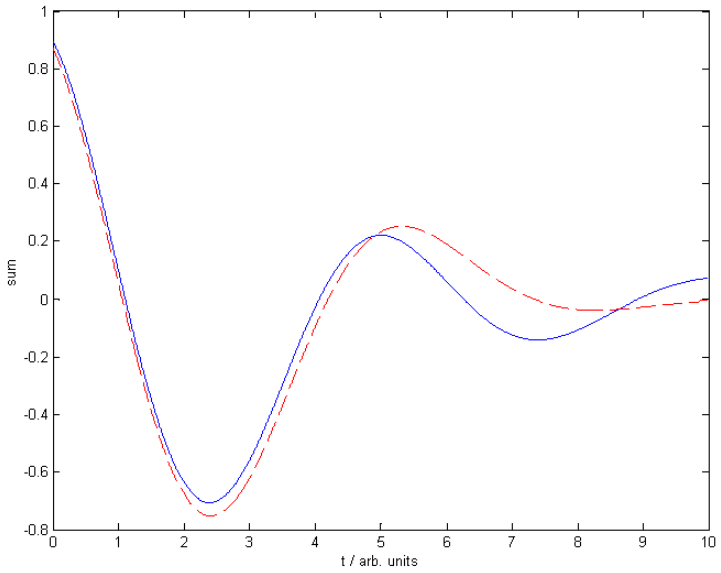

Figure 5.19: Response of 50 oscillators desynchronising at $t = 0$ (solid). Suitably scaled Morlet wavelet (dashed)

DEMIRALP *et al* (1998) demonstrated the usefulness of the Prony method in the analysis of similar potentials. This method adaptively fits damped sinusoids to data and performs better than Fourier analysis for signals that are transient by nature. Another alternative technique might be to use wavelets to perform a multi-resolution decomposition. Wavelet theory, in general, does not provide a method for selecting a suitable basis function for a given data set, and is usually done by trial and error against a library of wavelets (MC DARBY, 2000). Our model seems to suggest a possible choice for an optimum wavelet for evoked potential decomposition.

5.4 Stochastic Synchronisation

The proposed model in the preceding section for evoked potential generation based on the transient synchronisation of coupled oscillators assumes a deterministic system for the oscillators in question. In Chapter III, we introduced a model of the EEG based on stochastic limit cycle oscillators to capture the aperiodic variability seen in the actual EEG. It is interesting to consider employing stochastic oscillators to represent the activity of a neural population in an effort to model the variability of activity in real populations. The question naturally arises as to whether stochastic oscillators can synchronise and appears new. PECORA and CARROLL (1990) observed that certain subsystems of nonlinear chaotic subsystems can be made to synchronise and has spurred great interest in the application of chaos in communication systems (KOLUMBÁN *et al*, 1997; KOLUMBÁN *et al*, 1998). We conclude this chapter with a brief demonstration study of two coupled stochastic limit cycle oscillators of identical frequency. We show that synchronisation is indeed possible.

Consider the stochastic ellipse limit cycle oscillator (using (3.20)) written here for two coupled oscillators

$$dX = \begin{pmatrix} \frac{2X_2}{b^2} - \lambda\left(\frac{X_1^2}{a^2}+\frac{X_2^2}{b^2}-1\right)\frac{2X_1}{a^2} - \alpha\left(X_1 - X_3\right) \\ -\frac{2X_1}{a^2} - \lambda\left(\frac{X_1^2}{a^2}+\frac{X_2^2}{b^2}-1\right)\frac{2X_2}{b^2} - \beta\left(X_2 - X_4\right) \\ \frac{2X_4}{b^2} - \lambda\left(\frac{X_3^2}{a^2}+\frac{X_4^2}{b^2}-1\right)\frac{2X_3}{a^2} - \alpha\left(X_3 - X_1\right) \\ -\frac{2X_3}{a^2} - \lambda\left(\frac{X_3^2}{a^2}+\frac{X_4^2}{b^2}-1\right)\frac{2X_4}{b^2} - \beta\left(X_4 - X_2\right) \end{pmatrix} dt + \sigma \begin{pmatrix} -\lambda\frac{2X_1}{a^2} & 0 & 0 & 0 \\ 0 & -\lambda\frac{2X_2}{a^2} & 0 & 0 \\ 0 & 0 & -\lambda\frac{2X_3}{a^2} & 0 \\ 0 & 0 & 0 & -\lambda\frac{2X_4}{a^2} \end{pmatrix} dW$$

and

$$dX = \begin{pmatrix} dX_1 \\ dX_2 \\ dX_3 \\ dX_4 \end{pmatrix}, \quad dW = \begin{pmatrix} dW_1 \\ dW_2 \\ dW_3 \\ dW_4 \end{pmatrix}$$

where we have introduced coupling terms α and β similar to the approach used in Section 5.3.2. Figure 5.20 illustrates the first state variable for the two uncoupled oscillators (obtained by setting the coupling terms to 0).

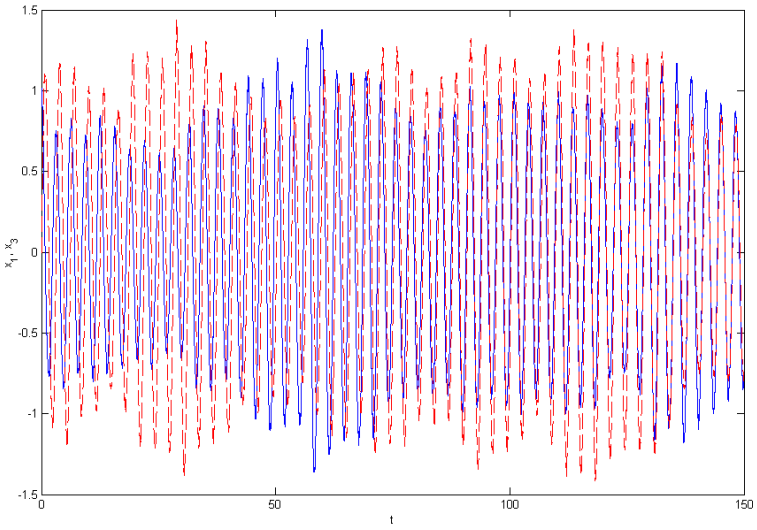

Figure 5.20: Plot of the first state variable for two uncoupled stochastic limit cycle oscillators with differing initial conditions. Simulation parameters: Runge-Kutta, $a = b = 1$, $\lambda=0.01$, $\sigma=4$, $\Delta t = 0.005$, $T = 150$

Figure 5.21 illustrates the effect for $\alpha = \beta = 0.5$ using the same Wiener process realisation for comparison.

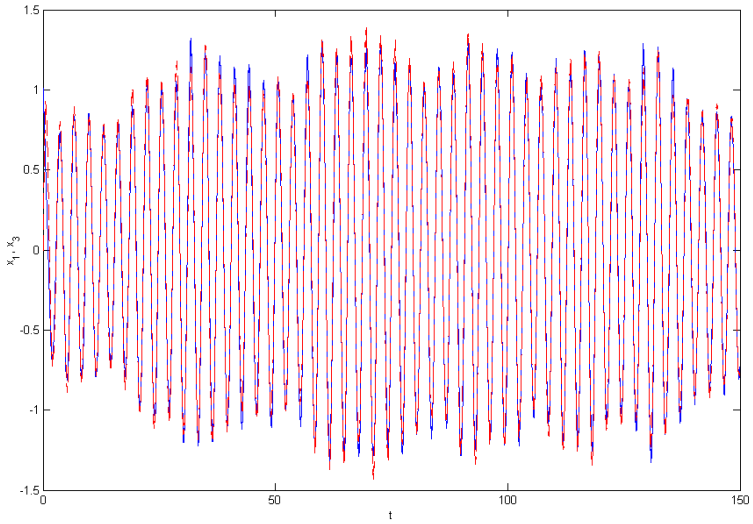

Figure 5.21: Plot of the first state variable for two coupled stochastic limit cycle oscillators ($\alpha = \beta = 0.5$) with differing initial conditions. Simulation parameters: Runge-Kutta, $a = b = 1$, $\lambda=0.01$, $\sigma=4$, $\Delta t = 0.005$, $T = 150$

Even for small values for the coupling terms, the oscillators synchronise both phase and amplitude immediately. The model for evoked potentials presented in the previous section could be extended to include stochastic dynamics. The introduction of stochastic effects would enable variability in amplitude and latency of the response to be modelled similar to that demonstrated with the stochastic Wilson-Cowan model in Section 3.7.3.

5.5 Summary

The key concept explored in this chapter is that afferent (sensory) stimulation naturally triggers synchronisation at both the microscopic (neuron) and macroscopic (neural population) levels. A small network of biologically realistic neurons with realistic synaptic coupling demonstrates the concept at the microscopic level. Coupled oscillators of the form introduced in Chapter II serves as the basic unit of the macroscopic model. Both phase resetting and phase synchronisation were shown to play important roles in the generation of evoked potentials with the (macroscopic) model capable of predicting anaesthesia related latency in evoked potentials. The suitability of the Morelet wavelet in analysis was also discussed. The chapter was concluded with a novel demonstration of how stochastic limit cycle oscillators can synchronise. Given the ubiquity of noise and synchronisation in the nervous system, this phenomenon appears quite likely to be important in brain processes.

Evoked (or indeed event-related) potentials can be viewed as a deterministic component in an otherwise stochastic signal. We exploit this in Chapter VI, where we discuss practical applications by introducing some novel algorithms for use in human-computer interfaces.

CHAPTER VI

DIRECT BRAIN INTERFACE APPLICATIONS

6.1 Introduction

Influenced by work in the preceding chapters, in particular the characterisation of the EEG as a stochastic process, we now explore the more pragmatic concerns of biosignal processing encountered in creating Direct Brain Interface (DBI) systems. A DBI typically operates by harnessing signals arising from processes within the brain, without depending on the brain's normal output pathways of peripheral nerves or muscles (WOLPAW *et al*, 2000). Whilst our main signal of interest in this chapter (and indeed thesis) is the EEG, the bigger picture of facilitating communication and control for people with severe physical disabilities should not be forgotten. To this end, we shall also briefly consider some alternatives to EEG-based DBIs. In fact, possessing a diversity of solutions at hand, with the most suitable solution fitted to the person's needs and not vice versa, is the hallmark of effective Occupational Therapy (FLYNN, 2001).

Central to direct brain interfacing strategies are the signal processing tasks of feature extraction and statistical pattern recognition, which are the focus of this chapter. After reviewing the state-of-the-art in EEG-based DBIs, we introduce the concepts of statistical pattern recognition and parametric modelling. A left/right self-paced typing exercise is analysed by novel extension of the common AutoRegressive (AR) model for EEG feature extraction to an AutoRegressive with Exogenous input (ARX) model for combined filtering and feature extraction (BURKE *et al*, 2002). Using this technique, event-related and evoked potential based DBIs may be constructed, for example the P300 based DBI by FINUCANE *et al* (2003). As a sidebar, we consider an alternative communication paradigm relevant for severe physical disabilities based on eye-blink detection employing image processing (BURKE *et al*, 2001). The fMRI signal is also investigated using the ARX feature extraction method to facilitate single-trial extraction of haemodynamic activation information in the brain, a possible future strategy applicable in constructing noninvasive DBIs.

6.2 A Brief Survey of EEG-Based Direct Brain Interfaces

The most common signal employed in DBIs is the non-invasive, scalp recorded EEG, which harnesses features originating from somatosensory, motor, or visual areas of the cortex (FARWELL and DONCHIN, 1988; WOLPAW *et al*, 1991; PFURTSCHELLER *et al*, 1996; SUTTER, 1992; MIDDENDORF *et al*, 2000; BIRBAUMER *et al*, 1999). Complex digital signal processing and pattern recognition techniques are often employed to extract pertinent information from the measured signals in real-time and subsequently used to drive communication and control applications. When designing a DBI, in addition to the technical methodology of recording, feature extraction, and pattern recognition, the choice of EEG feature and elucidation paradigm is paramount. Broadly speaking, most EEG-based DBIs fall into one of four categories depending on the specific activity patterns of interest including:

- oscillatory EEG components,
- slow cortical potentials,
- event-related potentials,
- evoked potentials.

A coarser categorisation may be made by noting that the former two refer to phenomena that require operant conditioning and/or behavioural modification to control. WOLPAW *et al* (1991) exploit behavioural modification of the 8-12 Hz mu[1] rhythm or 20-25 Hz central beta rhythm (both of which originate in the sensorimotor cortex) as a control signal for 1-dimensional and 2-dimensional cursor control. PFURTSCHELLER *et al* (1996) harness Event-Related Desynchronisation (ERD) and Event-Related Synchronisation (ERS) of the 8-12 Hz mu and central beta rhythms modified by imagined movements as control inputs for a DBI (contralateral ERD and subsequent ipsilateral ERS occur during and after preparation for movement respectively). BIRBAUMER *et al* (1999) have constructed a spelling device that trains locked-in patients to self-regulate slow cortical potentials in the EEG (although it takes several weeks to months to learn to operate satisfactorily). FARWELL and DONCHIN (1988) employ the P300 endogenous potential to facilitate users to select letters of the alphabet for communication. The P300 potential is elicited by subjects attending to a 'rare' event in an oddball sequence of stimuli (in this case, the subject attends to a

[1] The mu rhythm is the name given to the alpha-like rhythm located over the sensorimotor cortex. It is related to the usual alpha rhythm in frequency and amplitude but has a different physiological significance (NIEDERMEYER, 1999).

particular letter while other 'odd' letters are presented at random; the chosen letter is the rare event). Based on work by SUTTER (1992), MIDDENDORF *et al* (2000) created a simple DBI based on steady-state visual evoked potentials. Briefly, two icons flashing at different frequencies are presented on a computer screen. As the subject gazes at one icon (which may have a control or communication action associated with it), a visual evoked response at the stimulus frequency of the particular icon will ensue and thus one may determine which icon is being attended to by analysing the EEG signal. Notably, the signal evoked by the stimulus increases enormously the closer to the centre of the visual field it is located and is a direct result of the disproportionately large amount of visual cortex allocated to processing at the centre of our vision where acuity is greatest (SUTTER, 1992).

6.3 Statistical Pattern Recognition

Pattern recognition loosely refers to the concept of extracting meaning from data by the *classification* of information into a predefined set of *classes*. Familiar examples of pattern recognition include recognising faces, detecting familiar voices, and detecting smells, all of which are tasks that humans perform exceedingly well. Enabling machines to achieve comparable results is a challenge and the theoretical framework is based naturally on a probabilistic one, which reflects the fact that the information we wish to describe is usually described in a statistical sense and the results we seek are also often expressed in a statistical sense. In this section, we shall introduce sufficient pattern recognition material to present the results of this chapter; for a detailed introduction, the interested reader is referred to BISHOP (1995) and RIPLEY (1996).

Figure 6.1 outlines the typical pattern recognition approach consisting of two stages, namely a *feature extraction* stage and a *classification* stage.

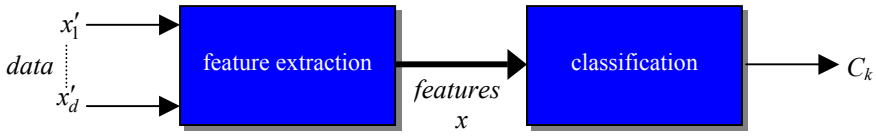

Figure 6.1: Typical architecture for pattern recognition

The former usually results in a transformation of the data $x'_1 \ldots x'_d$ in addition to a dimensionality reduction to produce a feature vector x. Reducing the amount of data entering the classification stage often increases the accuracy of the later stage (it might

be helpful to consider that higher dimensional data can result in an ineffective sparse representation of all classes when the available data is limited). As we shall see later, effective pattern recognition is critically dependent on a good feature extraction strategy.

The goal of the classification stage is, given the feature vector x, classify the data into one of a predetermined set of classes C_k. The classifier may be trained by a *supervised* or *unsupervised* paradigm. For the supervised case, target values are supplied, i.e. the features x belonging to class C_k are available. For the unsupervised case the aim might be to separate the data into distinct clusters. In general, the overall scheme of Figure 6.1 may be viewed as a mapping of a set of input variables $x'_1 \ldots x'_d$ to a set of output variables $y_1 \ldots y_c$ representing the class labels

$$y_k = y_k(x', w) \tag{6.1}$$

where x' refers to the vector of inputs and w represents the vector of parameters that may be adjusted to fit (6.1) to the data. The functions in (6.1) may be very general, for example, neural network classifiers may be represented by nonlinear functions, where w are interpreted as weights of the network. There is a trade-off in the complexity of the functions in (6.1). The central goal of (6.1) is to support *generalisation*, that is, the ability to make good predictions on previously unseen data. If the structure of the functions is deficient, the equations are unable to model the data well and the model is said to possess large *bias*. On the other hand, if the order of the model is too high for example, it might over fit the data resulting in high *variance* and poor generalisation.

To discuss the statistical approach to pattern recognition, it is instructive to introduce a fictional system, which we take to be a device capable of detecting smells. The detection process commences with an analysis of a portion of air and a resultant feature vector x is generated based on a chemical analysis of the sample. The feature vector, comprising the concentrations of key chemicals, is supplied to a classifier, which must determine which of k classes C_k the feature vector belongs to. The *prior* probabilities $P(C_k)$ correspond to the likelihood of class C_k. If we were forced to make a decision without seeing the feature vector, the class with the highest prior probability would result in an above-chance classification. However, if the feature vector data is

available, we have much more information to base a classification on. An example of useful information to ascertain are the class conditional probabilities as illustrated in Figure 6.2 (for a two class problem).

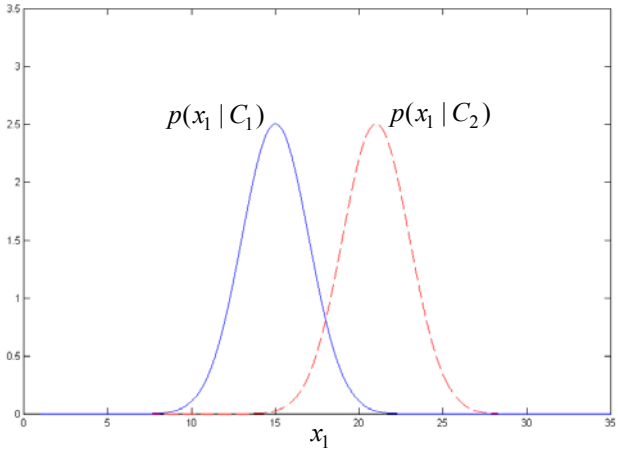

Figure 6.2: Overlapping class conditional probabilities. The probability of misclassification may be minimised by choosing the larger of the two probabilities given the feature vector component value x_1

If a component x_1 of the feature vector corresponds to the concentration of a particular chemical, the input may be classified as smell C_1 for a low concentration of the chemical and to smell C_2 for a high concentration of the chemical. Determining which class a feature belongs to for an intermediate concentration requires more analysis but it can be shown that the intuitive result of choosing the class C_k for which $p(x_1 \mid C_k)$ is largest gives the minimal error of misclassification (BISHOP, 1995). Of course, the quantity of most interest in statistical pattern recognition is the conditional probability $P(C_k \mid x)$ i.e. the probability that given a feature vector x, the most likely corresponding class is C_k. Bayes' theorem is a useful result for manipulating such quantities and may be written as

$$P(C_k \mid x) = \frac{p(x \mid C_k)P(C_k)}{p(x)} \tag{6.2}$$

The usefulness of (6.2) stems from the fact that it is easier to evaluate the r.h.s. than the posterior probability on the l.h.s. In the context of statistical pattern recognition, (6.2) may be expressed as *the probability that a feature vector belongs to a class C_k is equal to the product of its class conditional probability and its prior probability subject to a normalisation quantity (the denominator term).*

6.3.1 Linear Discriminant Analysis

Given an input vector *x* of features, the classification problem may be formulated in terms of a set of discriminant functions $y_1(x), y_2(x), ..., y_c(x)$ where an input vector x is assigned to class C_k if

$$y_k(x) > y_j(x) \qquad for\ all\ \ j \neq k \tag{6.3}$$

i.e. choose the class for which the corresponding discriminant function is largest. We may derive a discriminant function in terms of Bayes' theorem that minimises the probability of misclassification (BISHOP, 1995) by letting

$$y_k(x) = P(C_k \mid x) \tag{6.4}$$

By omitting the class independent probability from (6.2), and taking the logarithm (valid according to (6.3) for any monotonic function) results in

$$y_k(x) = \ln(p(x \mid C_k) + \ln(P(C_k)) \tag{6.5}$$

In this chapter, we will employ parametric classifiers, that is, classifiers that assume a specific functional form for the probability density function. A natural choice is the normal density function (appearing ubiquitously in nature thanks to the central limit theorem), given in *d* dimensions as

$$p(x) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left\{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)\right\} \tag{6.6}$$

where the mean vector and covariance matrix are given respectively by

$$\mu = E(x)$$
$$\Sigma = E\left[(x-\mu)(x-\mu)^T\right]$$

By assuming that covariance matrices are identical for all classes $\Sigma_k = \Sigma$, and neglecting constant terms, one forms the linear discriminant function

$$y_k(x) = \mu_k^T \Sigma^{-1} x - \frac{1}{2}\mu_k^T \Sigma^{-1} \mu_k + \ln(P(C_k)) \tag{6.7}$$

In the sequel, we refer to application of (6.7) as Linear Discriminant Analysis (LDA).

6.3.2 Cross-validation

In practice, the quantity of labelled data available to train a classifier such as (6.7) is limited. The cross-validation procedure (BISHOP, 1995) is used in this chapter to obtain

an estimate of accuracy for a particular feature extraction and classification scheme. The order of a data set (consisting of feature vectors) is first randomly shuffled and subsequently divided into N distinct segments. N-1 segments are used to train a LDA classifier and the remaining segment is used as the test set for which a classification accuracy is determined. This process is repeated for each of the N possible test sets and the mean test set accuracy is computed. Finally, the complete process is repeated M times with a new random shuffle each time to yield an overall mean accuracy. Figure 6.3 illustrates the cross-validation process.

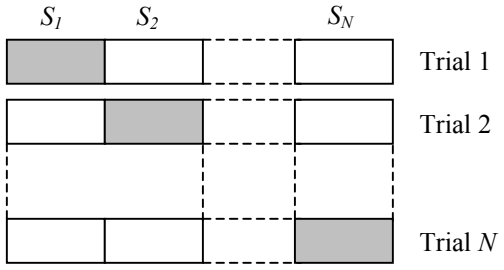

Figure 6.3: Cross-validation procedure. For each trial a segment is reserved as the test set (shaded) and the remainder used to train the classifier. See text for details. Reproduced from BISHOP (1995)

6.4 ARX Based Feature Extraction

In this section, we explore a novel method for feature extraction and subsequent classification for application in DBIs based on evoked or event-related potentials. Before presenting the results of applying the technique, we review some of the more popular methods for feature extraction employed in DBIs and introduce the basic theory of parametric modelling.

6.4.1 Popular Feature Extraction Methods for DBIs

Although it is possible to use the EEG data samples themselves as features in a DBI for classification (e.g. DONCHIN *et al*, 2000) computational efficiency and accuracy is usually greatly improved by feature extraction to achieve a transformation and dimensionality reduction of the data. The most common feature used for analysing oscillatory components of the EEG is the power in specific frequency bands. For example, WOLPAW *et al* (1991) control vertical cursor movement by $\Delta V = G(A - I)$, where G is gain, A is EEG amplitude in a specific band (square root of power in 8-12 Hz mu rhythm or 20-25 Hz central beta rhythm), and I is an intercept. A more advanced method for analysing oscillating components employs an AR model for

feature extraction (PENNY *et al*, 2000) or spectral estimation (WOLPAW *et al*, 2000). To circumvent problems with stationarity, PFURTSCHELLER *et al* (1996) employ an Adaptive AutoRegressive (AAR) model in the feature extraction stage for a DBI based on ERD and ERS.

DBIs based on evoked or event-related potentials require special techniques for detection and measurement since the problem posed is somewhat different to other DBI types. Evoked and event-related potentials are of the order of a microvolt and must be extracted from the background EEG, which is of the order of 10 – 50 μV (CHILDERS *et al*, 1987). In clinical practice, the most common approach employs time domain averaging to extract the signal from the background noise (MISULIS, 1994). Clearly, from the perspective of creating a real-time communication and control interface, this method of averaging is unsuitable. A plethora of techniques for single-trial extraction have been published in the literature in the last 20 years (see CHILDERS *et al* (1987) for a review with application to event-related potentials), however the primary approach used in DBIs to date appears to be based mostly on cross-correlation. SUTTER (1992) employs cross-correlation of the EEG with a known response template with varying delays to determine which area of a computer screen a subject is attending to. MIDDENDORF *et al* (2000) employ a simplified version of this technique involving just two icons. PATMORE and KNAPP (1995) suggest a DBI that harnesses EOG and evoked potentials to track gaze, again by employing cross-correlation similar to Sutter (the EOG provides the primary signal for controlling a cursor and the evoked potential input is used to recalibrate errors due to drift of the EOG signal). HUGGINS *et al* (1999) use cross-correlation of a template electrocorticogram with the ongoing electrocorticogram originating from sensorimotor areas to detected repetitive voluntary motor movements. DONCHIN *et al* (2000) found a small, yet statistically significant increase in classification accuracy using coefficients of a discrete wavelet transform with a Daubechies wavelet over using the raw data samples for detecting the P300 event-related potential.

6.4.2 Parametric Models

By parametric models, we mean the family of linear transfer function models ranging from the simplest structures of the AutoRegressive (AR) and AutoRegressive with eXogenous input (ARX) models to the more general Box-Jenkins (BJ) model (LJUNG,

1987). The parametric modelling technique fits one of these structures to a sampled signal by adjusting its free parameters. Assuming the model provides a good approximation to the signal's observed behaviour, a wide range of applications may be found for the model including spectral estimation, data compression, and feature extraction.

Following LJUNG (1987), a general linear time discrete model can be written as

$$y(t) = \eta(t) + w(t)$$

where $\eta(t)$ is the noise-free output and $w(t)$ is the disturbance term. Assuming, without loss of generality a sampling interval of unity, the noise-free output may be written as

$$\eta(t) = G(q,\theta)s(t)$$

where *s(t)* is the input and $G(q,\theta)$ is a rational function of the shift operator *q*

$$G(q,\theta) = \frac{B(q)}{F(q)} = \frac{b_1 q^{-k} + b_2 q^{-k-1} + \ldots + b_{n_b} q^{-k-n_b+1}}{1 + f_1 q^{-1} + \ldots + f_{n_f} q^{-n_f}}$$

where we have assumed a delay of *k* samples. Similarly, the disturbance term may be written as

$$w(t) = H(q,\theta)n(t)$$

where *n(t)* is white noise and the rational function $H(q,\theta)$ is

$$H(q,\theta) = \frac{C(q)}{D(q)} = \frac{1 + c_1 q^{-1} + \ldots + c_{n_c} q^{-n_c}}{1 + d_1 q^{-1} + \ldots + d_{n_d} q^{-n_d}}$$

The general model may thus be written as

$$y(t) = G(q,\theta)s(t) + H(q,\theta)n(t)$$

and is known as the Box-Jenkins (BJ) model. Figure 6.4 illustrates the structure.

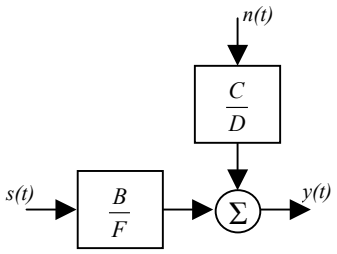

Figure 6.4: Box-Jenkins model structure

The parameter vector θ contains the coefficients b_i, c_i, d_i, and f_i. These parameters are chosen to fit the model to data after selecting the model orders n_b, n_c, n_d, n_f and delay k. In the next two sections we discuss relevant variants of the BJ model, and in each case, techniques to find the parameter vector θ that results in the optimum fit of the model to data.

6.4.2.1 AR Model

The AR model structure is the simplest of the Box-Jenkins family of models and is illustrated in Figure 6.5.

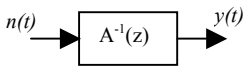

Figure 6.5: AR model structure

The time domain equation for the AR model is

$$y(t) = -a_1 y(t-1) - \ldots - a_n y(t-n_a) + n(t) \tag{6.8}$$

and the predicted value of *y(t)* is simply (drop the zero mean *n(t)* term)

$$\hat{y}(t) = -a_1 y(t-1) - \ldots - a_n y(t-n_a)$$

The forward prediction error may be written as

$$e(t,\theta) = y(t) - \hat{y}(t,\theta) \tag{6.9}$$

where, as before, θ is the vector of model parameters. To fit the model to the data *y(t)* collected over a period *t = 1, ..., N*, we choose $\hat{\theta}$ such that it minimises

$$E(\theta) = \frac{1}{N} \sum_{t=1}^{N} e^2(t,\theta) \tag{6.10}$$

which may be achieved by setting

$$\frac{\partial E}{\partial a_i} = 0$$

for $1 \le i \le n_a$, resulting in a set of n_a equations in n_a unknowns. Solving for the model coefficients and substituting the expressions back into (6.10) results in the Yule-Walker equations

$$\begin{bmatrix} R_0 & R_1 & R_2 & \ldots & R_{n_a} \\ R_1 & R_0 & R_1 & \ldots & R_{n_a-1} \\ R_2 & R_1 & R_0 & \ldots & R_{n_a-2} \\ \vdots & \vdots & \vdots & & \vdots \\ R_{n_a} & R_{n_a-1} & R_{n_a-2} & \ldots & R_0 \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \\ \vdots \\ a_{n_a} \end{bmatrix} = \begin{bmatrix} R_0 + \sum_{i=1}^{n_a} a_i R_i \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{6.11}$$

where R_τ is autocorrelation approximation

$$R_\tau = \frac{1}{N}\sum_{t=1}^{N} y(t)y(t-\tau)$$

The Yule-Walker equations (6.11) may easily be solved by standard Gaussian elimination or an optimised recursive technique by Levinson-Durbin (PARDEY *et al*, 1996).

The AR model may be conveniently used for spectral estimation. Taking the *z*-transform of (6.9) yields

$$E(z) = A(z)Y(z)$$

where

$$A(z) = 1 + \sum_{i=1}^{n_a} a_i z^{-i}$$

The transfer function of the AR model is thus

$$H(z) = \frac{Y(z)}{E(z)} = \frac{1}{1 + \sum_{i=1}^{n_a} a_i z^{-i}}$$

The frequency response is obtained by letting $z = e^{jwT}$ where T is the sampling period (taken to be unity throughout for simplicity). Assuming *E(z)* approximates white noise well, the spectrum will be flat and hence may be described by a constant $E(jw) = \varepsilon$. The estimate of the output spectrum is given by (assuming a sampling interval of unity)

$$P_y(w) = \frac{\varepsilon^2}{\left|1 + \sum_{i=1}^{n_a} a_i z^{-jwi}\right|^2} \qquad (6.12)$$

Equation (6.12) provides for the frequency domain interpretation of the AR model, namely an all-pole filter driven by flat-spectrum white noise. This representation fits well with the characterisation of the EEG as a noisy signal with certain preferred frequency bands. Moreover, the spectral estimates obtained by (6.12) are constructed using a continuous function of frequency and hence a significant improvement in frequency resolution over traditional periodogram methods may be achieved with this technique (PARDEY *et al*, 1996).

One of the earliest applications of AR modelling to the EEG was by ZETTERBERG (1969), incorporating it in an effective technique for obtaining neurophysiological relevant spectral parameters (see ISAKSSON and WENNBERG (1975) for a practical example). AR feature extraction was employed by PENNY *et al* (2000) to distinguish between motor imagery and a baseline, and motor imagery and a math task for use in a DBI. Typically, the AR model is appropriate when fitted to short, quasi-stationary segments of approximately 1 second in duration (PARDEY *et al*, 1996). PFURTSCHELLER *et al* (1996) analyse longer segments of EEG by using adaptive models to update the model parameters with the arrival of each new data sample. AR spectral estimation is also becoming popular for the reasons previously mentioned and was employed in a DBI by WOLPAW *et al* (2000).

6.4.2.2 ARX Model

The ARX model structure is obtained by including an exogenous input filtered by a transfer function with both numerator and denominator (AutoRegressive Moving Average) in (6.8). By assuming the noise enters early, the denominator polynomial can be assumed to be common, resulting in the ARX structure which is illustrated in Figure 6.6.

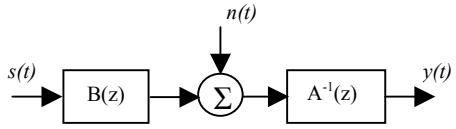

Figure 6.6: ARX model structure

The time domain equation for the ARX model is

$$\begin{aligned} y(t) = & -a_1 y(t-1) - \ldots - a_{n_a} y(t-n_a) \\ & + b_1 s(t-k) + \ldots + b_{n_b} s(t-k-n_b+1) \\ & + n(t) \end{aligned} \tag{6.13}$$

where s is the exogenous input and k is the delay. The predicted value of $y(t)$ may be written as

$$\hat{y}(t \mid \theta) = \theta^T \varphi(t)$$

where the column vector θ contains the model parameters

$$\theta^T = \begin{bmatrix} a_1 & a_2 & \ldots & a_{n_a} & b_1 & b_2 & \ldots & b_{n_b} \end{bmatrix}$$

and $\varphi(t)$ is the regression vector

$$\varphi(t) = \left[-y(t-1) \quad \ldots \quad -y(t-n_a) \quad s(t-k) \quad \ldots \quad s(t-k-n_b+1)\right]^T$$

The forward prediction error is

$$e(t,\theta) = y(t) - \theta^T \varphi(t) \tag{6.14}$$

and to fit the model to data we choose $\hat{\theta}$ such that it minimises (6.10). Inserting (6.14) into (6.10) results in

$$\begin{aligned} E(\theta) &= \frac{1}{N}\sum_{t=1}^{N}\left(y(t) - \theta^T \varphi(t)\right)^2 \\ &= \sum_{t=1}^{N} y^2(t) - 2\theta^T f + \theta^T R\theta \end{aligned} \tag{6.15}$$

where

$$f = \frac{1}{N}\sum_{t=1}^{N}\varphi(t)y(t)$$

$$R = \frac{1}{N}\sum_{t=1}^{N}\varphi(t)\varphi^T(t)$$

Assuming R is invertible, then (6.15) may be written as

$$\begin{aligned} E(\theta) = \frac{1}{N}\sum_{t=1}^{N} y^2(t) - f^T R^{-1} f \\ + \left(\theta - R^{-1}f\right)^T R\left(\theta - R^{-1}f\right) \end{aligned}$$

This expression is minimised when

$$\theta = \hat{\theta} = R^{-1}f$$

and provides an explicit equation for determining parameters which minimise the forward prediction error.

ARX models in the context of the EEG allow us to model a known deterministic signal such as an ensemble averaged evoked or event-related potential, linearly superimposed with a random background component (the ongoing EEG). CERUTTI *et al* (1988) employed the ARX model to filter Movement-Related Brain Macropotentials thereby drastically improving the signal to noise ratio of each single-trial for use in pathophysiological studies. The output of the filter was taken as the ARMA filtered exogenous input. In the present study, we make the previously unreported step in employing the ARX for feature extraction in a DBI. This technique extends the proven AR technique for EEG feature extraction (ZETTERBERG, 1969; PARDEY *et al*, 1996) by

including information pertaining to the known underlying endogenous or exogenously generated signal.

6.4.2.3 Model Order Estimation

We have thus far not discussed how to select model orders for the presented parametric models. Selecting model orders is a trade-off exercise between *bias* error and *variance* error. Bias error refers to inadequacies of the model to fit the data (even with an infinite amount of noise-free data) due to deficiencies in model structure. Variance error occurs when an experiment is repeated with the exact same input but different model parameters result. The output is not the exact same since the noise contributions are different each time. Using larger data sequences can often reduce variance error.

There are a number of methods for selecting optimum model order for parametric models in the literature. A popular method is the Akaike Final Prediction Error (FPE), which penalises higher order models and is given by (AKAIKE, 1969)

$$FPE(\theta) = \frac{N + n_a + n_b + 1}{N - n_a - n_b - 1} E(\theta)$$

(where $n_b = 0$ for the AR case). Of course, one should not lose sight of the real goal here of achieving a DBI with the best classification accuracy and the ultimate model order choice should be based on classification accuracies. Hence, we will use the FPE technique only as a guide to selecting model orders.

6.4.3 Comparison between ARX and AR features

In this section we compare the familiar AR feature extraction technique traditionally applied in direct brain interfacing to the ARX case where an evoked or event-related potential is employed as the exogenous input. A left/right self-paced typing exercise is analysed with six subjects. The ensemble averaged Bereitschaftspotential (an event related potential preceding the onset of movement) forms the exogenous input to the ARX model. The results show that the ARX case of modelling both the signal and noise was found to be considerably more effective than modelling the noise alone with the AR method yielding a classification accuracy of 52.8 ± 4.8% and the ARX method an accuracy of 79.1 ± 3.9 % across subjects. The results strongly suggest a role for ARX-based feature extraction in DBIs based on evoked and event-related potentials.

6.4.3.1 Experiment paradigm

A left/right self-paced typing exercise paradigm was employed with subjects sitting upright and fingers in the standard typing position at a Qwerty keyboard. The task consisted of pressing the left or right "home keys" with the corresponding fingers in a self-chosen order and timing. No feedback was given. The EEG was recorded using two bipolar leads with Grass P511 amplifiers (with 50 Hz line filter) from channels C3 and C4 of the 10-20 system, corresponding to the left and right primary motor areas respectively. The horizontal EOG from the right eye was also recorded. The signals were digitised at a sampling frequency of 512 Hz, low-pass filtered with a cutoff of 64 Hz and down-sampled to 128 samples per second. Six subjects took part in the trials (all male, aged between 22-25 years, in full health) in a single session per subject each lasting 10 minutes, resulting in a total of 60 minutes of data. Key presses were made on average every 2-3 seconds resulting in 200 - 300 epochs of 1500 ms length, ending 120 ms before the keystroke, thus avoiding effects of EMG activity masquerading as control signals. Trials where significant EOG activity took place (eye blinks) were omitted (< 2% of the collected data) from analysis (achieved automatically by linearly detrending and removing those time series whose maximum, rectified EOG amplitude exceeded a threshold).

6.4.3.2 Bereitschaftspotential

The experiment paradigm just described produces an event-related potential known as the Bereitschaftspotential (BP) - a gradually rising negative potential occurring about 1000 ms preceding the onset of movement (MISULIS, 1994). DEECKE *et al* (1969) identified an additional 3 components of the BP: a pronounced contralateral negativity over the precentral and parietal areas starting about 500 to 200 ms prior to movement, a small positive deflection beginning around 90 ms prior to movement, and a smaller negative potential starting about 50 ms prior to movement predominant over the primary motor cortex. Figure 6.7 illustrates the ensemble averages for C3 and C4 recordings for left and right self-paced movements from a single subject. Note the contralateral negativity commencing approximately 500 ms before the onset of movement at 0 s.

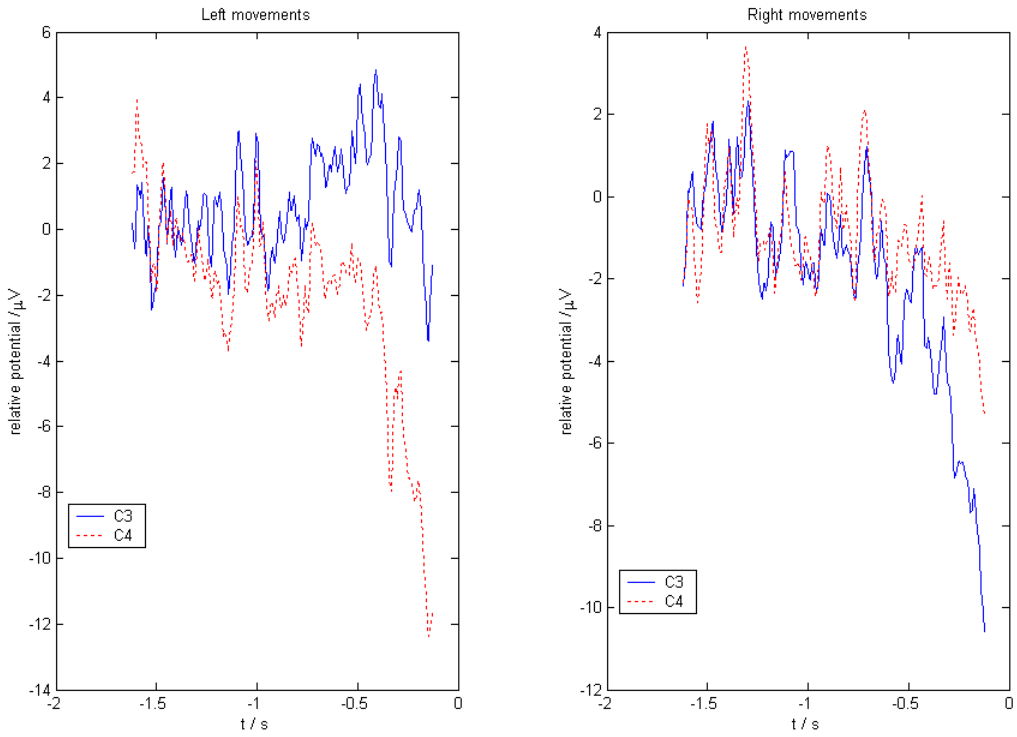

Figure 6.7: Ensemble average of left movements and right movements for electrodes C3 and C4 (movement occurs at 0 s). Note the contralateral negativity starting about 500 ms preceding movement

6.4.3.3 Results

To obtain an estimate of accuracy, a cross-validation procedure was employed (see Section 6.3.2). The trials for each session (consisting of parametric model coefficients) were first randomly shuffled and subsequently divided into N distinct segments. N-1 segments were used to train a LDA classifier and the remaining segment used as the test set. This process was repeated for each N possible test sets and the mean test set accuracy computed. Finally, the complete process with a new random shuffle was repeated M times to yield a mean accuracy and standard deviation. The data was uniformly processed with $M = 20$ and $N = 10$.

Table 6.1 displays the LDA classification accuracy for each subject in addition to the parametric model orders that yielded optimum accuracy for the ARX model (the Akaike FPE was used as a starting point for selection of model orders). In all subjects, the optimum classification accuracy was obtained for the same order n_a in both the ARX and AR models.

TABLE 6.1: PARAMETRIC MODEL ORDERS AND CLASSIFICATION ACCURACY

	n_a	n_b	ARX acc. %	AR acc. %
Subject 1	4	2	82.1	55.1
Subject 2	4	2	77.8	55.0
Subject 3	3	2	79.5	60.1
Subject 4	4	2	71.8	47.7
Subject 5	3	2	81.1	49.1
Subject 6	3	2	82.2	49.5
		Mean:	79.1	52.8

Figure 6.8 illustrates the mean accuracy and standard deviation over each shuffle for each subject for the AR and ARX models in the feature extraction stage.

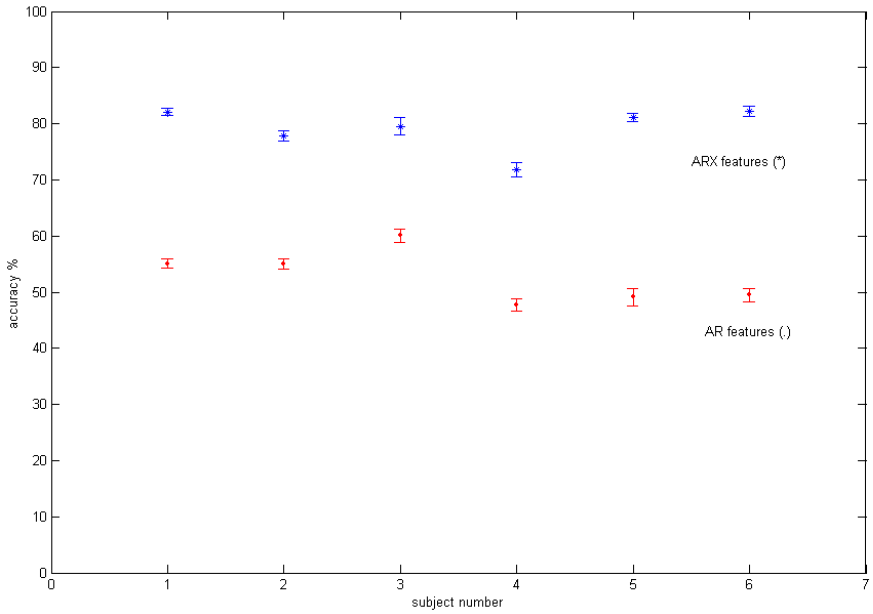

Figure 6.8: Mean accuracy % over $M = 20$ random shuffles for each subject for AR and ARX features

The mean accuracy over all subjects for ARX features is 79.1 ± 3.9% whilst for AR features is 52.8 ± 4.8%. The results show clearly that for each subject the ARX set of features yields a higher classification accuracy than the AR features.

6.4.3.4 Discussion

The common denominator structure of the ARX model, implying that noise enters early in the modelled process seems physiologically reasonable as ongoing EEG contributions from neighbouring neural populations to those firing synchronously contribute mostly to the background noise (MISULIS, 1994). Modelling the contributions of both the known signal (ensemble averaged evoked or event-related potential) and the noise (EEG) yields better features for the classifier and is an intuitively satisfying result.

The experimental paradigm employs short, quasi-stationary epochs of 1.5 seconds for analysis terminating 120 ms before the onset of movement. It is quite likely that longer epochs, for example those used in PFURTSCHELLER *et al* (1996), would yield higher accuracies for the AR model by also extracting pertinent features of event-related desynchronisation. The results suggest that the ARX feature extraction approach should prove useful for DBI systems based on evoked and event-related potentials e.g. SUTTER (1992), FARWELL and DONCHIN (1988). The model orders in (6.13) for the b_i coefficients are larger for more complex morphological waveforms such as the P300 (as demonstrated in FINUCANE *et al* (2003)).

Parametric feature extraction and classification stages are simple to implement, computationally efficient, and thus are suitable for implementation in real-time for DBI systems. With any DBI, speed and accuracy are important metrics for comparison with generally an inverse relationship between them. Evoked and event-related potentials can be elicited on a suitably short time scale and with sufficient classification accuracy, are well suited to the task at hand. The choice of employing an LDA was prompted by the emphasis of the research being on the extraction of pertinent features rather than on the classifier itself, nevertheless LDA classifiers are a natural choice providing computationally simple discriminant functions that minimise the error of misclassification. While application of higher order classification strategies such as Quadratic Discriminant Analysis (assumes non-uniform covariance matrices) did not improve classification results, it is possible that selective feature approaches such as Linear Selective Stepwise Analysis might yield better results (CHILDERS *et al*, 1987).

6.5 *SIDEBAR*: An Eye-Blink Interface

Most EEG-based DBIs harness electrical activity originating from somatosensory, motor, or visual areas of the cortex. These areas may be severely damaged in people with stroke or neurogenerative disease and hence other methods of communication or control may be sought. In this sidebar, we consider an alternative interface for communication and control based on image processing of blinking action (BURKE *et al*, 2001). Sophisticated mathematical algorithms (such as differencing, thresholding, aggregation and statistical analysis of skin colours) are used to compare successive

frames of computer-captured images of the face. From these, changes in state of the eyes are determined and are used to detect blinks. A recognition performance of 83.74 ± 0.03% is achieved over five subjects. A logical decision rule identifies purposeful blinks and applies them to control either a custom-designed communication package or an external device.

6.5.1 Background

There are non-contact systems available to detect eye gaze based on infrared, e.g. the Eyegaze System, CLEVELAND and DOYLE (1992). While effective, these systems require specialist equipment and suffer from the requirement that the user must calibrate the system, which necessitates some level of cognition to be first established with disabled patients. In contrast to marker-based systems such as MIYAZAKI *et al* (2000), the philosophy adopted here was to develop an entirely non-contact, marker-free, solution, using an inexpensive digital camera ('a webcam'), a desktop PC (Pentium III, 500 MHz), and a judicious choice of sophisticated techniques from digital image processing. Digital image processing is a very wide field but for present purposes can be regarded as a collection of mathematical algorithms applied to computer stored visual images in order to extract features, which can be used to accomplish some desirable end. By dispensing with the requirement of markers, the system presents itself as unobtrusive and simple to use to both therapist and patient alike. Based on a development of ideas described in CLARKE *et al* (1998), the mathematical techniques are summarised in the next section.

6.5.2 Methods

The basis of our blink detection method invokes the classic Euclidean distance formula to measure the *colour difference* between successive frames of images. Each pixel has a colour associated with it represented by its discrete *tristimulus* values (red (*R*), green (*G*), and blue (*B*) parameters corresponding to intensities ranging from 0 to 255). The colour distance for a pixel *(i, j)* between frame *n* and *n-1* is calculated from

$$d_{ij}[n] = \sqrt{\left(R_{ij}[n] - R_{ij}[n-1]\right)^2 + \left(G_{ij}[n] - G_{ij}[n-1]\right)^2 + \left(B_{ij}[n] - B_{ij}[n-1]\right)^2}$$

By applying a threshold to the above equation, we can extract only those pixels that have changed significantly in intensity between two successive frames, i.e. if $d_{ij}[n] >$

$d_{threhold}$ then set pixel *(i,j)* colour to blue say, otherwise set to white. Figure 6.9 illustrates the effect for a blink.

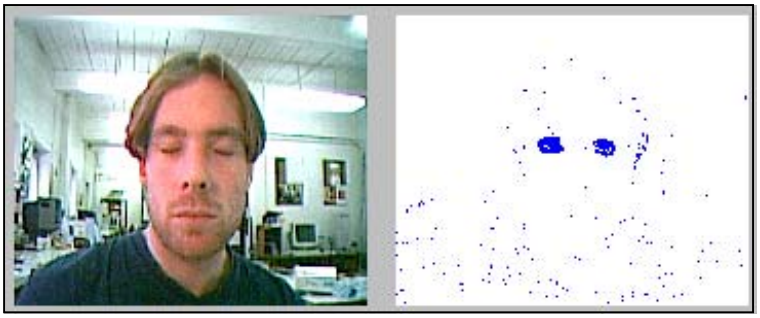

Figure 6.9: Detection of blinks by the colour distance method

Next, we determine the two aggregations of pixels having the largest areas. We accomplish this by using an algorithm that traces the perimeter of all groups of adjacent pixels in the frame and reports the respective areas. Eyes are classified as detected when the horizontal and vertical separations between the largest two aggregations are verified to be between obvious anatomical bounds (assuming an approximate vertical orientation of the head).

Although the techniques presented thus far can be used to detect eye blinks, they still require some refinement to remove false detections due to saccadic movements of the eyes (or in the case of one patient, to filter the effect of nystagmus). To differentiate eye movements from actual eye blinks, we make use of the fact that on average pixels representing closed eyes will be predominately flesh coloured. On detection of the first blink (or a saccadic eye movement as the case might be), we sample a rectangular area surrounding the eyes (proportional to the inter-eye separation) to generate a statistical histogram of flesh colours. To simplify computation, and remove brightness information, we normalise the tristimulus values, producing the *trichromatic coefficients* GONALES and WINTZ (1987)

$$r = \frac{R}{R+G+B}$$

$$g = \frac{G}{R+G+B}$$

$$b = \frac{B}{R+G+B}$$

Since

$$r + g + b = 1$$

we can now evidently describe *chromaticity* by just two values e.g. *r* and *g*. Figure 6.10 illustrates an example of a histogram of flesh colours sampled in an area surrounding the eyes.

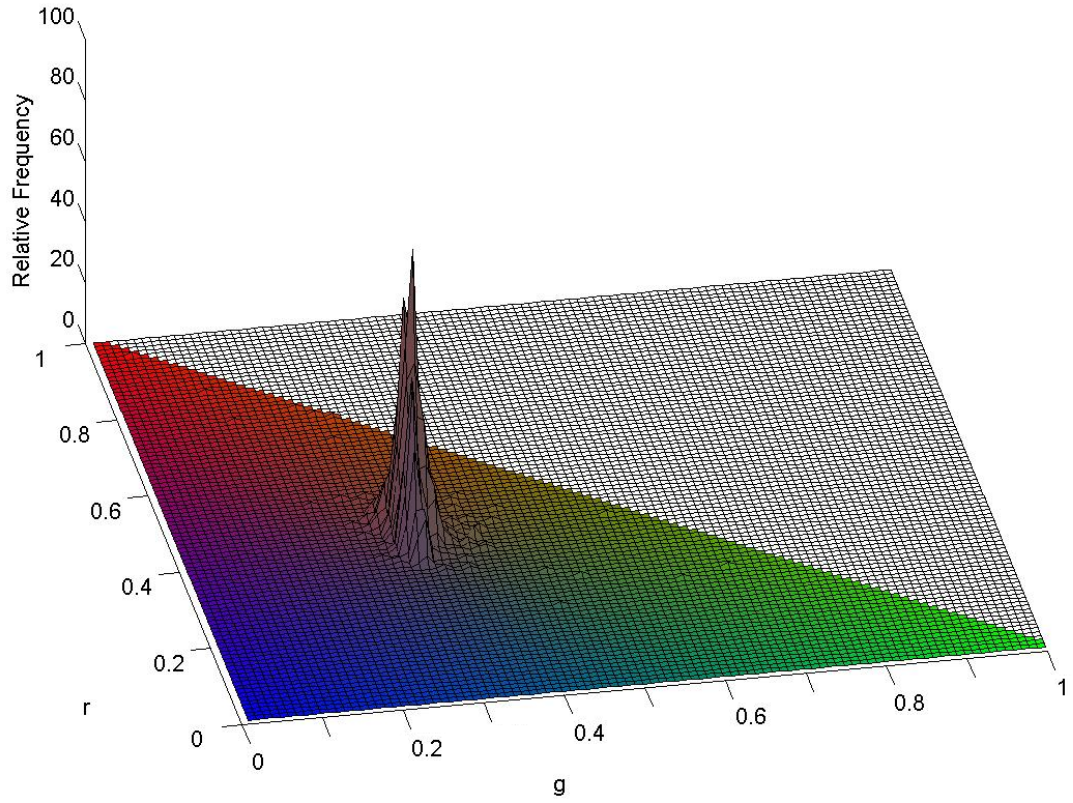

Figure 6.10: Distribution of the trichromatic coefficients of a flesh sample in an area surrounding the eyes

The brightness has been normalised and we are now viewing just the chromatic (*hue* and *saturation*) information. Once the histogram has been generated, the trichromatic coefficients of pixels from subsequent frames are compared to the stored distribution and the relative frequency of each measured chromaticity is determined. Frequencies above a threshold are characterised as flesh colour. Figure 6.11 illustrates the flesh colour filtering mechanism (flesh colours are shown in blue). With successive blinks, we compute the average probability that the detected pair of pixel aggregates is flesh coloured (i.e. eyelids are closed) and probabilities above a certain value are classed as true blinks.

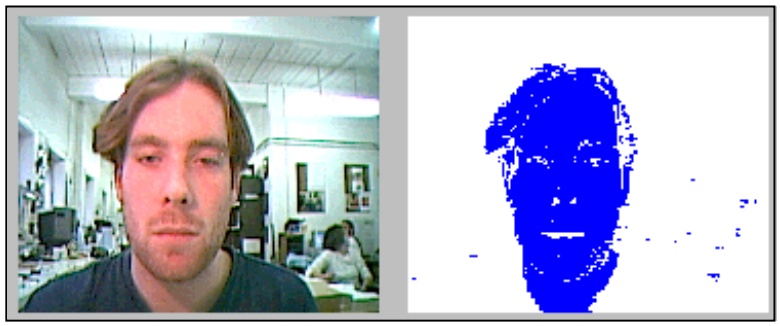

Figure 6.11: Detection of flesh colours

6.5.3 Results

Five volunteer subjects (2 male, 3 female, aged between 21 and 25) each performed a set of five experimental blocks lasting two minutes each. This translates to 50 minutes (5*x*5*x*2) of analysis. During each experimental block, subjects were placed in a seated position 0.5 metres from the digital camera and asked to look in the direction of the camera, while remaining reasonably still. A second camera recorded the subject's face and its output was routed to a video/audio capture board in a separate PC. Each detected blink resulted in an audio track to toggle between play and pause and this audio was recorded simultaneously by the capture board. The captured audio and video were then manually analysed offline and the fraction of true positives and false positives determined for each experimental block for each subject (Table 6.2).

TABLE 6.2: ACCURACIES (A: ACTUAL BLINKS; D: DETECTED BLINKS; F: FALSE POSITIVES)

	Trial 1			**Trial 2**			**Trial 3**			**Trial 4**			**Trial 5**		
	A	D	F	A	D	F	A	D	F	A	D	F	A	D	F
Subject 1	54	48	0	48	39	2	48	42	0	45	38	1	57	52	1
Subject 2	25	20	0	13	10	0	11	10	1	40	35	0	32	27	0
Subject 3	11	9	1	23	19	2	16	15	0	12	8	0	21	19	0
Subject 4	42	34	0	29	22	2	49	38	1	17	15	0	23	21	2
Subject 5	33	27	0	34	27	2	33	26	0	36	30	0	33	29	3

Figure 6.12 presents the results (mean and standard deviation) for each subject. Performance averaged over all subjects yields 83.74 ± 0.03% for true positives and 2.71 ± 0.01% for false positives.

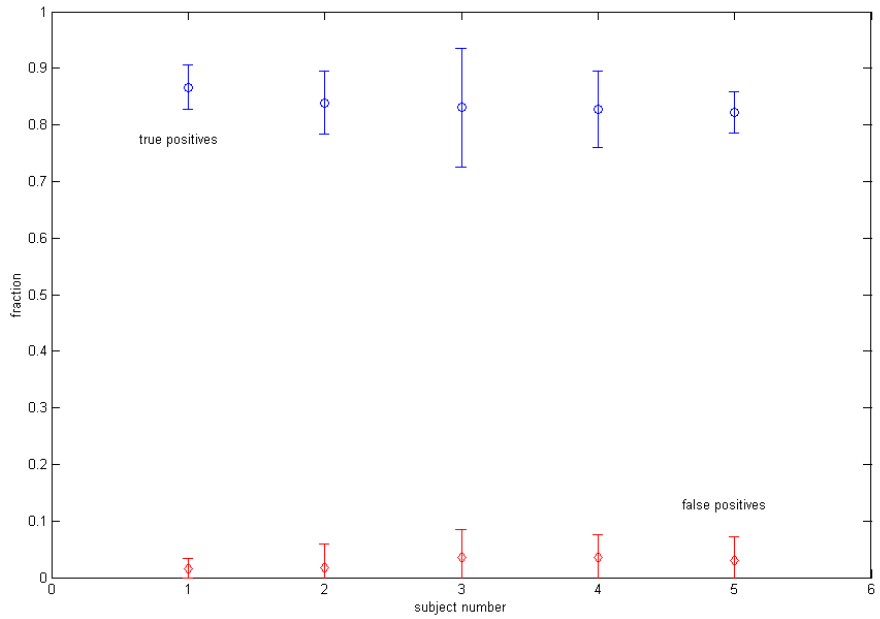

Figure 6.12: Fraction of true positives and false positives for each subject

6.5.4 Discussion

The preliminary results on the use of real-time digital image processing in the implementation of a non-contact human-computer interface are promising, indicating that this type of technology is sufficiently robust to be used in rehabilitation applications. A component architecture has been developed where the blink detection system runs in a separate process and inter process communication (IPC) allows different applications to be effectively 'plugged-in'. The first application uses purposeful blinks to control the switching action of a CD player, common in the average personal computer. The second application utilises blinks to control a custom designed communication program, allowing users to build up sentences letter by letter, and produce synthesised speech. With each application, an adaptable logical decision rule operates allowing different blink rates to trigger the corresponding control action. Only the on-response is detected since this allows us to filter eye movements, which can sometimes give erroneous results. Where large head movements occur, the system does not attempt to detect blinks (since determining the position of the eyes becomes problematic) although this has not been an issue with the type of patients the system was built for. An improvement in the blink detection algorithms might account for large head movements by tracking the position of the eyes (by first locating the face by the flesh colour filtering mechanism) and performing the difference operation on the pixels marked as representing the eyes.

6.6 fMRI Direct Brain Interface

Functional Magnetic Resonance Imaging (fMRI) is a recent imaging technique for determining which parts of the brain are activated in response to a stimulus or event (KWONG *et al*, 1992). In contrast to techniques such as the EEG, fMRI does not measure electrical activity of the brain but rather changes in local blood flow, volume, and oxygenation. These *haemodynamic* changes typically occur within a few seconds of neural activity (COHEN and BOOKHEIMER, 1994). There are already ample results in the literature that demonstrate the ability to detect phenomena conducive to building DBIs including, for example, imagined movement (KOSSLYN *et al*, 2001), steady-state visual responses (KASHIKURA *et al*, 2002), oddball paradigms and P300-related responses (HOROVITZ *et al*, 2002; MCCARTHY *et al*, 1997). Detecting haemodynamic phenomena is analogous to detecting evoked and event-related potentials due to the poor signal to noise ratio (WEISSKOFF *et al*, 1993) and in this section we apply the ARX feature extraction and LDA classification scheme to the problem of single-trial event-related fMRI detection.

The primary advantage of fMRI over EEG is in its millimetre spatial resolution and accuracy (COHEN and BOOKHEIMER, 1994). Functional MRI, like the EEG, is a non-invasive technique but suffers from the disadvantage that it operates on a fundamentally longer time-scale and requires complex equipment to produce data, in particularly the requirement of a large magnetic field to operate the imaging mechanism. Nevertheless, work is underway to combine the temporal resolution of EEG with the spatial accuracy of fMRI (HUANG-HELLINGER *et al*, 1995) and this may one day pave the way for the ultimate data source for DBIs. For completeness, we outline briefly the mechanisms underlying MRI and fMRI in the next section. A study of single-trial detection of event-related fMRI at 1.5 Tesla is presented that results in the ability to classify single-trial events with an accuracy of 69.3 ± 0.1% using the raw data samples as features and with an accuracy of 73.3 ± 2.0% when employing the ARX feature extraction strategy.

6.6.1 MRI and fMRI Basics

MRI is a complex technique and we only describe the basics here such that the results in the subsequent sections can be understood – the interested reader is referred to

ROSEN *et al* (1998), COHEN and BOOKHEIMER (1994), and KWONG *et al* (1992) for good starting points to the literature.

6.6.1.1 Nuclear Magnetic Resonance

MRI is based on the physics of Nuclear Magnetic Resonance (NMR). Protons such as those found in the hydrogen nucleus in water are slightly magnetic. Under the influence of a large magnetic field $\vec{B}$, there is a tendency for the protons to align themselves with the field. Magnetisation density refers to the amount per unit volume that an object in $\vec{B}$ is magnetised at any given location and is denoted by $\vec{M}$. After a short period of time, the $\vec{B}$ and $\vec{M}$ vectors will become aligned and the condition is called *fully relaxed magnetisation*. Application of radio frequency (RF) radiation at the correct resonance frequency will disturb the magnetisation. If the RF signal is removed, the protons relax back to align with the external field. However, the resultant path of the vector $\vec{M}$ is not simple, rather the phase-synchronised angular momentum of the nuclei (spin) will cause the resultant vector to *precess* (rotate) about $\vec{B}$ and emit RF signals as it does so (see Figure 6.13). The emitted RF may be detected by conventional means, its frequency being dependent on the strength of the magnetic field at the location of its origin. By employing orthogonal gradients in the externally applied field $\vec{B}$ and detecting the emitted RF over a range of frequencies, an image of RF intensity over an object may be obtained.

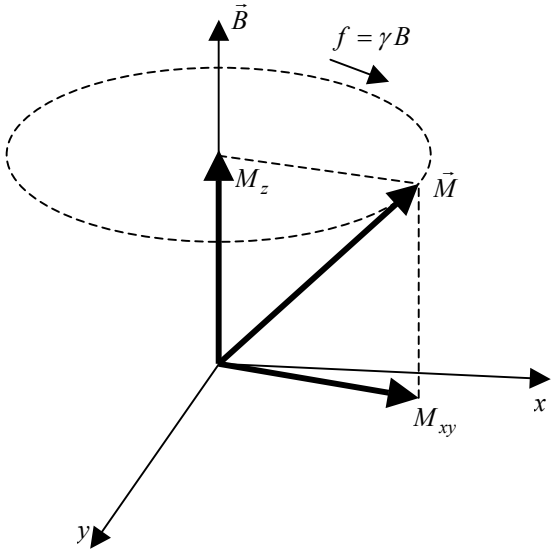

Figure 6.13: When disturbed from the relaxed state, $\vec{M}$ precesses about $\vec{B}$ in a clockwise direction at the Larmor frequency, $f = \gamma B$, while emitting RF radiation. The precession rate is much larger than the rate at which $\vec{M}$ relaxes back to the $\vec{B}$ direction

Different tissues and structures display different characteristic intensity relaxations and thus static images consisting of many volume elements called *voxels* may be obtained.

6.6.1.2 fMRI Principles

Functional MRI is concerned with obtaining time course information of NMR signals from voxels of interest by acquiring a series of rapid MRI images under different task conditions. The signal changes of interest (contrast) are those occurring as a result of neural activity to blood flow, volume, and oxygenation and are collectively known as haemodynamic effects. It should be noted that fMRI is a complex, indirect measure of neural activity (neural activity → blood flow → NMR signal) and caution should be applied when interpreting the resultant signal. Haemodynamic changes are considerably slower than EEG measures, for example changes related to oxygenation take 4-5 seconds to reach a peak, another 5 seconds to return to baseline, and are followed by an undershoot for a subsequent 10 seconds (BOYNTON *et al*, 1996).

There are four specific haemodynamic signals that influence the NMR signal (the exact physiological mechanisms of how neural activity signals haemodynamic changes is yet to be elucidated) including:

- increased blood flow (NMR signal ↑)
- oversupply of oxyhaemoglobin (NMR signal ↑)
- increased blood volume (NMR signal ↓)
- increased capillary flow (NMR signal ↑)

Local changes in arterial blood flow into a neurally active region will influence the NMR signal because of an increase in supply of water into the volume. These changes occur proximal (upstream) to the true site of activation and are highly dependent on the arterial supply to the region in question. Oxygen supply to capillaries does not remain constant during neural activation and in fact an over-supply of oxyhaemoglobin occurs. This is detectable distal (downstream) to the active region and results in an increase in the NMR signal. This phenomenon is known as the BOLD (Blood Oxygen Level Dependence) effect and is the basis for the majority (including ours) of fMRI studies (COHEN and BOOKHEIMER, 1994). Increased blood volume in response to neural activation occurs as a result of a ballooning effect of veins distal to the site of activation. After passage through capillaries, oxygen is removed resulting in a build up of deoxyhaemoglobin, which reduces the resulting NMR signal. Finally,

increased blood flow can be detected in capillaries by virtue of the fact that most water molecules that flow into capillaries do not flow out at the other end. By labelling the water molecules (using proton magnetisation), it is possible to detect flow into capillaries. This method is not yet widespread but gives better spatial localisation than the BOLD technique (ROSEN *et al*, 1998).

The fMRI BOLD response may be modelled by the γ-variate model (COHEN, 1997; GARAVAN *et al*, 1999) given by

$$y = kt^r \exp\left(\frac{-t}{b}\right) \tag{6.16}$$

Figure 6.14 illustrates a modelled response with realistic parameter values (GARAVAN *et al*, 1999) as per the figure caption.

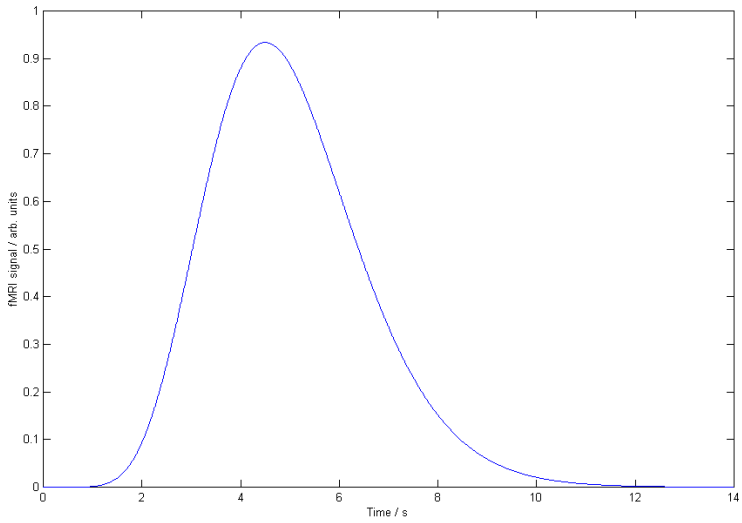

Figure 6.14: γ-variate function model of a typical haemodynamic response. Parameter values: $k = 0.01$, $r = 9$, $b = 0.5$

COHEN (1997) showed that the form of fMRI response to stimuli of freely varied timing is satisfactorily modelled by convolution of the impulse response (6.16) and the behavioural stimulus.

As mentioned at the outset, detecting responses of the form of (6.16) is difficult due to the low signal to noise ratio. At 1.5 Tesla, the changes that occur to the NMR signal during activation are only about 2-5% the amplitude of the background signal (COHEN and BOOKHEIMER, 1994). Magnetic field inhomogeneities can cause fMRI images to contain artefacts and 'ghosts' – faint duplicates of portions of an image displaced to

incorrect locations. Aside from the inhomogeneities inherent in construction of the scanner's electromagnets (although these can be somewhat calibrated electronically in a process known as *shimming*), placing an object inside the magnet's bore will alter the field due to the material's magnetic susceptibility. Gross motion of the subject being scanned also affects the resultant NMR signal, with small motions affecting the signal greatly. Motion outside of the scanned region e.g. swallowing or speaking will also affect the NMR signal due to magnetic susceptibility altering the magnetic field within the brain. Other physiological motions e.g. breathing will have similar effects.

6.6.2 Single-Trial Detection of Event-Related fMRI

In this section we describe preliminary results that test the ability of the ARX feature extraction and LDA classification strategy to detect single-trial fMRI responses. The central assumption is that the various noise sources are uncorrelated with the fMRI response (see Section 6.6.1.2) and may be modelled as a white noise source filtered by an autoregressive model. An ARMA structure models trial-by-trial variability and the combined structure (assuming a common denominator polynomial) results in the ARX model (6.13) as before. A comparison of classification accuracy of using the raw data samples and the ARX features is made to verify its validity.

6.6.2.1 Methods

The experiment paradigm was that of GARAVAN *et al* (1999) and employed the same data. Subjects were presented with a stream of letters serially every 500 msec and instructed to make a response whenever certain target letters (X or Y) were presented. An additional rule stipulated that responses must alternate between the targets, such that if the most recent target was, for example, the letter X, then subjects should only respond to the next letter Y in the letter stream. GARAVAN *et al* (1999) identified regions (strongly lateralised to the right hemisphere), which were activated by subjects correctly withholding a motor response. In our experiment, 12 voxels from a single subject were identified for showing a strong response (6 voxels) and no response (6 voxels) respectively by performing a goodness of fit to the ideal haemodynamic shape (6.16). Three sessions (each 136 seconds duration) were performed for the subject and fMRI events were separated by an average of 20 seconds to reduce the possible effects of overlapping responses. A haemodynamic template was obtained by ensemble averaging over the events from the active voxels per session using a time window of

20 seconds. This template forms the exogenous signal to the ARX model. Model orders were selected on the basis of yielding optimum classification accuracy. To avoid problems with ill-conditioned matrices as a result of the coarse sampling period (2 seconds), the fMRI data was resampled at 5 times its original rate using low-pass interpolation.

6.6.2.2 Results

To obtain an estimate of accuracy, a cross-validation procedure was employed. The data was uniformly processed with $M = 30$ and $N = 15$ (see Section 6.3.2). Table 6.3 displays the LDA classification accuracy for the raw data samples and for the ARX features and its corresponding model orders (chosen for yielding optimal classification accuracy).

TABLE 6.3: MODEL ORDERS AND CLASSIFICATION ACCURACY

	n_a	n_b	ARX acc. %	Raw acc. %
Session 1	7	4	71.0	69.3
Session 2	7	3	74.2	69.3
Session 3	7	3	74.7	69.2
		Mean:	73.3	69.3

Figure 6.15 illustrates the mean accuracy and standard deviation over each shuffle for each session for the ARX model in the feature extraction stage and for using the raw samples as features.

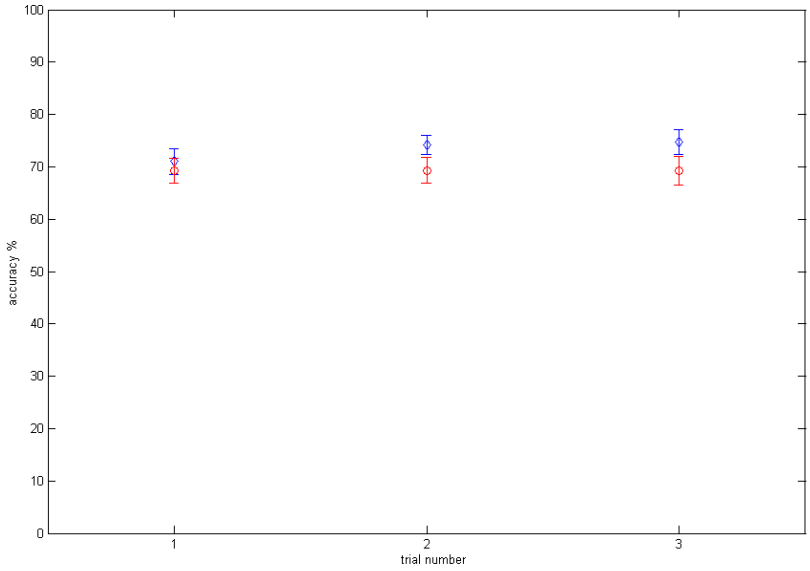

Figure 6.15: Classification accuracies for ARX features (diamonds) and raw samples as features (circles) over three sessions

The mean accuracy over all three sessions for ARX features is 73.3 ± 2.0% while for the raw samples as features is 69.3 ± 0.1%. In each session, the ARX set of features yielded higher classification accuracy than employing the raw data samples as features.

6.6.2.3 Discussion

Early fMRI experiment paradigms were based upon those used in Positron Emission Tomography (PET), i.e. employing a block design consisting of extended periods of "on" versus "off" activations (ROSEN *et al*, 1998). While radio-pharmaceutical based imaging such as PET mandate such paradigms to achieve quasi-equilibrium physiological states for about 1 minute, it has been shown that stimuli as short as 34 ms can elicit fMRI responses (ROSEN *et al*, 1998) and has led to the event-related fMRI paradigm. Event-related fMRI allows investigators to study new tasks not possible with the conventional block design e.g. unpredictable responses such as inhibition (GARAVAN *et al*, 1999), or the effect of novelty as in the oddball paradigm (MC CARTHY *et al*, 1997). Typically, event-related fMRI data is studied by fitting responses (often selectively averaged over many trials to increase signal to noise ratio) to a model and employing statistical tests (e.g. GARAVAN *et al*, 1999; JOSEPHS *et al*, 1997). A natural extension of this is single-trial detection of event-related fMRI responses analogous to event-related or evoked potential detection in the EEG. Accessing single-trial information in fMRI may:

- allow for Direct Brain Interfacing,
- facilitate new adaptive experiment designs,
- enable integration with other single-trial electrophysiological methods
- enable biofeedback applications,
- provide an objective measure of mental performance, and
- enable lie detection mechanisms.

The results for single-trial detection of responses at 1.5 Tesla suggest a role for parametric classification strategies and also parametric models for feature extraction in singe-trial event-related fMRI. Since the signal to noise ratio scales with magnetic strength (COHEN and BOOKHEIMER, 1994), improved results are expected at higher magnetic field strengths. While our main concern here has been detection of single-trial event-related responses, information about the unique aspects of an isolated event

might also be obtained by using the ARX model as a filter as in the original publication applied to EEG event-related potentials by CERUTTI *et al* (1988).

Creating an effective Direct Brain Interface requires a high information transfer rate and often a trade off between speed and accuracy. The temporal resolution of fMRI, fundamentally limited by the haemodynamic mechanisms that produce it, places an upper limit on the communication rate achievable by this method for DBI applications. However, as mentioned earlier, combining the spatial advantages of fMRI with the temporal advantages of simultaneous EEG measurement may soon yield a new intriguing source of data for DBIs. Our method goes some way to facilitating a tighter integration between fMRI and more traditional single-trial electrophysiological methods. In practice MRI scanners, requiring large magnetic fields to operate, have a long way to go before being able to compete with their electrophysiological counterpart methods in terms of portability but consequences of work such as HALPERIN *et al* (2001) on high-temperature superconductors may one day lead to practical, more cost effective scanners that could be used as part of a DBI solution.

6.7 Summary

Assuming a stochastic characterisation of the EEG, the AR model for feature extraction was extended to the ARX case for use in evoked and event-related potential based DBIs. The ARX model characterises the signal into deterministic (evoked or event-related potential) and stochastic (EEG) components. The increase in accuracy of the ARX method over the AR method is significant, and was verified over six subjects. This method has also recently proved useful for P300-based DBIs (FINUCANE *et al*, 2003). The fMRI signal was harnessed using the same paradigm, thus exploiting the analogous problem of single-trial haemodynamic response extraction. This technique should prove useful in combining fMRI and electrophysiological data where the spatial resolution benefits of the former complement the temporal resolution of the latter. As an aside, an alternative human-computer interface based on blink detection was presented. This non-contact interface may be appropriate for disabled people who have sustained damage to sensorimotor and visual cortical areas.

CHAPTER VII
CONCLUSION

Driven by the goal to create practical, direct brain interface applications, this thesis has sought to explore and characterise the EEG signal using a substrate of dynamical systems theory as the primary technical tool. The principal observation of this treatise is that macroscopic brain electrical activity manifests itself *not* as a system exhibiting chaotic dynamics but rather as a nonlinear, stochastic dynamical system. The body of work encompasses material ranging from system theoretic discussions and novel mathematical constructions through to successful practical applications influenced by the preceding work. As the work evolved, new avenues of study presented themselves in almost every major topic studied. The chapters are structured in such a way that, rather than conclude each with a short summary, the final section of the chapter discusses areas which the author feels deserves further study. We bring the thesis to a close here by providing a summary for each chapter in turn, with particular emphasis on possibilities for future study.

A novel method for prescribing deterministic limit cycle oscillators for given geometric planar curves is presented in Chapter II. La Salle's extension to Lyapunov's Direct Method (LA SALLE and LEFSCHETZ, 1961) provides the vehicle for proving asymptotic stability of the limit cycles. This construction is quite powerful, facilitating one to create oscillators with arbitrary prescribed geometries and is expected to have application in areas including electronic signal generation, modelling, and education. For the purposes of this thesis, it provides a solid basis for developing upon in later chapters. A second method for constructing oscillators is also presented that allows one to synthesise certain classes of periodic signals by using a derivative coordinate embedding. Initial results are promising but the current method introduces stiff equations. The author feels that further research is warranted to alleviate the disparate time scales resulting in the current proposal thus aiding practical implementation of the oscillators.

In Chapter III, the prescribed oscillator technique is extended using Itô calculus to produce *stochastic* limit cycle oscillators. Although dating back to the 1940s, Itô

stochastic differential equations have only recently started to gain popularity in the applications literature. In particular, biomedical studies involving such constructs are rare. To perform simulation studies, a new extensible software toolbox for solving stochastic differential equations has been created. As one might expect, stochastic stability theory is much less well developed than its deterministic counterpart but nevertheless a stochastic version of Lyapunov's Direct method does exist (HAS'MINSKII, 1980). Armed with this, we produce a definition of the prescribed stochastic limit cycle as an invariant set. A stochastic limit cycle oscillator model of the EEG is constructed and is demonstrated via numerical simulation to reliably reproduce the visually displayed patterns of the EEG, its spectra, and amplitude characteristics. The stochastic limit cycle oscillator model is suggested as an alternative mechanism to chaos to model the aperiodic behaviour seen in the actual EEG. The popular Wilson-Cowan model of interacting excitatory and inhibitory neurons is also extended to a stochastic differential form. An enviable characteristic trait of a chaotic system is its sensitivity to initial conditions, suggested by Freeman to be beneficial in the brain, for example to enable small microscopic inputs from sensory cells to elicit large-scale macroscopic responses across sensory cortex regions. The final section in Chapter III illustrates, in the context of dynamical systems, how this phenomenon might be achieved by employing an alternative mechanism. A stochastic dynamical system at a bifurcation point is illustrated which can display qualitatively different behaviour for a minimal expenditure of energy. This is somewhat similar to a theorem by HOPPENSTEADT and IZHIKEVICH (1997), whose modelling work is based on the concept that only those neurons participate nontrivially in the brain processes whose dynamics are at a bifurcation point. The study of different kinds of bifurcations (e.g. cusp, pitchfork, saddle-node etc) under the influence of stochastic 'blurring' appears to warrant further research, in particular its application to biological systems where noise is ubiquitous.

In Chapter IV, the focus is switched to analysing and modelling time series data. The goal here is to study the EEG time series using the tools of nonlinear time series analysis. In particular, we show that the stochastic limit cycle model of the EEG produces very similar dynamic characteristics to the actual EEG. The preliminary results of CASDAGLI (1991) are extended to *qualitatively different* EEG time series. For EEG originating from the eyes-closed relaxed state and deep sleep, little or no

evidence for nonlinearity is present in the data itself. In stark contrast, data obtained from a petit mal seizure appears quite nonlinear. These results were related to existing literature by employing the correlation dimension in a statistical framework involving phase-randomised surrogate data. The results of the correlation dimension test confirmed those of the Casdagli technique for each (qualitatively different) EEG data set tested. Next, the stochastic limit cycle model of the EEG is analysed and shows identical characteristics to actual (non-pathological) EEG data. This is particularly important, as the nonlinear mechanisms underlying the dynamics of the model do *not* manifest themselves in nonlinear time series analysis. Thus, while nonlinear dynamics almost certainly underlie the genesis of the actual EEG, nonlinear analysis techniques do not appear warranted by the data itself, with the possible exception of pathological data. A second order stochastic differential equation based on the Ornstein-Uhlenbeck equation is simulated and processed through a static nonlinearity. Locally linear (hence globally nonlinear) modelling is shown to provide the optimal predictor for the data. It is quite possible that the nonlinearity evident in pathological EEG data might be ascribed to this kind of non-dynamical nonlinearity (indeed this is backed up by HERNÁNDEZ *et al* (1996) who conclude that petit mal EEG data is best described by stochastic disturbed limit cycle behaviour). A further avenue of research here is to investigate under what conditions our stochastic limit cycle model exhibits increased nonlinearity in the data. Finally, we briefly review a recent technique by SIEGERT *et al* (1998) for determining the drift and diffusion terms for an underlying stochastic differential equation from the data itself. There are two developments of this technique that come to mind. Firstly, this technique suggests an intriguing way of parameterising a stochastic limit cycle oscillator similar to the one presented in Chapter III. Secondly, the parameter values themselves might well be used as features in pattern recognition applications.

Chapter V is concerned with modelling evoked potentials on the basis of stimulus-triggered phase reordering consisting of both phase synchronisation and phase resetting of individual oscillators. Stimulus-triggered synchronisation is shown to occur naturally in a network of biologically realistic model neurons. Next, a new macroscopic model for the generation of evoked potentials is presented based on the transient synchronisation of coupled oscillators representing neural populations. The model predicts latency under anaesthesia and generates a response that is shown to

approximate the analytic form of the Morlet wavelet – suggesting its application in the analysis of evoked potentials as an alternative to Fourier techniques (which employ sinusoidal basis functions). The chapter is concluded with a brief demonstration of stochastic synchronisation. PECORA and CARROLL (1990) recently observed that certain subsystems of nonlinear chaotic systems could be made to synchronise. We extend this observation to the case of stochastic limit cycle oscillators. Since both synchronisation and noise are ubiquitous in the brain, it appears that this phenomenon might well deserve further study.

The penultimate chapter is concerned with practical applications. Based on the results of Chapter IV, one may conclude that the AutoRegressive (AR) model might be the most appropriate approach to characterising the EEG for pattern recognition stages in direct brain interface systems. We extend the AR method for feature extraction to the ARX case thus facilitating the reduction of data to a set of coefficients for paradigms involving ensemble averaged evoked or event-related potentials as exogenous inputs. An exploration of the technique is performed on a number of subjects with promising results. As a sidebar, an alternative communication mechanism is presented for use with locked-in patients (developed during the work of this thesis (BURKE *et al*, 2001)). An EEG based direct brain interface may not always be possible (e.g. if there is damage to the relevant cortical areas) and this technique, based on blink detection, can give comparable and sometimes better results. Chapter VI is concluded with a study of a direct brain interface based on fMRI data. Again the results are promising, and suggest a novel way of performing tighter integration with 'traditional' electrophysiological techniques. Indeed, combining the spatial advantages of fMRI with the temporal advantages of simultaneous EEG measurement may soon yield a new intriguing source of data for direct brain interfacing.

Dynamical system modelling and analysis techniques have been shown to be powerful tools for studying electrical activity originating from the brain. In particular, this treatise suggests that macroscopic brain electrical activity is optimally modelled by the dynamics described by stochastic differential equations. This formulation, where the noise influences actually interact with the dynamics is substantially different to the usual definition of measurement noise and ought to lead to new ways of modelling other electrophysiological data.

REFERENCES

ABARBANEL, H. D. I. (1996): 'Analysis of Observed Chaotic Data', (Springer, New York, London)

ABARBANEL, H. D. I., BROWN, R., SIDOROWICH, J. J., and TSIMRING, L. S. (1993): 'The analysis of observed chaotic data in physical systems', *Reviews of Modern Physics*, 65, pp. 1331 - 1392

ABARBANEL, H. D. I., HUERTA, R., RABINOVICH, M. I., RULKOV, N. F. ROWAT, P. F., and SELVERSTON, A. I. (1996): 'Synchronized action of synaptically coupled chaotic model neurons', *Neural Communication*, 8, pp. 1567 - 1602

AKAIKE, H. (1969): 'Fitting autoregressions for prediction', *Ann. Inst. Stat. Math.*, 21, pp. 243-247

ANDRONOV, A. A., VITT, A. A., and KHAIKIN S. E. (1987): 'Theory of Oscillators', (Dover, original text 1966)

BABLOYANTZ, A. (1985): 'Strange attractors in the dynamics of brain activity', in HAKEN, H. (ed), 'Complex Systems Operational Approaches', pp. 116-123, (Springer-Verlag, Berlin)

BABLOYANTZ, A., and DESTEXHE, A. (1986): 'Low dimensional chaos in an instance of epilepsy', *Proc. Natl. Acad. Sci.*, 83, pp. 3513 - 2517

BARLOW, J. S. (1993): 'The Electroencephalogram: Its Patterns and Origins', (MIT Press, Cambridge, MA)

BAŞAR, E. (1983): 'Synergetics of neuronal populations', in BASAR, E., FLOHR, H., HAKEN, H., and MANDELL, A. J. (eds): 'Synergetics of the Brain', (Springer, New York)

BAŞAR, E. (1990): 'Chaos in Brain Function', (Springer-Verlag, Berlin, Heidelberg, New York)

BILNOWSKA, K. J., and MALINOWSKI, M. (1991): 'Non-linear and linear forecasting of the EEG time series', *Biol. Cybern.*, 66, pp. 159 - 165

BIRBAUMER, N., *et al* (1999): 'A spelling device for the completely paralyzed', *Nature*, 398, pp. 297-298

BISHOP, C. M. (1995): 'Neural Networks for Pattern Recognition', (Clarendon Press, Oxford)

BLACK, F., and SCHOLES, M. (1973): 'The pricing of options and corporate liabilities', *J. Political Economy*, 81, pp. 637-659

BOWER, J. M., and BEEMAN, D. (1998): 'The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System', (2nd ed., Springer-Verlag, NY)

BOX, G. E. P, and MULLER, M.E. (1958): 'A note on the generation of random normal deviates,' *Annals Math. Stat.*, 29, pp. 610-611

BOYNTON, G. M., ENGEL, S. A., GLOVER, G. H., and HEEGER, D. J. (1996): 'Linear systems analysis of functional magnetic resonance imaging in human VI', *J. Neurosci.*, 16, pp. 4207-4221

BURKE, D., WARD, T., and DE PAOR, A. (2001): 'Image processing used to harness blinking as a channel of communication and control for physically disabled people', *Med. Biol. Eng. Comput.*, 39, pp. 285 – 287

BURKE, D. and DE PAOR, A. (2002): 'An evoked potential model based on the transient synchronisation of coupled oscillators', in Proceedings of the 2nd European Medical & Biological Engineering Conference, EMBEC'02, Vienna, Dec, pp. 1434 – 1437

BURKE, D., KELLY, S., DE CHAZAL, P., and REILLY, R. (2002): 'A simultaneous filtering and feature extraction strategy for direct brain interfacing', in Proceedings of the 24th Annual International Conference of the Engineering in Medicine and Biology Society (EMBS-BMES), Houston, Texas (on CDROM)

CASDAGLI, M. (1991): 'Chaos and deterministic versus stochastic non-linear modelling', *J. R. Statist. Soc. B.*, 54, pp. 303 – 328

CASDAGLI, M. C., IASEMIDIS, L. D., SAVIT, R. S., GILMORE, R. L., ROPER, S. N., and SACKELLARES, J. C. (1997): 'Non-linearity in invasive EEG recordings from patients with temporal lope epilepsy', *Electroenceph. Clin. Neurophysiol.*, 102, pp. 98 – 105

CERUTTI, S., CHIARENZA, G., LIBERATI, D., MASCELLANI, P., and PAVESI, G. (1988): 'A parametric method of identification of single-trial event-related potentials in the brain', *IEEE Trans. Biomed. Eng.*, 35, pp. 701-711

CHILDERS, D. G., PERRY, N. W., FISCHLER, I. A., BOAZ, T., and ARROYO, A. A. (1987): 'Event-related potentials: A critical review of methods for single-trial detection', *CRC Critical Review in Biomedical Engineering*, 14, pp, 185 – 200

CLARKE, L., HARPER, P., and REILLY, R., (1998): 'Video Based Gesture Recognition for Augmentative Communication', in TIDE Third Congress Papers, Helsinki, July

CLEVELAND, N. R., and DOYLE, P. G., (1992): 'Eyegaze Communication System, How Does It Work? Who Can Use It?', in Closing The Gap – 10th Annual Conference, Minneapolis, Minnesota, Oct 22-24

COHEN, M. S., and BOOKHEIMER, S. Y. (1994): 'Localization of brain function using magnetic resonance imaging', *Trends in Neurosci.*, 17, pp. 268-277

COHEN, M. S. (1997): 'Parametric analysis of fMRI data using linear systems methods', *NeuroImage*, 6, pp. 83 – 103

DEECKE, L. SCHEID, R. P., and KORNHUBER, H. H. (1969): 'Distribution of readiness potential, pre-motion positivity and motor potential of the human cerebral cortex preceding voluntary finger movement', *Exp. Brain Res*., 7, pp. 158-168

DEFELIPE, J., and FARINAS, I. (1992): 'The pyramidal neuron of the cerebral cortex: morphological and chemical characteristics of the synaptic inputs', *Prog. Neurobiol.*, 39, pp. 563 - 607

DEMIRALP, T., ADEMOGLU, A., ISTEFANOPULOS, Y., and GULCUR, H. O. (1998): 'Analysis of event-related potentials (ERP) by damped sinusoids', *Biol. Cybern*., 78, pp. 487-493

DEWAN, E. M. (1964): 'Nonlinear Oscillations and Electroencephalography', *J. Theoret. Biol*., 7 pp. 141-159

DONCHIN E., SPENCER, K. M., and WIJESINGHE, R. (2000): 'The mental prosthesis: Assessing the speed of a P300-based brain-computer interface', *IEEE Trans. Rehab. Eng.*, 8, pp. 174 – 179

DUMERMUTH, G., and MOLINARI, L. (1987): 'Spectral analysis of EEG background activity', in Gevins, A. S., Remond, A. (eds): 'Methods of Analysis of Brain Electrical and Magnetic Signals. EEG Handbook', (Vol 1, Elsevier Science Publishers, B. V.)

ECKHORN, R. (1994): 'Oscillatory and non-oscillatory synchronizations in the visual cortex and their possible roles in associations of visual features', *Prog. Brain Res*., 102, pp. 405-426

ECKMANN, J. P., KAMPHORST, S. O., RUELLE, D., and CILIBERTO, S. (1986): 'Lyapunov exponents from time series', *Phys. Rev. A*, 34, pp. 4971 - 4979

ELBERT, T., RAY, W. J., KOWALIK, Z. J., SKINNER, J. E., GRAF, K. E., and BIRBAUMER, N. (1994): 'Chaos and physiology: Deterministic chaos in excitable cell assemblies', *Phys. Rev.*, 74, pp. 1 – 47

ERMENTROUT, G. B. (1981): 'n:m Phase-Locking of Weakly Coupled Oscillators', *J. Math. Biol.*, 12, pp. 327 – 342

ERMENTROUT, G. B. (1996): 'Type I membranes, phase resetting curves, and synchrony', *Neural Comput.*, 8, pp. 979 – 1001

ERMENTROUT, G. B., and KOPELL, N. (1984): 'Frequency plateaus in a chain of weakly coupled oscillators', *SIAM J. Math. Anal.*, 15, pp. 215-237

ERMENTROUT, G. B., and KOPELL, N. (1990): 'Oscillator death in a system of coupled neural oscillators', *SIAM J. Appl. Math.* 50, pp. 125-146

ERMENTROUT, G. B., and KOPELL, N. (1991): 'Multiple pulse interactions and averaging in systems of coupled neural oscillators', *J. Math. Biol.*, 29, pp. 195-217

FARMER, J. D., and SIDOROWICH, J. J. (1987): 'Predicting chaotic time series', *Phys. Rev. Lett.*, 59, pp. 845 - 848

FARWELL, L. A., and DONCHIN, E., (1988): 'Talking off the top of your head: A mental prosthesis utilizing event-related brain potentials', *Electroenceph. Clin. Neurophysiol.*, 70, pp. 355-372

FINUCANE, C,. BURKE, D., and DE PAOR, A. (2003): 'Investigations into a dual modality P300 based Brain-Computer Interface', in Proceedings of the 10th International Conference on Human-Computer Interaction, Crete, Greece, June

FIRSCHING, R. (1998): 'Moral dilemmas of tetraplegia; the 'locked-in' syndrome, the persistent vegetative state and brain death', *Spinal Cord*, 36, pp. 741-743

FLYNN, S. (2001): *Personal communication*, Occupational Therapy Department, National Rehabilitation Hospital, Dun Laoghaire, Co. Dublin, Ireland

FOWLER, M., and SCOTT, K. (2000): 'UML Distilled: A Brief Guide to the Standard Object Modeling Language', (Addison-Wesley Professional)

FREEMAN, W. J. (1975): 'Mass Action in the Nervous System', (Academic Press, New York)

FREEMAN, W. J. (1979): 'Nonlinear gain mediating cortical stimulus-response relations', *Biol. Cybern.*, 33, pp.237-247

FREEMAN, W. J. (1992): 'Tutorial on Neurobiology: From Single Neurons to Brain Chaos', *International Journal of Bifurcation and Chaos*, 2, pp. 451-482

FREEMAN, W. J. (1999): 'Noise-induced first-order phase transitions in chaotic brain activity', *International Journal of Bifurcation and Chaos*, 9, pp. 2215-2218

FREEMAN, W. J. (2000a): 'Neurodynamics; An Exploration of Mesoscopic Brain Dynamics', (Springer-Verlag, London UK)

FREEMAN, W. J. (2000b): 'A proposed name for aperiodic brain activity: stochastic chaos', *Neural Networks*, 13, 11-13

FREEMAN, W. J. (2000c): 'Spatial spectral analysis of human electrocorticograms including the alpha and gamma bands', *Journal of Neuroscience Methods*, 95, pp. 111-121

FREEMAN, W. J., and JAKUBITH, S. (1993): 'Bifurcation analysis of continuous time dynamics of oscillatory neural networks', in AERTSEN, A.: 'Brain Theory', (Elsevier Science Publishers)

GAMMA, E., HELM, R., JOHNSON, R., and VLISSIDES, J. (1995): 'Design Patterns: Elements of Reusable Object-Oriented Software', (Addison-Wesley Professional)

GARAVAN, H., ROSS, T. J., and STEIN, E. A. (1999): 'Right hemispheric dominance of inhibitory control: An event-related functional MRI study', *Proc. Natl. Acad. Sci.*, 86, pp. 8301 – 8306

GARD, T. C. (1988): 'Introduction to Stochastic Differential Equations', (Marcel Dekker, New York)

GOLOMB, D., WANG, X. J., and RINZEL, J. (1996): 'Propagation of spindle waves in a thalamic slice model', *J. Neurphysiol*. 5, pp. 750-769

GONALES, R. C., and WINTZ, P., (1987): 'Digital Image Processing', (Addison-Wesley, 2nd edition)

GRADISEK, J., SIEGERT, S., FREIDRICH, R., and GRABEC, I. (2000): 'Analysis of time series from stochastic processes', *Phys. Rev. E*, 62, pp. 3146 - 3155

GRASSBERGER, P., and PROCACCIA, I. (1983): 'Characterisation of strange attractors', *Phys. Rev. Lett*., 50, pp. 346

GUCKENHEIMER, J., and HOLMES, P. (1983): 'Nonlinear oscillations, dynamical systems, and bifurcations of vector fields', (Springer-Verlag, New York)

GUGER, C., SCHLÖGL, A., NEUPER, C., WALTERSPACHER, D., STREIN, T., and PFURTSCHELLER, G. (2001): 'Rapid prototyping of an EEG-Based Brain-Computer Interface (BCI)', *IEEE Trans. Rehab. Eng*., 9, pp. 49-58

GUPTA, L., MOLFESE, D. L., TAMMANA, R., and SIMOS, P. G. (1996): 'Nonlinear alignment and averaging for estimating the evoked potential', *IEEE Trans. Biomed. Eng.*, 43, pp. 348 – 355

GUYTON, A. C. (1991): 'Textbook of Medical Physiology', (8th ed., Saunders W. B. Co.)

HALPERIN, W. P., REYES, A. P., KUHNS, P., and MOULTON, W. G. (2001): 'Spatially resolved electronic structure inside and outside the vortex cores of a high-temperature superconductor', *Nature*, 413, pp. 501 – 504

HANSEL, D., MATO, M., and MEUNIER, C. (1995): 'Synchrony in excitatory neural networks', *Neural Comput.*, 7, pp. 307-337

HAS'MINSKII, R. Z. (1980): 'Stochastic Stability of Differential Equations', (Sijthoff and Noordhoff, Alphen aan den Rijn, Netherlands)

HEGGER, R., KANTZ, H., and SCHREIBER, T. (1999): 'Practical implementation of nonlinear time series methods: The TISEAN package', *Chaos*, 9, pp. 413

HERNÁNDEZ, J. L., VALDÉS, P. A., and VILA P (1996): 'EEG spike and wave modelled by a stochastic limit cycle', *NeuroReport*, 7, pp. 2246-2250

HIGHAM, D. J. (2001): 'An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations', *SIAM Rev.*, 43, 3, pp. 525-546

HINDMARSH, J. L., and ROSE, R. M. (1982): 'A model of the nerve impulse using two first order differential equations', *Nature*, 296, pp. 162-164

HINDMARSH, J. L., and ROSE, R. M. (1984): 'A model of neuronal bursting using three coupled first order diff equations', *Proc. R. Soc. Lond*, B221, pp. 87 - 102

HODGKIN, A. L., and HUXLEY, A. F. (1952): 'A quantitative description of membrane current and its application to conduction and excitation of nerve', *J. Physiol.*, 117, pp. 500 – 544

HOPPENSTEADT, F. C., and IZHIKEVICH, E. (1997): 'Weakly Connected Neural Networks', (Springer)

HOROVITZ, S., SKUDLARSKI, P., and GORE, J. (2002): 'Correlations and dissociations between BOLD signal and P300 amplitude in an auditory oddball task: a parametric approach to combining fMRI and ERP', *Magn. Reson. Imaging*, 20, pp. 319

HUANG-HELLINGER, F. R., BREITER H. C., MCCORMACK, G., COHEN, M. S., KWONG, K. K., SUTTON, J. P., SAVOY, R. L., WEISSKOFF, R. M., DAVIS, T. L., BAKER, J. R.,

BELLIVEAU, J. W., and ROSEN, B. R. (1995): 'Simultaneous functional magnetic resonance imaging and electrophysiological recording', *Human Brain Mapping*, 3, pp.13-25

HUGGINS, J. E. *et al* (1999): 'Detection of event-related potentials for development of a direct brain interface', *J. Clin. Neurophysiol*., 16, pp. 448-455

ISAKSSON, A., and WENNBERG, A. (1975): 'Visual Evaluation and Computer Analysis of the EEG – A Comparison', *Electroenceph. Clin. Neurophysiol*., 38, pp. 79-86

JALIFE, J. (1984): 'Mutual entrainment and electrical coupling as mechanisms for synchronous firing of rabbit sinoatrial pacemaker cells', *J. Physiol*., 356, pp. 221-243

JANSEN, B. H., ZOURIDAKIS, G., and BRANDT, M. E. (1993): 'A neurophysiologically-based mathematical model of flash visual evoked potentials', *Biol. Cybern*., 68, pp. 275 - 283

JANSEN, B. H, and RIT, V. G. (1995): 'Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns', *Biol. Cybern.*, 73, pp. 357 – 366

JASPER, H. H. (1936): 'Cortical excitatory state and variability in human brain rhythms', *Science*, 83, pp. 259 – 260

JOSEPHS, O., TURNER, R., and FRISTON, K. (1997): 'Event-related fMRI', *Human Brain Mapping*, 5, pp. 243-248

KANTZ, H. (1994): 'A robust method to estimate the maximal Lyapunov exponent of a time series', *Phys. Lett*., A185, pp. 77

KANTZ, H. (2001): 'Time series analysis in reconstructed state spaces', *Stochastics and Dynamics*, 1, pp. 85 - 111

KANTZ, H., and SCHREIBER, T. (1997): 'Nonlinear time series analysis', (Cambridge University Press, Cambridge)

KASHIKURA, K., SATO, T., TOYODA, H. NISHIZAWA, S., FUJIBAYASHI, Y., and YONEKURA, Y. (2002): 'The frequency dependence of BOLD fMRI response induced in human V1 is influenced by different photic patterns', in Proceedings of the 8th International Conference on Functional Mapping of the Human Brain, June 2-6, Sendai, Japan

KEENER, J., and SNEYD, J. (1998): 'Mathematical Physiology', (Springer-Verlag, New York, Berlin, Heidelberg)

KLOEDEN, P. E., PLATEN E., and SCHURZ, H. (1991): 'The Numerical Solution of Nonlinear Stochastic Dynamical Systems: A Brief Introduction', *International Journal of Bifurcation and Chaos*, 1, pp. 277-286

KLOEDEN, P. E., and PLATEN, E. (1999): 'Numerical Solution of Stochastic Differential Equations', (3rd ed., Springer-Verlag)

KOCH, C., and SEGEV, I. (1998): 'Methods in Neuronal Modelling: From Ions to Networks', (2nd ed., MIT Press)

KOLUMBÁN, G., KENNEDY, M. P., and CHUA, L. O. (1997): 'The role of synchronization in digital communications using chaos - Part I: Fundamentals of Digital Communications', *IEEE Trans. Circuits Syst*., 44, pp. 927-935

KOLUMBÁN, G., KENNEDY, M. P., and CHUA, L. O. (1998): 'The role of synchronization in digital communications using chaos - Part II: Chaotic modulation and chaotic synchronization', *IEEE Trans. Circuits Syst*., 45, pp. 1129-1140

KÖNIG, P., and ENGEL, A. K. (1995): 'Correlated firing in sensory-motor systems', *Curr. Opin. Neurobiol*., 5, pp. 511-519

KOPELL, N. (1988): 'Toward a theory of modelling central pattern generators', in COHEN, A. H., 'Neural Control of Rhythmic Movements in Vertebrates', (Wiley, New York)

KOPELL, N., and ERMENTROUT, G. B. (1986): 'Symmetry and phaselocking in chains of weakly coupled oscillator', *Comm. Pure and Appl. Math.,* 39, pp. 623 – 660

KOSSLYN, S. M., THOMPSON, W. L., WRAGA, M., and ALPERT, N. M. (2001): 'Imagining rotation by endogenous versus exogenous forces: Distinct neural mechanisms', *NeuroReport*, 12, pp. 2519-2525

KWONG, K. K., BELLIVEAU, J. W., CHESLER, D. A., GOLDBERG, I. E., WEISSKOFF, R. M., PONCELET, B. P., KENNEDY, D. N., HOPPEL, B. E., COHEN, M. S., TURNER, R. *et al* (1992): 'Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation', *Proc. Natl. Acad. Sci*, 89, pp. 5675 - 5679

LA SALLE, J., and LEFSCHETZ, S. (1961): 'Stability by Liapunov's Direct Method with Applications', (Academic Press, New York, London)

LANSKY, P., and LANSKA, V. (1987): 'Diffusion Approximation of the Neuronal Model with Synaptic Reversal Potentials', *Biol. Cybern*., 56, pp. 19-26

LEHNERTZ, K, and ELGER, C. E. (1998): 'Can epileptic seizures be predicted? Evidence from nonlinear time series analysis of brain electrical activity', *Phys. Rev. Lett*., 80, pp. 5019 - 5022

LEHNERTZ, K., ARNHOLD, J., GRASSBERGER, P., and ELGER, C. E. (1999): 'Proceedings of the Workshop on Chaos in Brain?', (University of Bonn, Germany)

LILEY, D. T. J., ALEXANDER, D. M., WRIGHT, J. J., and ALDOUS, M. D. (1999): 'Alpha rhythm emerges from large-scale networks of realistically coupled multicompartmental model cortical neurons', *Network: Comput. Neural Syst*., 10, pp. 79 – 92

LJUNG, L. (1987): 'System Identification: Theory for the User', (Prentice Hall)

LOPES DA SILVA, F., HOEKS, A., SMITS, H., and ZETTERBERG, L. H. (1974): 'Model of brain rhythmic activity', *Kybernetik*, 15, pp. 27-37

LOPES DA SILVA, F. H., PIJN, J. P. M., GORTER, J. A., VLIET, E. V., DAALMAN, E. W., and BLANES, W. (1999): 'Is nonlinearity evident in time series of brain electrical activity', in LEHNERTZ, K., ARNHOLD, J., GRASSBERGER, P., and ELGER, C. E.: 'Proceedings of the Workshop on Chaos in Brain?', (University of Bonn, Germany)

MALSBURG VON DER, C. (1981): 'The correlation theory of brain functions', Internal Report 81-2, Max-Planck-Institute for Biophysical Chemistry, Göttingen, FRG

MANDELBROT, B. B. (1985): 'The Fractal Geometry of Nature', (Freeman, San Francisco)

MATSUMOTO, M., and NISHIMURA, T. (1998): 'Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator', *ACM Transactions on Modeling and Computer Simulation*, 8, pp. 3

MC CARTHY, G., LUBY, M., GORE, J., and GOLDMAN-RAKIC, P. (1997): 'Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI' *J. Neurophysiol*., 77, pp. 1630 – 1634

MC DARBY, G. (2000): 'Making wavelets: Custom designing 2-channel filter banks with associated biorthogonal wavelet bases in L2(R)', PhD thesis, University of New South Wales, Australia

MERKIN, D. R. (1997): 'Introduction to the Theory of Stability', (Springer-Verlag, New York)

MIDDENDORF, M., MCMILLAN, G., CALHOUN, G., and JONES, K. S. (2000): 'Brain-computer interfaces based on the steady-state visual-evoked response', *IEEE Trans. Rehab. Eng.*, 8, pp. 211-214

MILNER, P. M. (1974): 'A model for visual shape recognition', *Psychol. Rev*., 81, pp. 521-535

MIROLLO, R. E., and STROGATZ, S. H. (1990): 'Synchronization of pulse-coupled biological oscillators', *SIAM J. Appl. Math*., 50, pp. 1645-1662

MISULIS, K. E. (1994): 'Spehlmann's evoked potential primer: visual, auditory, and somatosensory evoked potentials in clinical diagnosis', (Butterworth-Heinemann, MA)

MIYAZAKI, S., ISHIDA, A., KOMATSUZAKI, A, (2000): 'A Clinically Oriented Video-Based System for Quantification of Eyelid Movements', *IEEE Trans. Biomed. Eng*., 47, pp. 1088 – 1096

MURRAY, J. D. (1990): 'Mathematical Biology', (Biomathematics, 19, Springer-Verlag, New York)

NIEDERMEYER, E. (1999): 'The Normal EEG of Waking Adult', in NIEDERMEYER, E., LOPES DA SILVA, F. (eds): 'Electroencephalography: Basic Principles, Clinical Applications, and Related Fields', (4th ed., Lippincott Williams & Wilkins, Baltimore)

NIEDERMEYER, E., and LOPES DA SILVA, F. (1999): 'Electroencephalography: Basic Principles, Clinical Applications, and Related Fields', (4th ed., Lippincott Williams & Wilkins, Baltimore)

NOGAWA, T., KATAYAMA, K., OKUDA, H., and UCHIDA, M. (1991): 'Changes in the latency of the maximum positive peak of visual evoked potentials during anaesthesia', *Nippon Geka Hokan* (Archives of Japanese Surgery), 60, pp. 143-153

NUNEZ, P. L. (1995): 'Neocortical dynamics and human EEG rhythms', (Oxford University Press, New York)

ØKSENDAL, B. (1998): 'Stochastic Differential Equations: An Introduction with Applications', (5th ed., Universitext, Springer)

PAOR DE, A. M. (1994): 'Liapunov, Poincare-Bendixson and Cardiac Fibrillation', *Lekar A Technika*, 25, pp. 75 – 79

PARDEY, J., ROBERTS, S., and TARASSENKO, L. (1996): 'A review of parametric modelling techniques for EEG analysis', *Med. Eng. Phys.*, 18, pp. 2-11

PATMORE, D. W., and KNAPP, R. B. (1995): 'A cursor controller using evoked potentials and EOG', RESNA, June

PECORA, L. M., and CARROLL, T. L. (1990): 'Synchronization in chaotic systems', *Phys. Rev. Lett.*, 64, pp. 821 - 824

PENNY, W. D., ROBERTS, S. J., CURRAN, E. A., and STOKES, M. J. (2000): 'EEG-based communication: A pattern recognition approach', *IEEE Trans. Rehab. Eng.*, 8, pp. 214-215

PFURTSCHELLER, G., KALCHER, J., NEUPER, C. , FLOTZINGER, D., and PREGENZER, M. (1996): 'On-line EEG classification during externally-paced hand movements using a neural network-based classifier', *Electroenceph. Clin. Neurophysiol.*, 99, pp. 416-425

PIJN, J. P., VAN NEERVEN, J., NOEST, A., and LOPES DA SILVA, F. H. (1991): 'Chaos or noise in EEG signals: Dependence on state and brain site', *Electroenceph. Clin. Neurophysiol.*, 79, pp. 371-381

PLATEN, E. (1999): 'An introduction to the numerical methods for stochastic differential equations', *Acta. Numerica.*, pp. 197-246

PLONSEY, R., and BARR, R. C. (1998): 'Bioelectricity: A Quantitative Approach', (Plenum Press)

PRAST, J. W. (1949): 'An interpretation of certain EEG patterns as transient responses of a transmission system', *Electroenceph. Clin. Neurophysio*., 1, pp. 370

PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., and FLANNERY, B. P. (1992): 'Numerical Recipes in C: The Art of Scientific Computing', (2nd ed, Cambridge University Press, New York)

PRITCHARD, W. S., and DUKE, D. W. (1992): 'Measuring chaos in the brain: A tutorial review of nonlinear dynamical EEG analysis', *Intern. J. Neuroscience*, 67, pp. 31 – 80

PROTOPAPAS, A., and BOWER, J. M. (1998): 'Dynamics of Cerebral Cortical Networks' in Bower, J. M., Beeman, D. (eds), 'The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System', (2nd Ed., TELOS, Springer-Verlag, New York)

PROTOPAPAS, A., D., VANIER, M., and BOWER, J. (1998): 'Simulating Large Networks of Neurons', in KOCH, C., and SEGEV, I., 'Methods in Neuronal Modelling: From Ions to Networks', (2nd Ed., MIT Press)

PROVENZALE, A., SMITH, L. A., VIO, R., and MURANTE, G. (1992): 'Distinguishing between low-dimensional dynamics and randomness in measured time series', *Physica D*, 58, pp. 31 - 49

PUTTEN VAN, M. J. A. M., and STAM, C. J. (2001): 'Is the EEG really chaotic in hypsarrhythmia?', *IEEE Engineering in Medicine and Biology*, Sept / Oct, pp. 72 – 79

RAND, R. H., COHEN, A. H., and HOLMES, P. J. (1988): 'Systems of coupled oscillators as models of central pattern generators', in COHEN, A. H., 'Neural Control of Rhythmic Movements in Vertebrates', (Wiley, New York)

REILLY, R. B., WARD, T. E., and DE PAOR, A. M. (1996): 'An uncoupled oscillator model for evoked potential dynamical modeling', in Proceedings of the 18th Annual International Conference IEEE/EMBS, Amsterdam

REKTORYS, K. (1969): 'Survey of Applicable Mathematics', (Iliffe Books Ltd, London)

RIPLEY, B. D. (1996): 'Pattern Recognition and Neural Networks', (Cambridge University Press)

ROMBOUTS, S. A. R. B., KEUNEN, R. W. M., and STAM, C. J. (1995): 'Investigation of nonlinear structure in multichannel EEG', *Phys. Lett. A.*, 202, pp. 352 – 358

ROSEN, B. R., BUCKNER, R. L., and DALE, A. M. (1998): 'Event-related functional MRI: Past, present, and future', *Proc. Natl. Acad. Sci*., 95, pp. 773-780

ROTTERDAM VAN, A., LOPES DA SILVA, F., ENDE VAN DEN, J., VIERGEVER, M. A., and HERMANS, A. J. (1982): 'A model of the spatial-temporal characteristics of the alpha rhythm', *Bull. Math. Biol*., 44, pp. 283-305

SAUER, T., YORKE, J., and CASDAGLI, M. (1991): 'Embedology', *J. Stat. Phys*. 65, pp. 579

SCHREIBER, T. (1999): 'Is nonlinearity evident in time series of brain electrical activity', in LEHNERTZ, K., ARNHOLD, J., GRASSBERGER, P., and ELGER, C. E.: 'Proceedings of the Workshop on Chaos in Brain?', (University of Bonn, Germany)

SCHREIBER, T., and SCHMITZ, A. (2000): 'Surrogate time series', *Physica D*, 142, pp. 346 - 382

SIEGERT, S., FRIEDRICH, R., and PEINKE, J. (1998): 'Analysis of data sets of stochastic systems', *Phys. Lett. A*, 243, pp. 275 - 280

SINGER, W., and GRAY, C. M. (1995): 'Visual feature integration and the temporal correlation hypothesis', *Ann. Rev. Neurosci.*, 18, pp. 555-586

SOONG, T. T. (1973): 'Random Differential Equations in Science and Engineering', (Academic Press, New York)

SPECKMANN, E. J., and ELGER, C. E. (1999): 'Cellular Substrates of Brain Rhythms', in NIEDERMEYER, E., and LOPES DA SILVA, F. (eds): 'Electroencephalography: Basic Principles, Clinical Applications, and Related Fields', (4th ed., Lippincott Williams & Wilkins, Baltimore)

STAM, C. J., PIJN, J. P. M., SUFFCZYNSKI, P., and LOPES DA SILVA, F. H. (1999): 'Dynamics of the human alpha rhythm: Evidence for nonlinearity?', *Clin. Neurphysiol*., 110, pp. 1801 – 1813

STERIADE, M. (1999): 'Cellular Substrates of Brain Rhythms', in NIEDERMEYER, E., and LOPES DA SILVA, F. (EDS): 'Electroencephalography: Basic Principles, Clinical Applications, and Related Fields', (4th ed., Lippincott Williams & Wilkins, Baltimore)

STURM, A. K., and KÖNIG, P. (2001): 'Mechanisms to synchronize neuronal activity', *Biol. Cybern*., 84, pp. 153-172

SUGIHARA, G., and MAY, R. M. (1990): 'Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series', *Nature*, 344, pp. 734 - 741

SUTTER, E. (1992): 'The brain response interface: Communication through visually-induced electrical brain response', *J. Microcomput. Appl*., 15, pp. 31-45

TAKENS, F. (1981): 'Detecting strange attractors in turbulence', *Lecture Notes in Mathematics*, 898, (Springer)

THEILER, J. (1986): 'Spurious dimensions from correlation algorithms applied to limited time-series data', *Phys. Rev*, A34, pp. 2427

THEILER, J., EUBANK, S., LONGTIN, A., GALDRIKIAN, B., and FARMER, J. D. (1992): 'Testing for nonlinearity in time series: The method of surrogate data', *Physica D*, 58, pp. 77 – 94

THEILER, J., and RAPP, P. E. (1996): 'Re-examination of the evidence for low-dimensional, nonlinear structure in the human electroencephalogram', *Electroenceph. Clin. Neurophysiol*., 98, pp. 213-222

THOMSON, J. M. T., and STEWART, H. B. (1986): 'Nonlinear Dynamics and Chaos', (John Wiley and Sons)

THORNTON, C. (1991): 'Evoked potentials in anaesthesia', *European Journal of Anaesthesiology*, 8, pp. 89-107

TRAUB, R. D., and WONG, R. K. S. (1982): 'Cellular mechanism of neuronal synchronisation in epilepsy', *Science*, 216, pp. 745-747

TRAUB, R. D., MILES, R., and WONG, R. K. S. (1989): 'Model of the origin of rhythmic population oscillations in the hippocampal slice', *Science*, 243, pp. 1319-1325

WALTER, W. G. (1962) in Grenell, R. (eds): 'Neural Physiopathology', (Harper and Row)

WARD, T. E. (1999): 'Evoked potential modelling and biological information processing based on nonlinear interactions in communities of neurons', PhD thesis, National University of Ireland, Dublin

WEISSKOFF, R. M., BAKER, J. R., BELLIVEAU, J. W., DAVIS, T. L., KWONG, K. K., COHEN, M. S., *et al* (1993): 'Power spectrum analysis of functionally-weighted MR

data: What's in the noise?' in Proceedings of the 12th Annual Conference of the Society of Magnetic Resonance in Medicine', New York

WENNBERG, A., and ZETTERBERG, L. H. (1971): 'Application of a computer based model of the EEG analysis', *Electroenceph. Clin. Neurophysiol.*, 31, pp. 457-468

WILSON, M. A., and BOWER, J. M. (1989): 'The simulation of large scale neural networks' in KOCH, C., and SEGEV, I.: 'Methods in Neuronal Modelling: From Ions to Networks', (1st Ed., MIT Press)

WILSON, H. R., and COWAN, J. D. (1972): 'Excitatory and inhibitory interactions in localized populations of model neurons', *Biophysical Journal*, 12, pp. 1-24

WILSON, H. R., and COWAN, J. D. (1973): 'A Mathematical Theory of the Functional Dynamics of Cortical and Thalamic Nervous Tissue', *Kybernetik*, 13, pp. 55-80

WINFREE, A. T. (1967): 'Biological rhythms and the behaviour of populations of coupled oscillators', *J. Theoret. Biol.*, 16, pp. 15-42

WINFREE, A. T. (1980): 'The Geometry of Biological Time', (Springer-Verlag, New York)

WOLF, A., SWIFT, J. B., SWINNEY, H. L., and VASTANO, J. A. (1985): 'Determining Lyapunov exponents from a time series', *Physica D*, 16, pp. 285-317

WOLPAW, J. R., MCFARLAND, D. J., NEAT, G. W., and FORNERIS, C. A. (1991): 'An EEG-based brain-computer interface for cursor control', *Electroenceph. Clin. Neurophysiol.*, 78, pp. 252-259

WOLPAW, J. R., *et al* (2000): 'Brain-computer interface technology: a review of the first international meeting', *IEEE Trans. Rehab. Eng.*, 8, pp. 164-173

WRIGHT, J. J., and LILEY, D. T. J. (1996): 'Dynamics of the brain at global and microscopic scales: Neural networks and the EEG', *Behav. Brain Sci*., 19, pp. 285-320

YU, X., HE, Z. Y., and ZHANG, Y. S. (1994): 'Peak component latency-corrected average method for evoked potential waveform estimation', *IEEE Trans. Biomed. Eng*., 41, pp. 1072-1082

ZETTERBERG, L. H. (1969): 'Estimation of Parameters for a Linear Difference Equation with Application to EEG Analysis', *Mathematical Biosciences*, 5, pp. 227-275

ZETTERBERG, L. H. (1977): 'Means and methods for processing of physiological signals with emphasis on EEG analysis', in Lawrence, J. H., *et al*: 'Advances in Biology and Medical Physics', 16, pp. 41-91, (Academic Press, New York)

APPENDIX A1

ELECTRONIC CIRCUIT REALISATION OF A PRESCRIBED OSCILLATOR

The method for prescribing oscillators presented in Chapter II may easily be implemented in hardware. Analogue electronic circuit realisations may be used as a function generator for multiple purposes such as testing, modelling, and education. Here we present an electronic version of the Cassinian Ovals limit cycle oscillator (Example 2.4).

The prescribed oscillator system is given by

$$\dot{x}_1 = \frac{\partial H}{\partial x_2} - \lambda H \frac{\partial H}{\partial x_1}$$
$$\dot{x}_2 = -\frac{\partial H}{\partial x_1} - \lambda H \frac{\partial H}{\partial x_2}, \quad \lambda > 0$$

where for the Cassinian Ovals we require

$$H(x_1, x_2) = \left[(x_1 - b)^2 + x_2^2 \right] \left[(x_1 - c)^2 + x_2^2 \right] - k = 0$$

and hence the partial derivatives are

$$\frac{\partial H}{\partial x_1} = 2(x_1 - b)\left((x_1 - c)^2 + x_2^2 \right) + 2(x_1 - b)\left((x_1 - b)^2 + x_2^2 \right)$$
$$\frac{\partial H}{\partial x_2} = 2x_2 \left((x_1 - c)^2 + (x_1 - b)^2 + 2x_2^2 \right)$$

Figure A1.1 illustrates the analogue circuit diagram. Constants *b*, *c*, *k* were generated via zener reference diodes (1.2V). Quad opamp ICs (LM348N) and multiplier ICs (AD633) were employed.

Figure A1.2 illustrates an oscilloscope image of the system where (0,0) forms a saddle node with two homoclinic trajectories (b = 1.0, c = -1.0, k = 1.0, λ = 1). The oscillator can be tuned (via k) to randomly oscillate between lobes in this configuration. (Although maintaining perfect randomness is not possible).

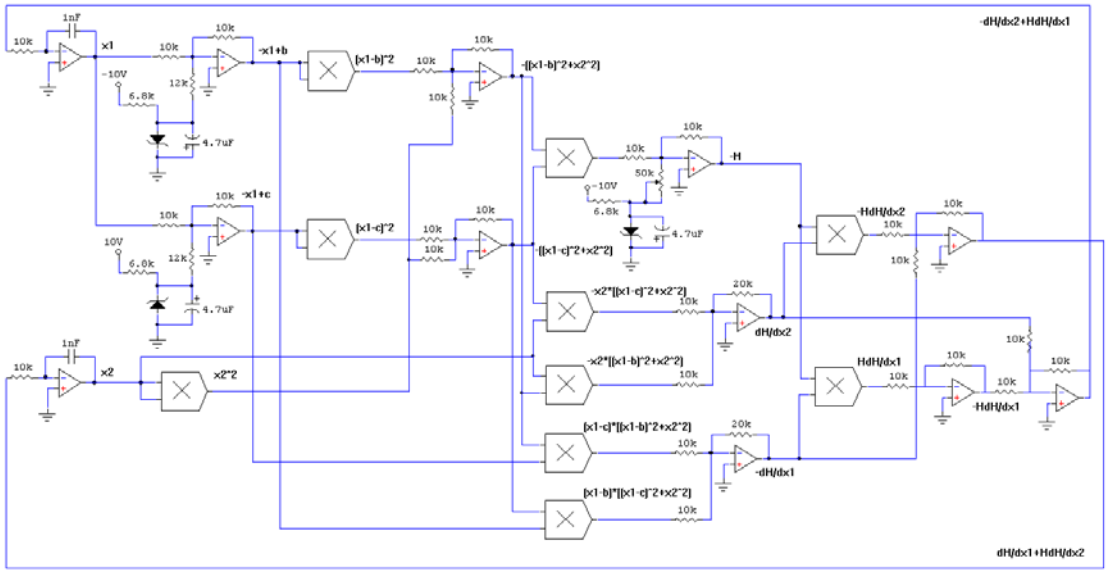

Figure A1.1: Analogue circuit diagram for the Cassinian Ovals oscillator

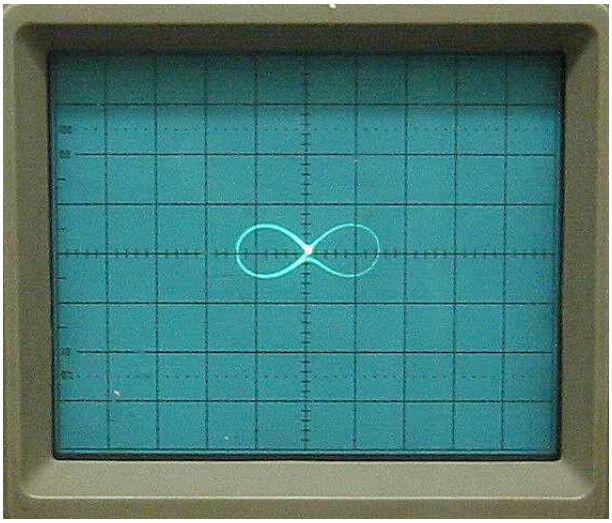

Figure A1.2: Oscilloscope X-Y trace for the Cassinian Ovals oscillator with $b = 1$, $c = -1$, $k = 1.0$, $\lambda = 1$