Quantitative Continuity Feature for Preterm Neonatal EEG Signal Analysis

by

Lisa Wong

Supervisor:

Dr. Waleed Abdulla

Co-Supervisor:

Dr. Mark Andrews

Medical Adviser:

Dr. Terrie Inder

In Association With:

BrainZ Instruments Ltd.

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in

Computer Systems Engineering

The University of Auckland

ABSTRACT

Electroencephalography (EEG) is an electrical signal recorded from a person's scalp, and is used to monitor the neurological state of the patient. This thesis proposes a quantified continuity feature to aid preterm neonatal EEG analysis. The continuity of EEG signals for preterm infants refers to the variation of the EEG amplitude, and is affected by the conceptional age of the infants. Currently, the continuity of the signal is determined largely by visual examination of the raw EEG signal, or by using general guidelines on amplitude-integrated EEG (aEEG), which is a compressed plot of the estimated signal envelope.

The proposed parametric feature embodies the statistical distribution parameters of the signal amplitudes. The signal is first segmented into pseudo-stationary segments using Generalized Likelihood Ratio (GLR). These segments are used to construct a vector of amplitude, the distribution of which can be modelled using a log-normal distribution. The mean and standard deviation of the log-normal distribution are used as the continuity feature. This feature is less prone to the effects of local transient activities than the aEEG.

This investigation has demonstrated that the degree of continuity corresponds to the major axis of the feature distribution in the feature space, and the minor axis roughly corresponds to the age of the infants in healthy files. Principal component analysis was performed on the feature, with the first coefficient used as a continuity index and the second coefficient as a maturation index. In this research, classifiers were developed to use the continuity feature to produce a qualitative continuity label. It was found that using a linear discriminant analysis based classifier, labelled data can be used as training data to produce labels consistent across all recordings. It was also found that unsupervised classifiers can assist in identifying the intrinsic clusters occurring in the recordings.

It was concluded that the proposed continuity feature can be used to aid further research in neonatal EEG analysis. Further work should focus on using the continuity information to provide a context for further feature extraction and analysis.

ACKNOWLEDGEMENTS

I was very fortunate to have received help from quite a number of people, without whom this thesis would not be possible. I would like to thank my academic supervisor, Dr. Waleed Abdulla, for his constant guidance, advice and support during this PhD research. I would also like to thank my co-supervisor, Dr Mark Andrews, and my Advisory Committee, Prof. Zoran Salcic and Dr. Catherine Watson, for providing valuable suggestions during various stages of the research.

This thesis would not have been possible without the help of the team at Brainz Instruments Ltd. Not only did they provide the valuable data used in this research, but they also assisted financially with research trips, and offered suggestions for the research direction. In particular I would like to thank Gordon Macdonald, Mark Stewart, and Michael Navakatikyan.

I would also like to thank A/Prof Terrie Inder, who has not only given valuable insight into the medical implications of this research, but also helped organise an opportunity for me to work with the wonderful team at the University of Washington in St. Louis (WUSTL) for 2 months to help make this research more relevant to the medical work in the field. I would also like to thank Prof Jeff Neil, Dr Divyen Shah, and the rest of the WUNDER team at WUSTL for being welcoming and offering their insights into this fascinating topic. Thank you also to the staff from the EEG lab, especially Dr John Zemple and Dr Mike Morrissey for explaining the procedure for EEG interpretation and giving advice on the project.

A big thank you is also due to my fellow researchers in the signal processing team at the University of Auckland. Not only did they offer suggestions, but their friendship and support also helped me survive the often frustrating research process. A heartfelt thank you to Kevin Wang, Tina Lin, Yushi Zhang, Octavian Cheng, and Kelvin Zeng.

Friends and family have been there for me every step of the way, and for that I am truly thankful. A special thank you to my husband, Peter Ravlich, who has not only supported me from the very beginning and tolerated my constant rambling of research progress, but also helped proofread this thesis and offered his professional editing service. I promise we will go on that honeymoon at some point.

Contents

List of Figures

List of Tables

List of Acronyms

aEEG	Amplitude Integrated Electroencephalogram
ANN	Artificial Neural Network
AP	Absolute Power
AR	Autoregressive
ARMA	Autoregressive Moving Average
BMU	Best Matched Unit
CA	Conceptional Age
CFM	Cerebral Function Monitoring
CT	Computed Tomography
CUS	Cranial Ultrasound
CWD	Choi-Williams Distribution
ECG	Electrocardiogram
ED	Exponential Distribution
EEG	Electroencephalogram
EMG	Electromyography
EOG	Electrooculography
FAR	False Alarm Rate
FFT	Fast Fourier Transform
FM	Frequency Modulation
FWE	Frequency-Weighted Energy
GA	Gestational Age
GLR	Generalized Likelihood Ratio
GMM	Gaussian Mixture Model
HIE	Hypoxic Ischaemic Encephalopathy
IBI	Inter-Burst Interval
ICU	Intensive Care Unit
IF	Instantaneous Frequency
LDA	Linear Discriminant Analysis

MBD	Modified B-Distribution
MIEF	Mutual Information Evaluation Function
MRI	Magnetic Resonance Imaging
NLEO	Nonlinear Energy Operator
PCA	Principal Component Analysis
PRSW	Positive Rolandic Sharp Waves
rEEG	Range-Electroencephalogram
REM	Rapid Eye Movement
RID	Reduced Interference Distribution
RP	Relative Power
SDR	Seizure Detection Rate
SEM	Spectral Error Measurement
SOM	Self Organising Map
STFT	Short Time Fourier Transform
TFD	Time Frequency Distribution
TS	Temporal Sawtooth
TSE	Total Square Error
VIBeS	Victoria Infant Brain Study
w-NLEO	Windowed Nonlinear Energy Operator
WVD	Wigner-Ville Distribution

Glossary

10-20 Electrode Placement

An international standard for electrode placement in EEG recordings.

Artificial Neural Network

A mathematical model based on the biological model of neurons, used in machine learning systems.

Burst Suppresion

An EEG signal with short periods of very high amplitude activities and very low amplitude inactivities in-between.

Cerebral Palsy

An umbrella term for motor disablities caused by brain defect or lesion.

Computed Tomography (CT)

Method for constructing 3D biomedical images using rotating beams of x-rays.

Conceptional Age (CA)

Age of an infant, measured from time of conception

Continuity

The variation of amplitude in an EEG recording.

Continuous Signal

An EEG signal with a relatively constant envelope.

Cortex

Outer layer of the brain, formed by neurons. Also known as Grey Matter.

Cortical EEG

An invasive form of EEG where the electrodes are placed surgically on the cortex of the brain.

Cranial Ultrasound (CUS)

Medical imaging technique that uses ultrasound to visualise the brain area.

Discontinuous Signal

An EEG signal with regions of high amplitude activities with low ampltiude activities in-between

Dysmature EEG

EEG signals that do not show expected maturation signs, and exhibit the behaviour expected in a younger patient's EEG.

Electrocardiogram (ECG)

An electrical signal originating in the muscular movements of the heart.

Electroencephalogram (EEG)

An electrical signal originating in the brain, measured from the scalp, and used for brain function monitoring.

Electrooculogram (EOG)

An electrical signal originating in movements of the eye.

Gestational Age (GA)

Age of infant at birth, measured from time of conception.

Grey Matter

See **Cortex**

Hypoxia

A lack of oxygen supply.

Hypoxic-Ischaemic Encephalopathy (HIE)

Damages to brain cells due to lack of blood flow and oxygen.

Inter-Burst Interval (IBI)

Periods of inactivity between bursts in burst suppression EEG.

Magnetic Resonance Imaging(MRI)

Method of biomedical imaging that utilises a powerful magnetic field to visualise the internal structure of the body.

Non-REM Sleep

A stage of sleep where no rapid eye movement occurs, also known as Quiet Sleep in infants.

Positive Rolandic Sharp Wave (PRSW)

An EEG pattern on full channel EEG that is associated with brain injuries.

Preterm

A term used to describe infants born before 37 weeks GA.

Rapid Eye Movement (REM) Sleep

A stage of sleep where rapid eye movement occurs, also known as Active Sleep in infants.

Sleep-Wake Cycle

The alternating periods of sleep and wakefulness, as shown by alternating continuous and discontinuous signals in an EEG recording.

Symmetry

The similarity in EEG behaviour between the two hemispheres of the brain.

Synchrony

The similarity of EEG recordings from both hemispheres of the brain.

Temporal Sawtooth (TS)

A "sawtooth" pattern in EEG that occurs in temporal channels.

Time-Frequency Distribution

A 3D distribution of energy as expressed in the joint time-frequency domain.

Tracé Alternant

See **Discontinuous Signal**

Tracé Discontinu

See **Burst Suppression**

White Matter

Solid components of the brain, located under the brain surface, used for connecting the grey matter (on the brain surface) and and carrying the electrical signals that connect neurons.

Chapter 1

Introduction

Preterm neonatal electroencephalogram (EEG) analysis is a growing field, as medical advances have increased the survival rate of preterm infants, while they remain in the higher risk group for neurological problems later in life. This chapter outlines some background information on EEG, what it is and how it is measured. Methods of brain monitoring are mentioned, and compared with EEG monitoring. The scope of the research as well as the structure of the thesis is included at the end of this chapter.

1.1 Objectives

The objective of this research was to investigate the use of engineering signal processing techniques in preterm neonatal EEG analysis. In particular, the goal was to develop a novel way to analyse the signal and to ultimately establish new features that can be used for neonatal neurology reseach, as well as aiding monitoring and prognosis.

1.2 What is Electroencephalography?

Electroencephalography, or EEG for short, is a non-invasive way to monitor brain activities. The signal is measured by placing electrodes on a subject's scalp. The potential difference between the electrodes, caused by electrical signals emitted by neurons in the brain, is recorded [1]. Studies have shown that the state of the brain can be deduced by studying the signals [2, 3]. Compared to other methods such as the Computed Tomography (CT) scan or Magnetic Resonance Imaging (MRI) , which provide a snap shot of the brain structure, EEG gives a continuous stream of data that indicatees the neurological state of the patient. This makes EEG more suitable for long term monitoring of subjects and for viewing changes in neurological state during the time of recording, and hence EEG provides a very useful diagnostic tool.

1.3 Why Study Premature EEG?

Premature neonatal EEG refers to EEG signals measured from infants born earlier than 40 weeks after conception. Modern medical advances have greatly increased the survival rate of these premature infants [4]. However, they remain in the high risk group for neurological conditions, such as cerebral palsy or retardation of mental development [2, 3, 4]. Medical research shows that brain damage can be stopped before it becomes irreversible [5], thus minimising the permanent damage to an infant's brain. The EEG is therefore a powerful tool to detect, or even predict, brain damage that could affect infants. Compared with imaging techniques such as MRI and Cranial Ultrasound (CUS), EEG is much more suited to continuous monitoring, as well as tracking changes and the maturation process of the infant, because of the non-intrusive nature as well as its availability. EEG is also a more economical procedure than imaging processes such as MRI. Bedside monitoring systems, such as the BRM monitoring system from BrainZ Instruments Limited, are available to neonatal intensive care units (ICUs) to monitor infants constantly, using only four electrodes [6]. This helps medical staff to monitor the progress of an infant at all times, as well as providing feedback to the clinicians on any treatments administered to the infant.

1.3.1 Statistics of Preterm Neonatal Clinical Outcomes

With medical advances and the constantly developing knowledge of neonatal care, the survival rate of infants born prematurely has been steadily increasing [7]. However, recent studies showed that preterm infants are at a higher risk of developing neurological problems later in life, such as cerebral palsy and learning disablities. The likelihood of disablity is significantly higher than that for infants born at term. Preterm infants are 5-10% more likely to be affected by cerebral palsy, and 30-50% more likely to be affected by learning disabilities later in life [8].

1.3.2 Brain Monitoring and EEG

To understand the role of EEG in brain monitoring, various methods of monitoring the brain are discussed here. The methods discussed have different advantages and weaknesses, and should be used as complementary monitoring methods to help construct a more comprehensive study of patients' neurological states.

Brain Imaging

In order to non-intrusively examine the structure of an infant's brain, different scanning methods can be used to image the internal brain structure. The most common methods are Cranial Ultrasound (CUS), Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) [8]. These methods all have their respective strengths and weaknesses.

CUS is performed using a portable device, thus eliminating the need to move the infant. Any haemorrhage under ultrasound will show up as a bright area, and it can also show the outlines of ventricles as dark areas. However, CUS is not particularly sensitive to lesions within the brain, compared with the other methods. The area of visibility using CUS is also limited compared with CT and MRI.

The CT scan uses x-ray beams to reconstruct cross-sections of an infant's brain area. It does not have the same visibility constraints as CUS, and gives a more comprehensive picture of the infant's brain structure. Like conventional x-ray scans, different substances (for example bone material, and different areas of the brain such as grey matter and white matter) appear as different shades in a CT scan. The scanning time of a CT scan is also shorter than that of an MRI scan. However,

compared with CUS, it lacks portability and requires the infant to be physically moved to the CT scanner. It also exposes the infant's brain to radiation, which may increase the infant's risk of developing long term brain damage.

MRI scanning uses the magnetic resonance of the ^{1}H atoms in water to reconstruct a cross-sectional image of the infant's brain. It is therefore sensitive to changes in the water concentration within the brain, and is thereby useful for identifying any abnormality inside the infant's brain, as well as viewing the development of grey and white matter. Unlike the CT scan, which uses X-ray beams, MRI does not expose the infant's brain to any harmful radiation. However, like the CT scan, it also lacks the portability of CUS, and the scanning time of an MRI scan is longer than that of a CT scan. New methods are being developed to improve the quality of MRI and shorten the scanning time.

Electroencephalography

Electroencephalography (EEG) is a way to monitor the brain that does not require the patient to remain still for the duration of the scan. It also provides a means of continuously monitoring the infant's brain activity, rather than taking a snapshot. This enables clinicians to monitor changes in an infant's brain function, such as reaction to treatment. Equipment used for EEG is also more portable compared with the MRI or CT methods, and operates at a lower cost. This makes EEG more ideal for tasks such as long term continuous monitoring.

To perform an EEG reading, the EEG technician places electrodes at standard positions on an infant's head (as indicated in Figure 1.1), and the electrical potentials between electrodes are recorded [1, 9]. The positions of the electrodes in the international 10-20 electrode placement are calculated proportional to the patient's head dimensions. This ensures that the electrodes detect EEG from similar regions of the brain for every patient's recording. Although the physical meaning of these waveforms remains unknown, clinicians over the years have discovered that the general behaviours of the EEG traces change as the infant grows closer to term [10]. Currently, there are general guidelines for reading infant EEG traces, and criteria to determine whether the traces are "normal" or not. These guidelines are detailed in Sections 2.1 and 2.2. Conventional EEG recordings are also accompanied by other physiological measurements such as electrooculogram (EOG) for eye movement, electrocardiogram (ECG) from the heart, and measurements indicating respiratory activities. Video

recording in-synch with EEG recording is also used in EEG interpretation.

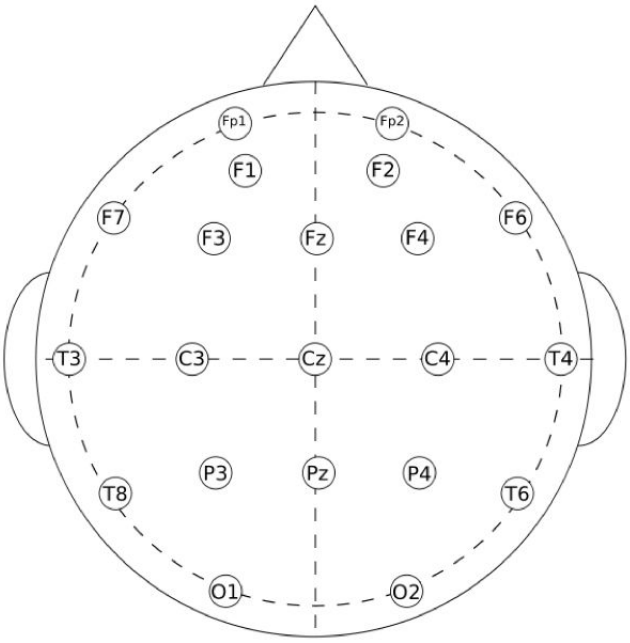

Figure 1.1: Diagram showing a birds-eye-view of the international 10-20 electrode placement adapted for infants.

To interpret EEG traces, clinicians generally examine 10 seconds of the recording at a time, and group the channels in ways that emphasise features of interest. Instead of reading the 10 second windows in detail, most clinicians opt for a "scanning" process where these 10 second windows are quickly scrolled through, to gain a general visual impression from the traces. Usually, after this stage, the clinician should have a fair idea about the state of the patient at a general level. The clinician will then return to the EEG traces and go through the recording more thoroughly, spending more time examining any areas of interest in detail.

Artifact is a problem for EEG, as muscle artifact and electrical pulses from other areas (e.g. ECG from the heart) can interfere with the EEG recording. Clinicians are trained to either ignore segments corrupted by artifacts, or to "look through" the artifacts. For example, if a segment of the EEG recording is corrupted by high frequency muscle artifact, the clinician can still read the underlying low frequency activities. Sometimes, however, it is difficult to determine whether some particular pattern is caused by artifact or has a cerebral origin. Usually, an educated guess is

formed by viewing the video recording as well as comparing the channel displaying the pattern in question with other neighbouring channels.

1.3.3 Bedside Monitoring

Conventional EEG can involve as many as 19 electrodes. The reading process is complicated and requires expert interpretation. Bedside monitoring systems are systems designed to perform "always-on" monitoring by simplifying the EEG recording. The BRM monitoring system from BrainZ Instruments uses 4 electrodes, placed on C3 and P3 on the left, and C4 and P4 on the right, as shown in Figure 1.1. This electrode arrangement focuses the area of monitoring in the central and parietal areas, since the frontal lobe in preterm infants is not developed yet. Two channels are recorded from the BRM monitors, one for each hemisphere. This minimal number of electrodes make the system more user friendly, while monitoring the two hemispheres at the same time. However, since there is only one channel per hemisphere, one loses the advantage of multiple channels, namely having the information of neighbouring channels to aid in understanding the area of the brain where the signal is generated, as well as to verify whether a pattern is caused by artifact or is cerebral in origin.

1.4 Contribution Summary

The main contribution of the work covered in this thesis is quantifying the continuity information in a preterm EEG recording. The following is a list, in order of the corresponding sections in this thesis, of the various contributions this work has made towards the study of preterm EEG.

- Analyses the clinical method of EEG interpretation and translates qualitative features into mathematical terms
- Investigates the use of time-frequency distributions in EEG signal analysis
- Evaluates methods of EEG segmentation by comparing segment boundaries with time-frequency distribution of the signal
- Presents a novel quantitative EEG continuity feature

- Quantifies continuity in neonatal EEG using principal components of EEG continuity feature
- Proposes a continuity feature line plot as an improved alternative to current aEEG display
- Presents a continuity-based maturation index for neonatal EEG
- Proposes the use of classifiers to use EEG continuity feature to improve the current definition of the continuity states

1.5 Research Scope

The scope of this research includes investigating the use of signal processing techniques in the task of preterm neonatal EEG processing: Specifically, finding features that can be correlated with maturation and/or brain injury and assist in clinical research. Seizure detection, although a popular field in engineering EEG processing, is not within the scope of the project.

1.6 Overview of Proposed System

The proposed system extracts features related to the amplitude of the preterm EEG recordings, and generates a quantified measurement of continuity and an estimate of maturation accordingly. The features can also be used as inputs for either supervised or unsupervised classifiers to produce qualitative labels.

The EEG signals used for this research have two channels: one for each hemisphere of the brain. The machine used for recording was the BrainZ Instruments BRM monitor [6], and the sampling rate was 64Hz. The conceptional age (CA) of the preterm infants at the time of recording ranged from 23 to 38 weeks. Their two-year neurological test results for language and motor skills, as well as white matter injury scoring from MRI scans at term, are available. A list of the infants and their neurological health assessments is presented in Appendix A.

In order to analyse the continuity of the signal, the signal is divided into psuedo-stationary segments using generalised likelihood ratio. This mimics the way a human mentally divides the EEG signal into blocks according to behavioural similarities.

Segmenting the EEG signal into psuedo-stationary segments also assists in further analysis of the features. Details of the segmentation method are discussed in Chapter 5. The mean absolute voltages of the segments are used to represent the amplitude of the signal, and a continuity feature related to the distribution of the amplitude is extracted from the vector of this amplitude measurement. This vector is referred to as the "amplitude vector" in this thesis and is an estimate of the signal amplitude envelope. The resultant continuity feature is the estimated mean and standard deviation of the log of the amplitude vector. This feature was chosen because it defines the statistical distribution of the amplitude vector. The feature is then used to calculate a quantitative measure for continuity. This is achieved by performing Principal Component Analysis (PCA) on the dataset. The details of the continuity feature and the quantitative measurement are presented in Chapter 6. The feature can be further analysed to correlate with maturation, and a maturation index can be calculated using the minor component from the PCA in the previous step. This is discussed in Chapter 7. Further classification using the feature can provide a qualitative label familiar to clinicians, and this is discussed in Chapter 8. Figure 1.2 shows a block diagram of the system.

1.7 Thesis Structure

Chapter 2 presents an overview of how preterm neonatal EEG is currently interpreted by clinicians. The chapter explains the various aspects that clinicians look for in neonatal EEG and how these relate to maturation and/or brain injury. Chapter 3 presents and discusses existing work in both the medical and engineering research fields regarding neonatal EEG analysis and EEG in general. The gap between the two fields is also discussed. Chapter 4 documents initial investigations carried out to understand the neonatal EEG signal using time-frequency analysis, and presents and discusses the findings. Chapter 5 discusses the task of segmentation of the EEG signal, and evaluates three existing methods for segmenting the preterm EEG signals. Chapter 6 presents a feature extraction method (using the segmentation method discussed in Chapter 5) that relates to the continuity of the EEG recording. Chapter 7 discusses how the continuity feature extracted can be correlated with maturation, and how brain injury can impact on the way maturation changes as the infant ages. Chapter 8 presents an investigation of the classification methods to be used with the continuity feature to provide the traditional qualitative labelling. Chapter 9

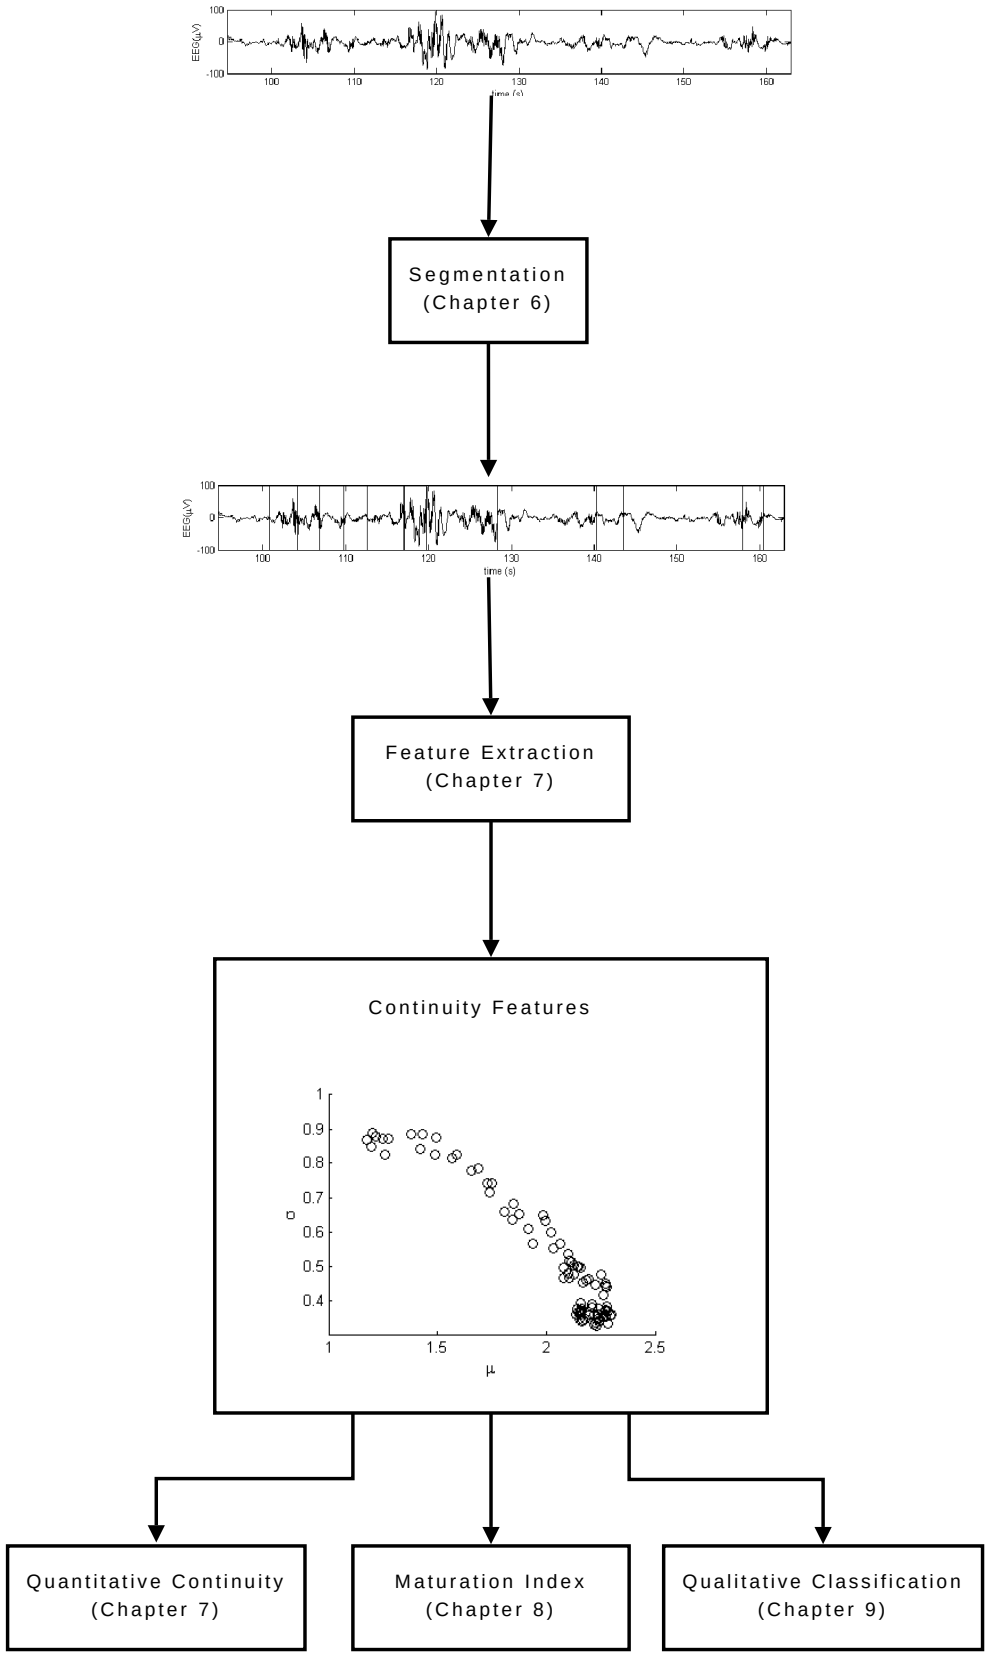

Figure 1.2: Block diagram of the proposed feature extraction algorithm

discusses results of the feature extraction and classification, comparing how they relate to existing systems as well as to current EEG continuity intepretation. The chapter also discusses the roles of both the quantified continuity measurement and the classification to translate this feature into the labels that clinicians are familiar with. Chapter 10 summarises the thesis and presents conclusions drawn from this research, and gives suggestions for further work based on the findings of this research.

Chapter 2

Current Clinical Practice for EEG Interpretation

EEG interpretation is a skill that requires years of training. Typically, clinicians learn by experience, gradually coming to recognise the tell-tale signs of maturation and abnormality in EEG signals. There are also general guidelines on the different aspects of EEG signals that should be looked for, and how these aspects are impacted by maturation and brain injury. This chapter presents some of the EEG qualities that clinicians look for, summarises the way human experts approach the task of EEG interpretation, and explores how the qualitative approach can be translated into a quantitative approach which can be automated.

2.1 Neonatal EEG and Maturation

The behaviour of EEG signals for a premature infant changes very rapidly as the infant grows closer to term. This is due to the speed of development of neurons during these critical weeks. There are a few certain aspects of the EEG that clinicians use to judge how well the EEG reflects the conceptional age (CA) of an infant. Dysmature EEG

traces, referring to EEG traces that appear to be recorded from a younger infant, may indicate neurological problems.

2.1.1 Continuity

Continuity, used in a clinical sense when describing EEG signal, refers to the variation of the EEG signal amplitude. A section of EEG signal where the envelope stays relatively constant is described as "continuous" signal. A signal consisting of periods of high amplitude signal as well as low voltage activities is referred to as "discontinuous". Continuity is usually described qualitatively. Figure 2.1 shows segments of EEG signal from the common categories of neonatal EEG continuity.

Continuous normal voltage signal refers to continuous signal where the voltage remains within the normal range and maintains a relatively constant amplitude. This pattern is the normal behaviour for near term infants. *Tracé alternant*, also referred to as discontinuous, is used to describe signal where regions of high and low amplitude can be easily identified. This pattern is normal for younger preterm infants. For very premature infants, EEG traces predominantly consist of the *tracé discontinu* pattern, appearing as alternating periods of high amplitude bursts and very low amplitude inactivity. In older infants, this pattern is referred to as "burst suppression" and is considered a sign of abnormality. Continuous low voltage refers to EEG with a relatively constant amplitude with consistent abnormally low voltage; while flat lining means that the EEG signal is almost non-existent.

As the infant grows closer to term, the amplitude of the bursts in the *tracé discontinu* pattern will decrease, and the period of the low amplitude inactivity will shorten, until around 34 weeks CA, when the EEG of an infant during wakeful periods becomes relatively continuous. Figure 2.2 shows an example of progression of EEG continuity from 26 to 30 weeks CA.

Because continuity is a subjective qualitative measurement, for this thesis, the different continuity labels need to be defined. Continuous EEG refers to EEG signal which shows relatively little variation in its amplitude, while discontinuous signal refers to signal where the variation in amplitude is noticible. Burst suppression, in this thesis, refers to both the abnormal pattern that appears in term infants and the normal pattern that appears in preterm infants referred to as *tracé discontinu*. Burst suppression can be seen as an extreme form of discontinuity, and the distinction

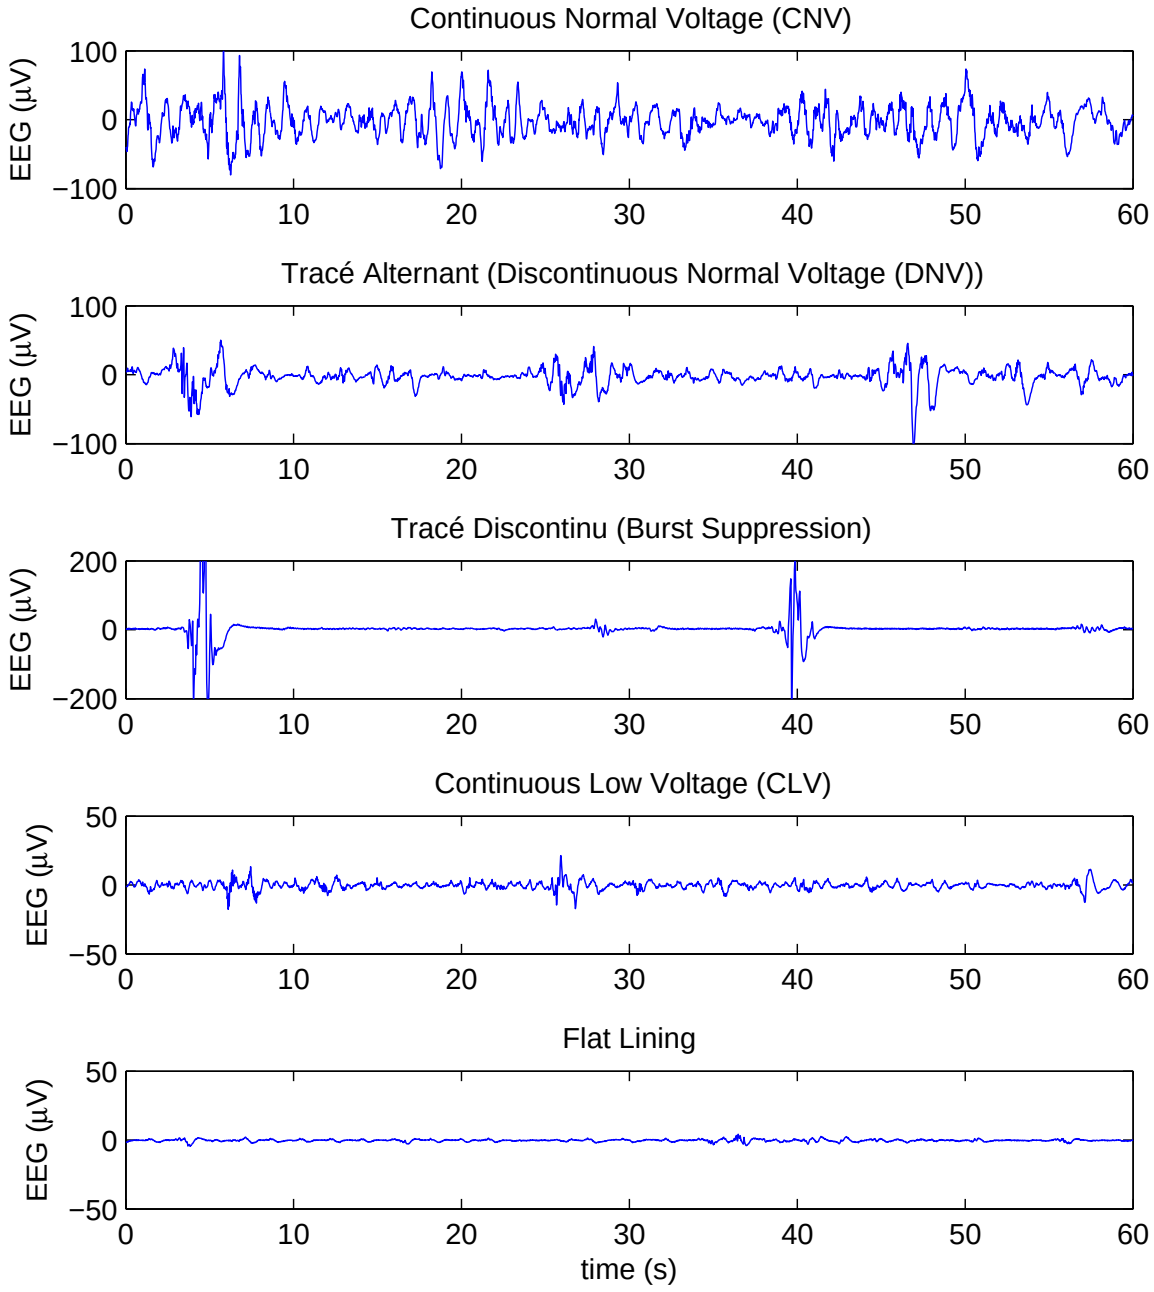

Figure 2.1: Examples of different background continuity patterns. Note the different y-axis scales.

between the two is not always obvious and relies on the context of the signal.

By looking at the continuity of the signal, one can compare the behaviour of

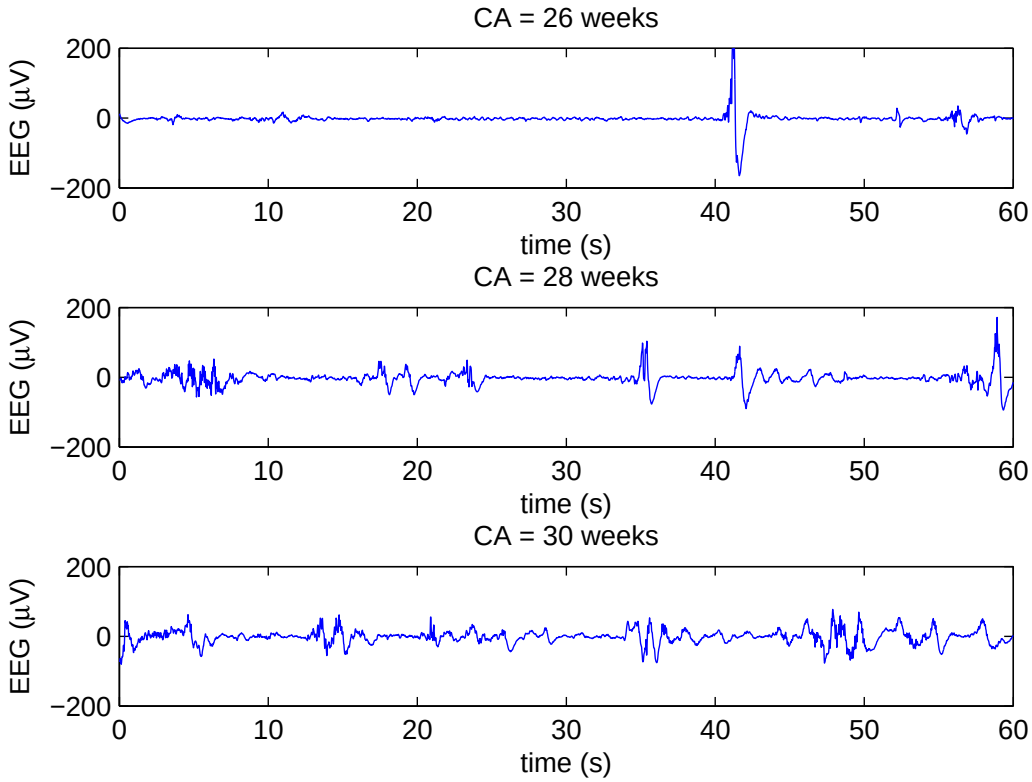

Figure 2.2: Examples of one-minute EEG segments from healthy infants with conceptional ages of 26, 28 and 30 weeks

the infant with its conceptional age. Tracings from infants that do not match the continuity criteria for their age are considered dysmature.

2.1.2 Sleep-Wake Cycle

Related to continuity is the sleep wake cycle. Starting from around 31 weeks CA, infants should display different patterns during sleep, that can be used to differentiate between the sleeping and wakeful periods. At around 33 weeks CA, rapid eye movement (REM) sleep and non-REM sleep should also be identifiable [11]. The absence of this sleep-wake cycle is considered an abnormality when the infant is over 33 weeks old. REM or active sleep refers to sleep periods where rapid eye movement can be observed. EEG recordings in this period are generally continuous. Non-REM or quiet sleep refers to sleep periods not involving rapid eye movement, and is generally discontinuous in nature. The discontinuous EEG pattern displayed during non-REM

Conceptional Age (weeks)	Continuity			Synchrony			EEG difference between awake and sleep
	Awake	Quiet sleep	Active sleep	Awake	Quiet sleep	Active sleep	
27 - 28	-	D	D	-	++++	++++	No
29 - 30	D	D	D	0	0	0	No
31 - 33	D	D	C	+	+	++	No
34 - 35	C	D	C	+++	+	+++	No
36 - 37	C	D	C	++++	++	++++	Yes
38 - 40	C	C	C	++++	+++	++++	Yes

Table 2.1: Summary of EEG behaviour during different sleep stages. Key: D - Discontinuous, C - Continuous, 0 - No Synchrony, ++++ - Total Synchrony

sleep is also referred to as *tracé alternant*.

A summary of the different behaviours of the sleep EEG patterns, from 24 weeks to term (40 weeks), is shown in table 2.1 (adapted from [11]).

The length of time an infant spends in the sleep or wakeful state is also of importance. However, this is often interrupted by caring procedures (e.g. feeding or changing) or medication (e.g. sedation given for seizures).

2.1.3 Synchrony

Synchrony of the EEG refers to whether activities that appear in one hemisphere also appear in the other hemisphere of the brain at the same time. In terms of preterm infants, there is no neural connection between the hemispheres until 32 weeks CA, therefore most patterns occurring in one hemisphere will not be reflected in the other hemisphere. After these connections are established at the age of 32 weeks CA, some level of synchrony is expected from the EEG traces [11]. The degree of synchrony is dependent on the infant's CA as well as the state of the infant (e.g. awake or REM sleep). Note also that before 28 weeks, the most prominent pattern visible is *tracé discontinu*, and during this age the high amplitude bursts occur in synchrony in the two hemispheres.

2.1.4 Other Landmark Patterns

Clinicians also look for specific patterns in the EEG tracings. Some patterns are expected for infants after a certain age, while others may indicate possible problems. One example of these additional patterns is sharp waves. Although the occasional presence of sharp waves in preterm infant EEG traces is not a sign of abnormality, the excessive and regular appearances of these waves usually indicates abnormality in the hemisphere where the waves occur. Sleep spindles, patterns where higher frequency waves are superimposed on low frequency waves that occur during sleep, are another example of landmark patterns that clinicians look for. After the infant reaches term, sleep spindles should be visible while the infant is asleep. A lack of sleep spindles may indicate neurological development problems.

2.2 Signs of Abnormal EEG

Severely abnormal EEG traces can usually be recognised easily because of their significantly different characteristics when compared with normal EEG recordings. However, abnormalities in the EEG recorded from infants with more subtle neurological problems can be a lot harder to detect. Section 2.1 outlined some of the patterns that clinicians examine to determine whether an EEG recording is considered normal for the infant's CA. If a recording does not match the infant's CA it can be considered abnormal. This section considers some other common patterns that are considered abnormal.

2.2.1 Symmetry

Symmetry refers to the similarity of the two hemispheres. This differs from synchrony as it does not require the activities to occur simultaneously in both hemispheres, but rather concerns the general pattern seen in the two hemispheres, and how they compare with each other. If a particular pattern (e.g. visible slow delta waves) appears in one hemisphere and not the other, the EEG trace is considered asymmetric. Symmetry is generally expected from EEG traces of all infants regardless of age or state. Asymmetry usually indicates brain injury and/or delayed development in one of the hemispheres.

2.2.2 Burst Suppression

Burst suppression is a common pattern occuring in preterm infants, consisting of bursts of high energy activity, separated by periods of low amplitude inactivity. For very premature infants (< 28 weeks), this pattern is known as *tracé discontinu*, and is normal for infants of that age [11]. However, the inactivity periods should gradually become higher in amplitude with a shorter duration, and eventually the *tracé discontinu* pattern should become more continuous. When the infant is of an age where the wakeful EEG should be continuous (around 34 weeks CA), yet still exhibits the burst suppression pattern, it is treated as an abnormality. Usually, for these infants, the suppression periods have even lower amplitudes than the *tracé discontinu* pattern found in younger infants. The suppression periods generally also last longer than the normal *tracé discontinu* pattern found in very young preterm infants.

In order to determine whether the burst suppression pattern is normal, the CA of the infant, as well as the amplitude and duration of burst and suppression periods should be taken into consideration.

2.2.3 Seizure

Seizures in infants usually appear as a series of regular slow waves which evolve though the duration of the seizure. The frequency and the shape of the wave usually changes during the seizure. Any seizure is considered an abnormality; however some infants recover from seizures on their own without medication, and without any seizure recurrence or adverse consequences.

Seizures can be episodic and isolated, episodic and repeated, or continuous (also known as *status epilepticus*). Preterm infants experiencing a seizure episode often only exhibit very subtle or even no clinical signs. Therefore, the use of EEG can be very helpful in detecting these subclinical seizures.

2.2.4 Dysmature EEG

Dysmature EEG refers to EEG recording that shows no obvious abnormalities, but exhibits the behaviour of EEG recorded from a younger infant. For example, if an EEG recording of a 33 week CA infant shows behaviour appropriate for a 28 week infant, the recording is said to be dysmature. Although dysmature EEG is considered

abnormal, it does not necessarily imply brain injury. Dysmature patterns can also be caused by delayed brain development.

2.3 Translating Clinical Knowledge into Mathematics

The main challenge to writing an algorithm to mimic the way humans interpret EEG is the qualitative nature of EEG interpretation. The current medical guidelines for EEG interpretation, while comprehensive, rely on experience, the ability to tell what artifact looks like, and reading the pattern in context. Certain properties (e.g. synchrony) are only defined qualitatively. In order to write an algorithm to analyse EEG signals, one must define quantitative terms that can represent these qualitative features used, and interpret these quantitative terms sensibly.

2.3.1 Normal Range

Preterm neonatal EEG can vary widely, and, compared with adult EEG, can exhibit a relatively wide range of behaviours that can be considered normal. In order to determine what can be considered normal, a range of EEG from minimal risk infants should be analysed, and compared against EEG for infants with adverse clinical outcomes.

2.3.2 Defining Continuity

Currently, clinicians measure the continuity of signals qualitatively. A signal is defined as continuous if the envelope of the signal remains relatively level, and labelled as discontinuous otherwise. Some attempts have been made to quantify continuity by setting a threshold and measuring the percentage of time the signal is above the threshold. While this gives an idea of how continuous the signal is, it is still a very rough measure.

Using segmentation algorithms one can segment the EEG trace into pseudo-stationary segments. The envelope of the EEG signal should be relatively even for the duration of the segment. The energy of the segment or its averaged absolute voltage

can be used with the length of the segment to give a more detailed measurement of the continuity of the signal as a whole.

2.3.3 Defining Synchrony

As described in Section 2.1.3, synchrony in the two channel system used for this project refers to whether activities occur in both hemispheres of the brain at the same time. This can be achieved by finding the correlation between the two channels. There are several signal processing measurements established to measure the synchrony of two signals, such as the use of correlation coefficient and cepstrum. One can apply different methods of synchrony measurement and decide which gives the most desirable results for EEG.

2.3.4 Defining Symmetry

Unlike synchrony, symmetry is difficult to define mathematically, since it is a subjective observation and involves the general behaviour of the signal rather than local behaviours. Therefore, symmetry should be defined by using a suitable window size, and comparing the two channels within this window space. The length of the window should be large enough to show the general behaviour of the signal rather than the local behaviours, and the state of the signal should also be taken into account when comparing the two hemispheres. For example, the behaviour of the sleep and wakeful states of the infant should be compared separately.

The features used for this comparison also need consideration. One obvious candidate is comparing the spectrum of the two channels in a predefined window. However, since clinicians also look for patterns that occur in the EEG tracing, it may be more appropriate to compare the general parameters obtained in the two hemispheres during the same state of the infant's EEG recording. For example, once the state of the infant is determined, the parametric distribution of the left and right EEG channels during each state can be compared with each other, and the symmetry of the EEG can be determined by looking at the how similar these distributions are.

2.3.5 EEG Pattern Recognition

There are two types of pattern recognition that clinicians perform during the EEG interpretation: local pattern recognition and global pattern recognition.

Patterns such as sharp wave and delta brushes are localised patterns. These patterns can be recognised by analysing the parameters of pseudo-stationary segments, and do not require the information from neighbouring segments. In order to automatically recognise these patterns, the parameters of segments known to contain patterns of interest can be analysed. Classification algorithms, such as artificial neural networks, can then be utilised to decide whether any segment contains a pattern of interest by analysing said parameters.

Global patterns are patterns such as sleep wake cycle and burst suppression. These patterns require the parameters of the segments as well as neighbouring segments to be analysed. These patterns can be seen as a change of state of the EEG recordings. To recognise these types of patterns, a classification system that utilises state information can be used. For example, the burst suppression pattern can be broken down into burst state and suppression state. The burst suppression pattern can be defined as alternating burst segments and suppression segments.

2.3.6 Artifacts

Artifact rejection is important in EEG analysis. Since some muscle artifact shares the same spectrum range as EEG signal, a simple band pass filter is not enough to filter the artifact. Since the human skull acts as a low pass filter and dampens the higher frequencies, any high frequency components with abnormally high intensity that registered in the EEG recording cannot be cerebral in origin. One can analyse segments of EEG with no muscle artifacts to study the normal behaviour of frequencies in EEG signals. Afterwards, signals will be processed using time-frequency analysis to show the different frequency components against time. Each component can then be compared with the profile obtained from artifact-free EEG signals to determine whether the component has too much energy for its frequency. Components with significantly more energy than components with the same frequency from the artifact-free EEG signal can then be assumed to be artifacts. The advantage of analysing the components rather than the signal as a whole is that the underlying signal can still be processed. This can be compared to the way clinicians mentally

ignore the high frequency artifact activities and analyse the underlying low frequency trend.

Figures 2.3 and 2.4 show the time frequency distribution (TFD) of a segment of EEG with mechanical and high frequency muscle artifacts, respectively. Mechanical noise can be caused by a variety of noise sources, including electromagnetic interference from the main power source (usually 50 or 60 Hz) and mechanical ventilator. The mechanical artifact appears as a straight line in the high frequency region, with the second harmonic also visible at twice the fundamental frequency. The muscle artifact appears as a random frequency component in the higher frequency region. Figure 2.5 shows the TFD of an EEG segment without noticeable artifact. When the high frequency components from Figure 2.4 and 2.5 are compared with each other, one can see that the muscle artifact in the higher frequency region exhibits more energy than the high frequency signal of cerebral origin.

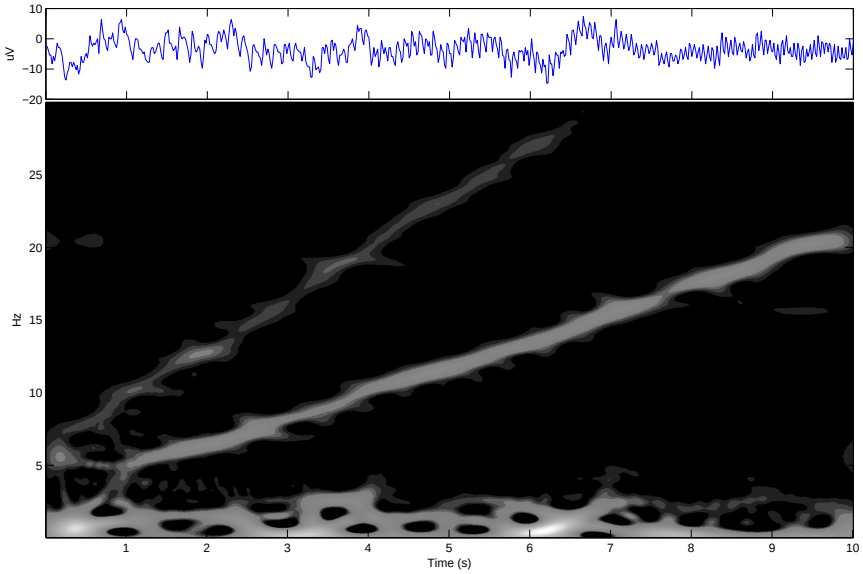

Figure 2.3: Diagram showing a segment of EEG recording with mechanical noise and its time-frequency energy distribution. Note the linear chirp above 5Hz.

2.3.7 The Big Picture Approach

Isolated incidents seen in the EEG tracings are not as important as the global behaviour of the EEG tracings. One example would be sharp waves in preterm infants,

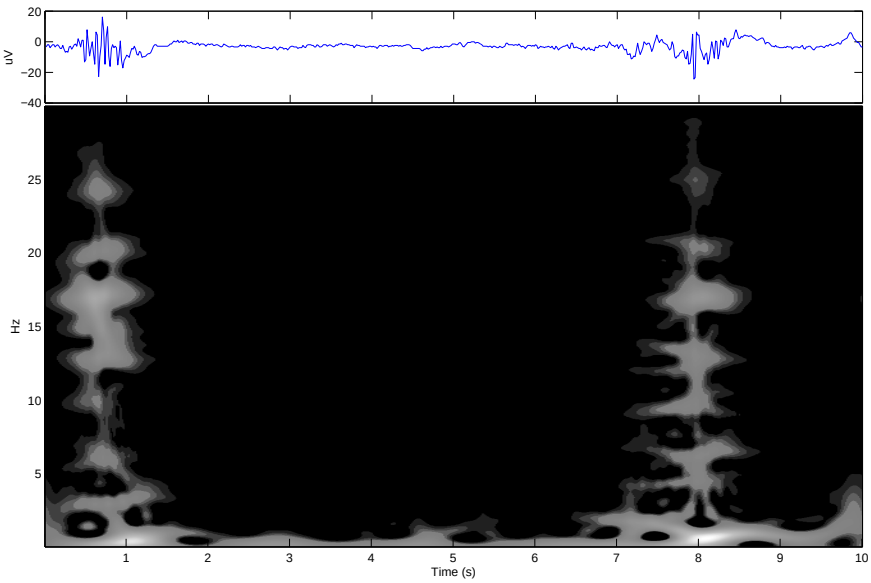

Figure 2.4: Diagram showing a segment of EEG recording with muscle artifact and its time-frequency energy distribution. Note the high frequency muscle noise at 1s and 8s.

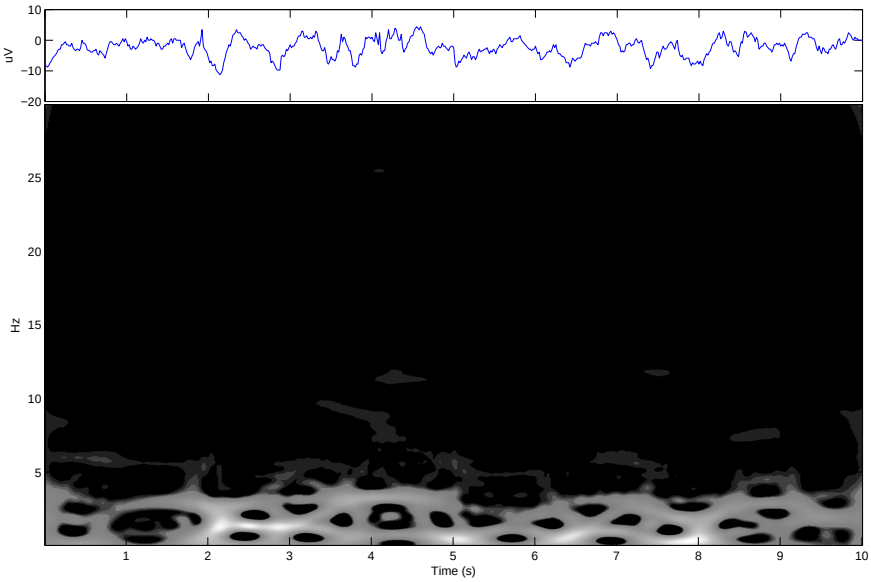

Figure 2.5: Diagram showing a segment of artifact-free EEG recording and its time-frequency energy distribution. Note the lack of high frequency energy above around 5Hz.

where an isolated sharp wave does not generally indicate abnormality, but a consistent occurrence of these sharp waves may indicate problems in that particular area of the brain. Patterns are also read in context. Some behaviour, such as the occurrence of certain frequencies or patterns, may be expected in some specific states (e.g. REM sleep) but not at other times.

In order to provide the EEG signal with context, some sort of state mechanism can be utilised to keep a record of the state of the signal. Variables can also be used to track other information, such as the number of occurrences of spikes, and these can be updated as more data are analysed. The state information and monitoring variables should be taken into account during the data analysis stage.

2.4 Discussion

The first consideration in determining the direction of an automated approach to analysing neonatal EEG, is to figure out which of the above criteria for clinical EEG interpretation is the best criterion for maturation, as well as whether implementing one criterion will assist future work on a quantified method of EEG analysis. Since the goal is to give a better representation of maturation, continuity is the logical choice for the first step. Not only is continuity directly related to maturation, it is also related to the sleep-wake cycle, a critical milestone in the neurological development of neonates. Continuity can also act as a way to provide context for further analysis. Since the behaviour of EEG signals can change dramatically during recording, knowing the continuity of the signal helps to provide some framework when deciding areas for comparison between different recordings.

Defining continuity can also be of assistance in automating the monitoring of EEG synchrony, since the degree of synchrony in normal neonatal EEG is dependent on the background continuity state as well as the age of the infant.

2.5 Summary

In this chapter the current clinical conventions of neonatal EEG interpretation were defined. In EEG analysis, clinicians look for signs of maturation and any indications of abnormality. The signs of maturation include an increasingly continuous

background state, the presence of sleep-wake cycles after 31 weeks CA, and changes in synchrony between the two hemispheres of the brain. The signs of abnormality include dysmature signal, burst suppression in older infants, and asymmetry. Continuity, synchrony and symmetry are three major criteria for the measurement of maturation and normality, with other signs either a measurement of the three criteria in context (e.g. sleep-wake cycle detection involves detecting alternating continuous and discontinuous segments of EEG signal) or detection of specific patterns. Since context is important in EEG recognition, detecting the background continuity state of the EEG is crucial. The continuity of the signal should be the first aspect considered, as it not only is an important aspect of maturation and normality, but it also provides the context needed for further analysis.

Chapter 3

Literature Review

Research on neonatal EEG analysis can be divided into two rough categories: medical studies on the prognostic and diagnostic value of neonatal EEG, and engineering research on pattern recognition in EEG. This chapter reviews some of the studies in both categories. In medical studies, the focus is on two main areas: signs of maturation and signs of abnormal development. In engineering studies, the focus is on various applications such as seizure detection, hypoxia detection, sleep detection and segmentation methods. The chapter conludes by comparing the two categories of work, and identitfying areas where engineering signal processing techniques can address medical research needs.

3.1 Medical EEG Studies

Neonatal EEG is gaining attention in the field of medicine, since advances in medicine have increased the survival rates of premature infants [4]. As these premature infants still remain in the high risk group for developing adverse neurological outcomes [2, 3, 4], e.g. cerebral palsy, retardation of mental development, more research is needed to identify infants with abnormal development, and intercept problems before

they become irreversible, when the window of opportunity is still present [5]. The conventions of preterm EEG interpretation by clinicians are discussed in detail in [10, 11, 1] and summarised in Chapter 2. In this section of the literature review, medical studies of preterm EEG feature analysis are reviewed and discussed. There are two main areas of interest: signs of maturation, and signs of abnormal development.

3.1.1 Signs of Maturation

Victor *et al.* conducted a study on EEG of very premature infants for the first four days after birth [12]. Using spectral analysis (Fourier transform), the absolute power (AP) and relative power (RP) of the frequency bands were noted, as well as the inter-burst interval (IBI). Artifact removal was carried out manually to reject sections of recording corrupted by noise. They found an increase of RP in delta (from 68% to 81%) from day 1 to day 4, and that the IBI shortens (from 14 s to 8 s) from day 1 to day 3, without any significant difference between day 3 and day 4. This study is one of the few that utilise frequency analysis, and combine it with features from the time domain (namely the IBI). The IBI can be viewed as a way to analyse the continuity of the signal, since more discontinuous signals have longer inactive periods between bursts. It is worth noting that they found no difference in RP of delta bands and IBI between healthy and unhealthy infants.

Biagioni *et al.* measured some maturation signs for preterm infants in their first two weeks of life, in order to discern any correlation between these features and neurological outcomes [13]. The maturation signs they measured included the minimum burst period, maximum IBI, IBI amplitude, maximum amplitude of delta wave in bursts, max amplitude of 8-20 Hz activities in burst, maximum amplitude of a "saw-tooth" pattern in the temporal regions called temporal sawtooth (TS), incidence of 8-20 Hz activities, and incidence of TS in bursts. They found that for infants with a higher gestational age (GA), the measured minimum burst period is longer than those infants with lower GA. The measured maximum interval period, on the other hand, is shorter in infants with a higher GA compared to infants with a lower GA. The amplitude of delta waves in bursts decreased with GA. Amplitudes of 8-20 Hz activities increased with GA. Amplitude of TS was high for 27-30 week GA infants and decreased afterward. The number of 8-20 Hz activity instances increased with GA. Incidence of TS was high for 27-30 weeks, and decreased afterwards, becoming

rare in 33-34 week infants. The study only analysed discontinous data, since not all infants exhibit continuous patterns, and therefore the maturation progress of continuous segments was not examined. The burst period, IBI and IBI amplitude are all related to continuity. The maximum amplitude of activities in various frequencies (identified visually) can be seen as localised spectral information. This study confirms the idea that continuity is related to maturation, and found that the continuity of the discontinuous section of the signal changes with maturation. Temporal sawtooth pattern requires a conventional standard EEG recording which monitors the temporal region of the scalp (T3-T8 on the left and T4-T6 on the right in the 10-20 international EEG electrode placement as shown in Figure 1.1). This is not always suitable for bedside monitoring measurement, as bedside monitoring systems use only the minimum number of electrodes specified, normally placed closer to the central region for better coverage. However, infants can be placed on a full channel recording for a limited time to observe the number of occurrences and the amplitude of the TS pattern. Since the study only looked at the maximum values of the continuity related features, only a global description of the continuity was given and no local variation of continuity was recorded.

West *et al.* conducted a study on preterm neonatal EEG using available quantitative measurements [14]. The study included preterm neonates with less than 32 weeks GA and analysed EEG recorded during their first week after birth. The infants included in the study were chosen on the basis of no evidence of white matter injury and intraventricular haemorrhage grades of 1 or lower, indicating no or very minimal brain injury. The quantitative measurement the study used targeted continuity, amplitude and frequency contents. Using a set of thresholds (10μV, 25μV, 50μV, and 100μV), the proportion of the EEG signal in a 60 minute epoch that lay under the thresholds was recorded as a measurement of continuity. The median amplitude of the 60 minute recording was also recorded, as well as the the spectral edge frequency (the 90^{th} percentile in the frequency spectrum). The median of the measurements was correlated with age. It was found that there is an increase in continuity measurement and amplitude, accompanied by a decrease in the spectral edge frequency. The decrease of spectral edge frequency may at first appear to contradict Biagioni's finding of an increase of activities in 8 to 20 Hz activities. However, considering that the spike-like high amplitude bursts decrease in amplitude as the infant matures, this affects the amount of energy in the high frequencies. Thus the lower spectral edge frequency should be viewed as an indication of the decrease of amplitude in the bursts.

The continuity measurements in this study were carried out with intensity thresholds, and while this can indicate the continuity state of the signal, it uses arbitary thresholds and only gives an idea of the proportion of the signal that falls below said thresholds.

Selton *et al.* studied the maturation of very premature neonates with normal two year neurological follow ups [4]. Criteria included in this study were IBI, burst duration, burst morphology, synchrony, symmetry, spatial organisation, and sleep states. The infants included in this study were aged between 25 to 28 weeks CA, and 17 infants with a normal clinical follow-up were included. The study compared the findings with existing literature on preterm neonatal EEG development and found them consistent. The paper placed an emphasis on the morphology of the bursts and focused on discussing the changes occuring in the burst period during 25-28 weeks. The IBI were discussed, and it was noted that there is no consensus for what values should be used as the threshold for the determination and definition of the interburst period. The paper discussed the values obtained from their studies (IBI between 1.0 to 46 seconds). A prolonged burst period as the infants matured between 25 to 28 weeks was also noted. Synchrony between the two hemispheres is present at this age. The authors commented that with the help of computerised analysis the local discontinuities can be analysed, rather than the global indicators, such as maxmimum IBI or burst amplitude. The paper concluded that EEG recordings of very premature infants display consistent patterns that can be used for prognostic purposes. The paper did not compare how the patterns exhibited by infants with abnormalities differ from the patterns discussed.

Burdjalov *et al.* studied the effect that maturation has on EEG recordings using cerebral function monitoring (CFM) [15]. CFM uses a single channel amplitude-integrated EEG (aEEG) that is derived from a 2-channel EEG recording to aid monitoring long periods of EEG recordings: It is, in essence, the envelope of the signal plotted on a compressed time scale (see appendix B for the definition and plotting convention of aEEG). In this study the authors attempted to quantify the assessment of EEG maturation by focusing on four different areas and giving a score for each of these aspects of aEEG. The four areas for scoring are continuity, cycling, amplitude of lower border and bandwidth span. Note that the term "bandwidth", when used in the context of aEEG, differs from the normal sense of the term in signal processing, and refers to the difference between the bottom border and the top border of the aEEG recording. The scoring was based on a combination of qualitative and quan-

titative values. Table 3.1 gives an overview of the scoring system [15]. The study scored aEEG recordings of 30 healthy infants ranging from 24 to 39 weeks GA. The scoring system scores four difference aspects of aEEG related to maturation. The sum of the four scores gives a total ranged from 0 (most immature) to 13 (most mature). The study showed that as the preterm infant grows towards term, the scoring on their aEEG increases. They also found that the areas most sensitive to maturation are the presence of cycles, continuity and the aEEG bandwidth. No mention was made of how the scores were determined on areas of continuity and cycling, besides relative terms. The scoring system attempted to quantify a maturation measurement of aEEG with some success but did not quantify some of the criteria used.

Score	Continuity	Cycling	Amplitude of Lower Border	Bandwidth Span and Amplitude of Lower Border
0	Discontinuous	None	Severely depressed ($< 3\mu$V)	Very depressed: low span ($\leq 15\mu$V) and low voltage (5μV)
1	Somewhat continuous	Waves first appear	Somewhat depressed ($3 - 5\mu$V)	Very immature: high span ($> 20\mu$V) or moderate span ($15 - 20\mu$V) and low voltage (5μV)
2	Continuous	Not definite, somewhat cycling	Elevated ($> 5\mu$V)	Immature: high span ($> 20\mu$V) and high voltage ($> 5\mu$V)
3		Definite cycling, but interrupted		Maturing: moderate span ($15 - 20\mu$V) and high voltage ($> 5\mu$V)
4		Definite cycling, noninterrupted		Mature: low span ($< 15\mu$V) and high voltage ($> 5\mu$V)
5		Regular and Mature cycling		

Table 3.1: Summary of the aEEG scoring system used with CFM proposed by Burdjalov *et al.*

3.1.2 Signs of Abnormal Development

Hellström-Wastas *et al.* published a study on the importance of EEG in brain injury of preterm neonates [16]. The paper discussed how the EEGs of preterm neonates with acute and chronic brain injury differ from the EEG of infants with healthy clinical outcomes. They concluded that signs pointing to abnormalities include depressed background activities, seizure, delayed maturation, the presence of Positive Rolandic Sharp Waves (PRSW) and changes in sleep cycles. Of the mentioned markers of abnormal neonatal EEG, delayed maturation, depressed background activities and changes in sleep cycles are related to continuity of the EEG signal. Delayed maturation, in particular, is more subtle compared to the other markers discussed. The paper suggested the use of aEEG to continuously monitor the infant, to complement the standard conventional EEG monitoring and help detect seizures and sleep cycles.

Marret *et al.* showed the prognostic value of neonatal EEG for infants less than 33 weeks of gestation age in [2]. The main features analysed were presence of PRSW, electrocephalographic seizures, and background abnormalities appearing in the EEG recordings. PRSW occur more frequently in infants that survived with motor development problems than in those without. However, their presence does not appear to be a precursor for infant death. PRSW occurring more than twice per minute appears more frequently in infants who did not survive than in those who did. It was also found that background abnormalities occurred more frequently in infants who did not survive than in those who did. Of the surviving infants, abnormalities occurred more frequently in those with adverse neurological outcomes. Although the study shows that EEG can be used as a prognostic tool in neonates, the features used for this study are very subjective. Features like PRSW and background abnormalities are very loosely defined, and different clinicians may categorise a signal differently. Identifying abnormalities in neonatal EEG also takes years of training, and is performed by EEG experts manually. This makes abnormal EEG incidents useful for medical research, but unsuitable to be adapted for automated EEG monitoring, unless an expert classifier is developed for the task. While this study served as a good example of neonatal EEG containing information that may lead to a neurological prognosis, the use of background abnormalities as features is difficult to automate.

Biagioni *et al.* published their study on discontinuous EEG patterns in full-term infants with hypoxic-ischaemic encephalopathy (HIE) [17] . They studied the minimum burst period, the maximum IBI, the interval amplitude, maximum amplitude

of delta components during bursts, maximum amplitude of higher frequency components (8-22 Hz) during bursts, and incidents of other EEG abnormality during the recording. They found that high maximum IBI and short minimum burst periods indicate poor outcomes, or a high HIE grade (measure of neurological damage due to HIE, with a high score signifying more severe damage). This means EEG recordings of infants affected by HIE are more discontinuous than those recorded from healthy infants. The authors suggested that these parameters can help define the degree of discontinuity better than the labels commonly used, namely the qualitative labels of "burst suppression" and "non-burst suppression" periods. The results supported findings from the studies discussed in the previous section, that continuity is related to maturation and can be used for prognostic purposes. This study also included incidents of abnormal EEG transient besides the continuity related features.

Watanabe *et al.* published a study on using neonatal EEG as a tool to assess brain damage [3]. They analysed abnormal neonatal EEG signals and categorised them as being either acute or chronic. Using qualitative features based on continuity, frequency and voltage, different grades were assigned which allowed deduction of the time of occurrence for any injury. They concluded that signs such as an increase in discontinuous epochs, attenuation of higher frequency and low voltage, and dysmature signals and disorganised signals (e.g. PRSW, mechanical brushes) can be used to grade the degree of abnormality of the signal. Their work agrees with [17], that increasing discontinuity (i.e. prolonged inter-burst period and/or short burst period) indicated abnormality. However, the study was qualitative, which means the subjectivity of different clinicians may lead to different categorisation of signals.

3.2 Engineering Projects

While the medical works focus on correlating features with maturation or clinical outcomes, engineering research, in general, focuses on the task of pattern recognition. Like most other pattern recognition tasks, the problem is to identify useful features that can be used in a classifier. The choice of features and the classification method varies depending on the application. EEG applications that involve pattern recognition include recognition of seizure patterns, HIE induced burst suppression, sleep detection, and segmentation of multi-channel EEG signals. This section of the literature review will focus on a few common applications that utilise EEG analysis

to aid clinicians to make better prognoses.

3.2.1 Seizure Detection

Seizures occur when neurons in the brain start firing in a steady rhythm continuously. In infants, seizures can occur without any clinical signs, besides an abnormal EEG tracing.

Boashash *et al.* have published a series of works on using joint time-frequency analysis to investigate neonatal seizure detection techniques [18, 19]. To aid this detection, Boashash and Mesbah have came up with a set of template patterns for the shape of seizures [18]. This set of six patterns models the changes within the components of the frequency domain with respect to time. The set consists of four seizure related patterns, and two non-seizure related patterns which are also of interest. Each of these patterns has a quasilinear instantaneous frequency (IF) law associated with it. Boashash and Mesbah proposed to use these patterns as masks and perform 2D correlation in order to detect any seizure activity [19]. The results presented suggest a good recognition rate (99.1% detection rate and 0.4% false alarm rate). However, the study was carried out using only simulated data, and not actual recorded neonatal EEG data.

Zarjam *et al.* presented a method to detect seizures using the idea of time-frequency divergence [20]. The idea behind this method was that the time-frequency distance between the EEG segment of interest and the patient's background EEG would indicate whether the patient was having a seizure. The study used three different reduced interference distributions (RIDs), namely Choi-Williams, modified B-Distribution, and spectrogram. Three divergence measures were investigated: they were Kullback-Leibler, Jensen difference, and Rényi divergence. The neonatal EEG signals were divided into segments of 6 seconds duration. The time-frequency distance between the TFD of the segment and background was calculated and distances above a predetermined threshold were seen as an indicator of seizure. Based on the seizure detection rate (SDR) and false alarm rate (FAR), the Kullback-Leibler divergence was seen as the most reliable divergence measure, and the modified B-distribution the most reliable TFD. The SDR and FAR with the optimal combination were 96.73% and 4.55% respectively. The Rényi divergence gave a poor performance in terms of SDR, but gave the best FAR (averaged at 1.77%). The drawback of this method is that it requires a sample of background EEG to detect the seizure, and the paper

suggested that the background EEG selection was performed manually at this stage. The background EEG of a patient would also need to be recorded and updated frequently, as the brains of neonates develop at a rapid rate and EEG patterns change significantly during the maturation process.

Saab and Gotman proposed a method of detecting the onset of seizures in adults using wavelet decomposition as a tool [21]. The idea was to use the coefficients from wavelet decomposition as features, and apply the Bayes formula to estimate the probability that the EEG signals contained seizure activity. The method was designed to detect seizures in adults. The detection rate was around 77%, with a detection delay of around 10 seconds. The detection rate was less satisfactory compared to some of the other methods reviewed here, and since it was designed mainly for adults, the possibility of using similar methods for neonates has yet to be investigated.

The use of artificial neural networks (ANN) is also popular in the task of seizure detection. Özdamar *et al.* proposed a method of identifying seizures by training an ANN to detect spikes within the EEG [22, 23, 24]. In this system, 16 EEG channels were needed and the signal was processed in epochs 5 seconds long. The ANN consisted of two layers: the first layer was responsible for detecting any spikes in the 16 channels, and the second layer combined the results from the first layer to come up with a final result. This system was implemented into a microcomputer system utilising a digital signal processor. The accuracy of the system was about 83%, with the false detection rate at around 20% [23]. This seizure detection method was less accurate than other systems presented. The system was designed for adult seizure detection, and therefore by detecting spikes alone it may be sufficient for the task. However in terms of neonatal seizure detection, a more sophisticated system may be required.

Kalayci *et al.* and Zarjam *et al.* have developed systems which use wavelet coefficients as inputs to an ANN, and hence the systems include information from the neonatal EEG signals in the time-scale domain [25, 26, 27, 28].

In the studies published by Kalayci *et al.*, they used 16 coefficients as the input to a three-layer perceptron network, trained using the back propagation method. The number of hidden neurons was decided using cross validation. Essentially, the network detected spike events rather than seizures, and had an accuracy of around 90% [26]. Like the system proposed by Özdamar *et al.*, this system was designed for adult seizure detection, and therefore focused on spike detection instead of seizure

detection.

Zarjam *et al.* proposed a system specifically designed for seizure detection in neonates. Because the features were pre-selected to include the maximum amount of information, the number of inputs to the system could be reduced [27, 28]. Similar to Kalayci *et al.*, the neural network they chose was a three-layer feed-forward preceptron network, trained with back propagation. The system utilised Mutual Information Evaluation Function (MIEF) [28], and 10 features were selected to be processed by the ANN. The system accuracy was 96.35% with a false detection rate of 6.2%. With only four features selected, accuracy was found to be around 93.5%. The elimination of less relevant features simplified the complexity of the system by reducing the number of inputs, with only a slight drop in accuracy. This shows that the inclusion of less relevant features does not significantly increase the accuracy of the system, and therefore the selected inputs should only include those most relevant to the task.

3.2.2 Hypoxia Detection

Hypoxia is a condition where a reduction occurs in the oxygen supply to the brain. As the brain is starved of oxygen, the EEG behaviour changes. Early detection of a change in brain behaviour can help clinicians to start treatment early and prevent or minimise permanent damage.

Hoyer *et al.* investigated the way EEG behaviour changes at the onset of hypoxia or ischemia (a reduction of bloody supply to the brain), both of which relate to a reduction in oxgyen supply to the brain [29]. Newborn piglet EEGs were used in the study, as the condition of the subject and the oxygen supply to the brain can be monitored and controlled. The piglets were anaesthetised, immobilised, and artifically ventillated. An arterial blood controller was used to managed the amount of blood flow into the brain of the piglet. Cortical EEG (EEG recorded from the cortex instead of through the scalp) was used, and the aim was to train an Artifical Neural Network (ANN) to recognise the onset of a change in brain activity. A band power estimator was used, where the power of frequency bands was estimated using the fast Fourier transform (FFT). These powers were used as inputs to the ANN in a supervised learning system. The system peformed well, with a correct recognition rate of 70% or better in various tests. Worth noting is the system adaptively learnt the ranges of the frequency bands, and it was discovered that the learnt frequency band ranges produced better recognition rates than the traditional frequency bands

in EEG analysis (e.g. delta band for 0-4 Hz data, theta band for 4-8 Hz). This raises questions about the use of conventional frequency bands used in EEG analysis, esepcially when applied to neonatal EEGs. Research should be carried out to identify bands of interest and group frequency ranges according to these, rather than using the conventional frequency bands used for adult EEG analysis.

Goel *et al.* analysed the hypoxia onset pattern by finding frequency information of piglet EEG data [30]. Using autoregressive modelling, the frequency content of the signal was estimated. The authors divided the spectrum into three different bands: "low frequency" for 1.0-5.5 Hz, "medium frequency" for 9.0-14.0 Hz and "high frequency" for 18.0 - 21.0Hz. They measured changes in power in the three bands after hypoxia. A new feature was proposed called the mean normalised separation, which is derived from the power in the three frequency bands. They found a correlation between the clinical outcomes of the piglets, and the mean normalised separation measurement.

Löfgren *et al.* also studied early detection of hypoxia by analysing EEG of piglets exposed to hypoxia [31]. They proposed using a spectral distance measure inspired by Itakura distances [32], a spectral distance measurement originally used for speech processing, extended to apply to the ARMA model. This uses frequency domain data so the spectral differences are measured as a result. References were needed to establish the spectral contents of normal EEG background, and EEG signal before the onset of hypoxia was used for reference purposes. It was found that this new spectral distance measurement is better than using the Itakura distances, shown previously to be an efficient measurement for the application of hypoxia detection. However, when compared with other measurements, an entropy based algorithm was shown to outperform the proposed algorithm.

In all three studies, piglet EEG data was used, where the medical condition of the piglets could be monitored and controlled. Piglets were immobilised to minimise the effect of muscle or movement artifact. Cortical EEG was also often used as it eliminates the low-pass filter effect of the skull. In a clinical situation, infants will not be immobilised and care procedures will also affect the EEG recording. The effect of these factors on the algorithms is yet to be explored. Some of the algorithms also rely on a reference signal representing normal background EEG activity. Without a human expert to determine what section of the signal is considered normal background EEG, the algorithm would have to also be able to select a section for reference purposes.

3.2.3 EEG Summary by Segment Clustering

Agarwal and Gotman proposed to use Nonlinear Energy Operator (NLEO) as the basis for a segmentation algorithm [33, 34, 35]. The idea was to segment the EEG data into segments where the contents would be quasi-stationary, and grouping them into clusters exhibiting similar patterns. The derivation of the NLEO method can be found in [34]. The main idea was to estimate the energy content of the signal, and detect any abrupt changes in the energy. The points where changes occurred would be selected as the places to segment the signals. These signals were then classified into clusters using three features extracted from them. The features were indications of the amplitude, the dominant rhythm, and the frequency-weighted energy (FWE). Using k-means clustering, the segments were classified into different groups with similar features. This method allowed long recording sections to be compressed into a one- or two-page summary. This means that time required for an EEG technician to review a record could be dramatically decreased. While there were still issues to be solved, such as seed selection (i.e. the selection of initial cluster centres for the k-means algorithm) and outlier data, the method had the potential to help EEG technicians by providing an overview of the EEG data for long sessions. However, because the features are not sensitive to the shape of the EEG wave, specific localised markers may be lost or clustered with similar segments without the presence of the marker. The clusters do not have labels associated with them and need to be analysed by clinicians. Such labels could help clinicians to target areas of interest and point them to the section of the raw trace that needs attention.

3.2.4 Background Continuity State Detection

Turnbull *et al.* published a study on automatic detection of *tracé alternant* pattern (see Figure 2.1) during infant sleep [36]. EEG recordings from full term infants were used in this study, and full channel EEGs were used. The EEG signals were analysed using wavelet decomposition, and the powers of different frequency bands were analysed. The power of the theta band, meaning signals with frequencies in the range of 4 - 8 Hz, was found to be a good indicator of the presence of *tracé alternant*. A Fourier transform of the theta band power was carried out and it was reported that the plot peaked at the region between 5 to 10 cycles per minute. The peak power of this region is used as an indicator of *tracé alternant*. On visual comparison, the peaks

of a plot of this feature correspond to regions of EEG marked by clinicians as *tracé alternant*. No automated algorithms were presented, but rather the Fourier transform of the theta band power was presented as a feasible feature for an automated system. Different infants showed different changes in this feature, where in some infants, the rising edge of the peaks corresponded to the *tracé alternant* region, while in other infants the falling edge corresponded to the *tracé alternant* region. The paper also pointed out that other applications, such as noise rejection algorithms, could help improve the feature further.

Navakatikyan *et al.* published a study where a new measure, called rEEG (for range-EEG), was proposed as a way to improve on the aEEG based continuity detection or classification method [37, 38]. It was suggested that the peak-to-peak value of the EEG in a 2 second window be used to construct this rEEG. Different amplitude bands were defined, and the proportion of the signal in each of these bands was calculated. These proportions were used to determine which background continuity state the signal belonged to. The algorithm performed well, with a 76% average agreement rate. The study also noted that for the human experts who rated the EEG for their background states, the average agreement rate amongst themselves was 68%. This highlights the issue of subjectivity in continuity detection.

3.3 Discussion

3.3.1 The Gap between the Fields

The literature review shows a division between the approaches in research carried out in the medical and engineering fields. Medical research, in general, focuses on correlating features with clinical outcomes or maturation. Engineering research, on the other hand, has tended to focus on the recognition of known patterns of interest, and few works address the features used by the medical researchers for their work in defining maturation and correlating signals with clinical outcomes.

3.3.2 Continuity as a Measure of Maturation

The medical research shows that two significant criteria when analysing preterm EEG, whether for maturation or prognosis, are continuity and specific landmark patterns.

Not all landmark patterns occur at the channels recorded by the BrainZ monitoring system used for data collection in this study. Specific markers also need to be read in context, a task that requires human experts years of training to carry out accurately. Continuity makes a more realistic target at this stage as an EEG feature that can be further analysed and quantified. Currently, the quantification of continuity includes primitive measures such as maximum IBI or intensity measurements. While these measurements show a correlation with maturation, they only offer a global view of the signal and cannot show how continuity changes within a recording. Showing the change of continuity within a recording can also help identify the sleep-wake cycle, another marker that clinicians look for when analysing preterm EEG [15].

3.3.3 The Importance of Context

The purpose of this research is to focus on extracting useful features for neonatal EEG that can help identify maturation or brain injury. Unlike applications such as hypoxia onset detection or seizure detection, there is no convenient way to identify some sort of marker or label for a section of EEG to be targeted for classification process. This causes problems in EEG analysis, as the behaviour of the signal changes dramatically during the recording, depending on the state of each infant. Without some form of label, features extracted will either be an averaged measurement taken from the whole duration of the recording (as in the intensity measurements), or a minimium or maximum measurement (as in maximum inter-burst interval). With recordings which only consist of one background state, this will not pose a problem. However, as infants grow towards term, the behaviour of the EEG will change, depending on the state each infant. If only one global feature is extracted, this may average out the effect of the different states; or even mask the information in one of the states, if a maximum or minimum is used. A continuity context can also help to label the area in an EEG signal that is being focused on (e.g. burst suppression), and provide a way to compare information from different EEG recordings to ensure that only signals of similar background states are being compared.

3.3.4 Using Engineering Methods to Help Medical Research

The review of medical and engineering research literature uncovered a gap in the area of continuity quantification. Although clinicians rely on continuity to determine the

maturation of an infant, and related features have been demonstrated to correlate with maturation and clinical outcomes, its quantification has not been thoroughly explored in engineering research. Background continuity pattern recognition attempts have been made using classifiers to assign labels to EEG signals. However, to help researchers to further explore the relationship between continuity and maturation or clinical outcomes, as well as aiding further development of automatic EEG analysis, a quantified measure is more valuable than qualitative labels. A quantitative index would better express the degree of continuity than a qualitative label, especially since the labelling is currently carried out with a degree of subjectivity.

3.4 Summary

In this chapter, a literature review of medical and engineering work on neonatal EEG analysis was presented. In medical research, works have focused on recognising markers that can be correlated with maturation or clinical outcomes for the infant. Various EEG features were used, with continuity related features like maximum interburst period being a good indicator of maturation and clinical outcome. In engineering, foci are on pattern recognition with labelled data for particular patterns such as seizure and hypoxia, where the tasks are usually carried out by training supervised learning systems to differentiate the pattern of interest from general background activities. Some work has been carried out on an automatic background state detection system, but so far no studies of a complete quantified measurement of continuity have been found. Since the goal of this study is to find features that can assist clinical researchers to study preterm infant maturation, and clinicians have successfully correlated continuity related features with maturation and clinical outcomes, a study in quantifying continuity can be seen to be beneficial.

Chapter 4

Time Frequency Analysis of EEG Signals

Time-frequency analysis is a way to represent the energy contents of a signal in the joint time-frequency domain. It provides a good visual way to separate the frequency contents of a multi-component signal, and display the changes of these components with respect to time. This chapter outlines the initial investigative work on neonatal EEG signals using time-frequency analysis. The Cohen's class distributions are discussed, and kernel optimisation for the Cohen's class distributions is outlined. Segments of EEG with different background continuity states are analysed using a Cohen's class distribution, and their characteristics are discussed. The findings contributed to the direction of the project and are discussed at the end of the chapter.

4.1 Time Frequency Analysis

Time-frequency analysis, in contrast with the traditional time or frequency domain analysis, analyses the signal in both the time and the frequency domains. It involves representing the signal in the time-frequency domain to show the changes in

its frequency components with respect to time. Using time-frequency analysis, the energy contents of an EEG signal can be visualised. This can aid in understanding the properties of the EEG in order to determine the best approach to further analyse and process the signal. For this study, the Cohen's class distribution is used, which expresses the distribution of energy in the joint time-frequency domain.

4.1.1 Overview of Cohen's Class Distributions

Cohen's class distributions are a class of energy distributions which show the distribution of the energy within a signal in the time-frequency domain. They are based on the Wigner-Ville distribution, with a kernel to help eliminate the cross-term artifact that is inherent to the bilinear nature of the Wigner-Ville distribution. The Cohen's class distributions differ from wavelet decomposition as they do not give a set of coefficients that can be used as features in later stages, although they do provide an energy distribution in the joint time-frequency domain with consistent resolution, whereas the wavelet decomposition gives coefficients of different time resolutions as the frequency (or scale) changes. Spectrogram and Wigner-Ville distributions can both be described as special cases of Cohen's class distributions.

Spectrogram

One of the simplest and most intuitive ways to represent the signal in the time frequency domain is the Short-Time Fourier Transform (STFT). This involves multiplying the signal by a window function, and applying the Fourier transform to the windowed function. This process is repeated to give the Fourier transform with windowed signals that emphasise the original signal at different times. The result of this is the short-time Fourier transform. If the magnitude of the STFT is squared, the result is the energy spectrum for the signal in the time-frequency domain, and is called a spectrogram.

While the spectrogram gives a picture of the change in frequency contents with respect to time, it has a trade-off between time and frequency resolution due to the uncertainty principle. In other words, increasing the resolution in the time direction will cause the resolution to decrease in the frequency direction, and vice versa [39].

Wigner-Ville Distribution

To overcome this trade-off, various distributions have been designed to represent the signal in the time-frequency domain. One of the first distributions proposed for this purpose is the Wigner-Ville distribution [39, 40, 41]. Because of its simplicity, the number of desirable properties, and its relationships to the more complicated distributions, several analyses and studies have been done on the Wigner-Ville distribution. The Wigner-Ville distribution is defined in equation 4.1.

$$W_z(t,f) = \int_{-\infty}^{\infty} z(u+\frac{\tau}{2})z^*(u-\frac{\tau}{2}))e^{-j2\pi f\tau}d\tau \tag{4.1}$$

where $z(t)$ is a complex signal and $z^*(t)$ is its complex conjugate.

To see how this equation can be derived [40], one can consider the time-frequency distribution, $W_z(t,f)$, expressed as the Fourier transform of a signal kernel, $K_z(t,\tau)$, in the time-lag domain, (t,τ), where τ represents the variable of the lag domain. The lag domain is defined here as the inverse Fourier transform of the frequency domain, f. The Fourier transform is defined as:

$$F(f) = \mathcal{F}_{t\to f}\left\{f(t)\right\} = \int_{-\infty}^{\infty} f(t)e^{-j2\pi ft}dt \tag{4.2}$$

where t is the time domain variable and f is the frequency domain variable. The inverse of the Fourier transform is defined as:

$$f(t) = \mathcal{F}^{-1}_{t\leftarrow f}\left\{F(f)\right\} = \int_{-\infty}^{\infty} F(f)e^{j2\pi ft}df \tag{4.3}$$

The domain transform is explicitly stated in the notation to avoid confusion when dealing with dual domain transformations, such as those involved in the joint time-frequency domain.

To reach the time-lag signal kernel, $K_z(t,\tau)$, consider a signal $z(t) = e^{j\phi(t)}$. The instantaneous frequency (IF) of $z(t)$, $f_i(t)$, is $\frac{\phi'(t)}{2\pi}$, where $\phi'(t)$ is the first derivative of $\phi(t)$. The time-frequency distribution at any given time t should be the instantaneous frequency, f_i. This can be expressed as:

$$W_z(t,f) = \delta(f - f_i(t)) \tag{4.4}$$

The signal kernel, $K_z(t,\tau)$, is therefore the inverse Fourier transform in the f

domain:

$$\begin{aligned} K_z(t,\tau) &= \mathcal{F}^{-1}_{\tau \leftarrow f}\{W_z(t,f)\} \\ &= \mathcal{F}^{-1}_{\tau \leftarrow f}\{\delta(f - f_i(t))\} \\ &= e^{j2\pi f_i(t)\tau} \\ &= e^{j\phi'(t)\tau} \end{aligned} \tag{4.5}$$

The derivative of $\phi(t)$, $\phi'(t)$, can be defined using first principles as follows:

$$\phi'(t) = \tau \overset{\lim}{\rightarrow} 0 \ \frac{\phi(t+\tau/2) - \phi(t-\tau/2)}{\tau} \tag{4.6}$$

Using the central finite-difference approximation, the derivative can be approximated as:

$$\phi'(t) \approx \frac{\phi(t+\tau/2) - \phi(t-\tau/2)}{\tau} \tag{4.7}$$

Substitute equation (4.7) into equation (4.5), $K_z(t,\tau)$ can be modelled as:

$$\begin{aligned} K_z(t,\tau) &= e^{j\phi'(t)\tau} \\ &= e^{j[\phi(t+\tau/2) - \phi(t-\tau/2)]} \\ &= e^{j\phi(t+\tau/2)} e^{-j\phi(t-\tau/2)} \end{aligned} \tag{4.8}$$

Recall $z(t) = e^{j\phi(t)}$,

$$K_z(t,\tau) = z(t+\tau/2)z^*(t-\tau/2) \tag{4.9}$$

$$W_z(t,f) = \int z(t+\tau/2)z^*(t-\tau/2)e^{j2\pi f\tau}d\tau \tag{4.10}$$

The Wigner-Ville distribution has a number of desirable properties which makes it ideal as a starting point to explore time-frequency distributions (TFDs). However, one of its biggest drawbacks is the existence of cross-terms (or artifacts). These

cross-terms arise in two situations: in non-linear FM mono-component signals, and in multi-component signals.

Cross-terms arising from non-linear FM mono-component signals are referred to as inner artifacts. The interference occurs because the central finite difference estimation of $\phi'(t)$ is no longer exact. The inner-artifacts oscillate in the direction normal to the IF law, $f_i(t)$. The interference can be reduced by using a windowed version of the Wigner-Ville distribution.

Cross-terms present in multi-component signals are referred to as outer artifacts. These occur due to the quadratic nature of the distribution. The outer-artifacts are also oscillatory in nature, and are found mid-way between the IF of the different components.

Cohen's Class Distribution

To compensate for the existence of cross-term artifacts, smoothing windows can be applied to the TFDs by taking advantage of the oscillatory nature of the cross-terms. The Cohen's class distribution is a class of TFDs which incorporate a kernel used for smoothing [39, 40, 41]. The Cohen's class TFD is defined as follows:

$$\rho_z(t,f) = \int\int G(t-u,\tau) z(u+\frac{\tau}{2}) z^*(u-\frac{\tau}{2}) e^{-j2\pi f\tau} du d\tau \tag{4.11}$$

where $G(t,\tau)$ is an arbitrary kernel that controls the time-frequency resolution and cross-term artifact attenuation. This kernel in the distribution is expressed in the time-lag domain, which is preferred for implementation purposes. To understand how the kernel can help with cross-term attenuation, $W_z(t,f)$ can be transformed into its ambiguity function, $A_z(\nu,\tau)$, in the Doppler-Lag domain:

$$A_z(\nu,\tau) = \mathcal{F}_{t\to\nu}\left\{K_z(t,\tau)\right\} = \int z(t+\frac{\tau}{2}) z^*(t-\frac{\tau}{2}) e^{j2\pi\nu t} dt \tag{4.12}$$

where ν in the variable in the Doppler domain.

The Cohen's Class distributions can therefore be expressed as:

$$\begin{aligned}\rho_z(t,f) &= \mathcal{F}_{t\to\nu}\left\{\mathcal{F}^{-1}_{\tau\leftarrow f}\left\{g(\nu,\tau)A_z(\nu,\tau)\right\}\right\} \\ &= \int\int\int g(\nu,\tau)z(u+\frac{\tau}{2})z^*(u-\frac{\tau}{2}) \\ &\quad e^{j2\pi(\nu u-\nu t-f\tau)}dud\nu d\tau \end{aligned} \tag{4.13}$$

where $g(\nu,\tau)$ is the Fourier transform of the time-lag domain kernel, $G(t,\tau)$, i.e.:

$$g(\nu,\tau) = \mathcal{F}_{t\to\nu}\left\{G(t,\tau)\right\} \tag{4.14}$$

Since cross-term artifacts in time-frequency distributions are oscillatory, they will appear in $A_z(\nu,\tau)$ as components away from the origin of the Doppler-lag domain. Thus, for a kernel to effectively attenuate cross-term artifacts, $g(\nu,\tau)$ should have a large value in the regions close to the origin of the Doppler-lag domain, and a small value in regions further from it.

Another way to view the Cohen class distributions is to look at the resultant distribution in terms of the Winger-Ville distribution and the time-frequency kernel. The time-frequency kernel is the inverse Fourier transform of the time-lag kernel in the lag direction:

$$\gamma(t,f) = \mathcal{F}^{-1}_{\tau\leftarrow f}(G(t,\tau)) \tag{4.15}$$

where $\gamma(t,f)$ is the time-frequency representation of the kernel as expressed by the time-lag kernel ($G(t,\tau)$). Since the multiplication in the Fourier transformed domain equates to the convoluted integral in the original domain, equation (5.6) can be expressed as:

$$\rho_z(t,f) = \gamma(t,f) **_{(t,f)} W_z(t,f) \tag{4.16}$$

where $**_{(t,f)}$ represent the 2-dimensional convoluted integral. The time-frequency kernel can therefore be seen as a 2D filter being applied to the Wigner-Ville distribution.

Several kernels have been proposed and studied previously, and a selection of examples is shown in table 4.1.

Other desirable properties can also be ensured by the design of the kernel [40].

Distribution	Parameters	$G(t,\tau)$	$g(\nu,\tau)$
Wigner-Ville [40]	N/A	$\delta(t)$	1
Windowed Wigner-Ville [40]	$w(\tau)$	$\delta(t)w(\tau)$	$w(\tau)$
Choi-Williams [42]	σ	$\frac{\sqrt{\pi\sigma}}{\lvert\tau\rvert}e^{-\pi^2\sigma t^2/\tau^2}$	$e^{-\nu^2\tau^2/\sigma}$
B-distribution [43]	β	$\lvert\tau\rvert^{\beta}(\cosh(t))^{2\beta}t$	$\frac{\lvert\tau\rvert^{\beta}\lvert\Gamma(\beta+j\pi\nu)\rvert^2}{2^{1-2\beta}\Gamma(2\beta)}$
Modified B distribution [44]	β	$\frac{(\cosh(t))^{-2\beta}}{\int_{-\infty}^{\infty}(\cosh(\xi))^{-2\beta}d\xi}$	$\frac{\lvert\Gamma(\beta+j\pi\nu)\rvert^2}{\Gamma^2(\beta)}$
ZAM [45]	$w(\tau)$,a	$w(\tau)\text{rect}\frac{t}{2\tau/a}$	$w(\tau)\frac{a}{2\lvert\tau\rvert}\text{sinc}frac2\nu\tau a$
Born-Jordan [40]	α	$\frac{1}{\lvert 2\alpha\tau\rvert}\text{rect}\frac{t}{2\alpha\tau}$	$\text{sinc}(2\alpha\nu\tau)$
Spectrogram [40]	$w(t)$	$w(t+\frac{\tau}{2})w(t-\frac{\tau}{2})$	$A_w(\nu,\tau)$

Table 4.1: Examples of common Cohen's class distribution kernels in time-lag and Doppler-lag domain. Note that $w(t)$ indicates a choice of window which is part of the distribution parameters.

4.2 Kernel Decision

The two main candidates for the kernel used to calculate the time-frequency distribution of the EEG signals were the Choi-Williams distribution and the Modified B-distribution. Both of these kernels were recommended and used in [46] and [20], and have been shown to be suitable for EEG signals. The kernels were implemented in Matlab (see Appendix B.1), based on code segments available from [40].

4.2.1 Choi-Williams Distribution

The Choi-Williams distribution (CWD) is one of the earliest reduced interference distributions (RID) proposed [42]. It is also known as the exponential distribution (ED) because of the exponential term in its formula. The kernel is defined, in the Doppler-lag domain as follows:

$$g_{\nu,\tau} = e^{-\nu^2\tau^2/\sigma} \tag{4.17}$$

where σ is a parameter that controls the cross-term attenuation. The shape of the kernel, with different values of σ, in the ambiguity domain can be seen in Figure 4.1.

As seen in Figure 4.1, the higher σ is, the closer the kernel is to unity. This means that for large values of σ, the Choi-Williams distribution resembles the Wigner-Ville distribution. The effect of σ can be seen in Figure 4.2

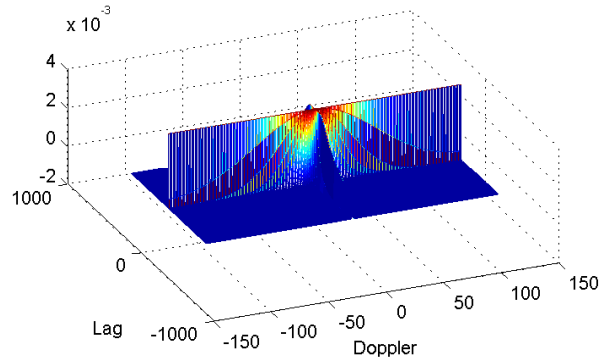

(a) $\sigma = 0.5$

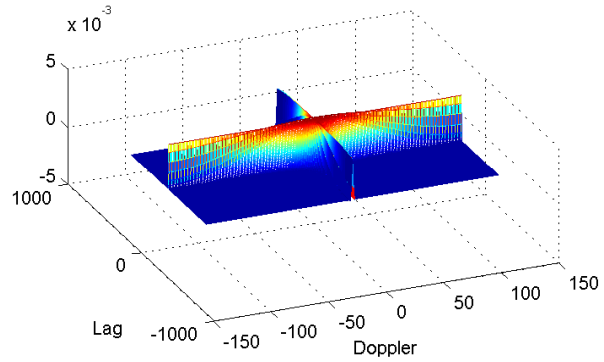

(b) $\sigma = 10$

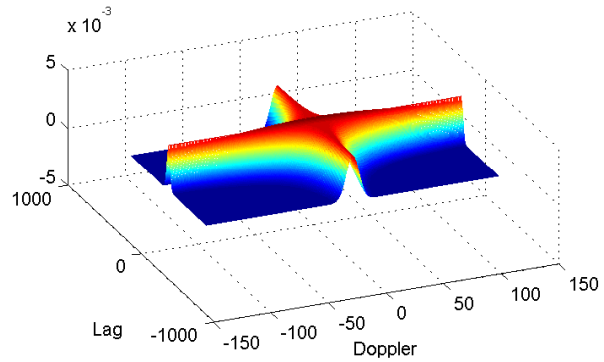

(c) $\sigma = 1000$

Figure 4.1: The Choi-Williams distribution kernel with different σ values in the Doppler-Lag domain

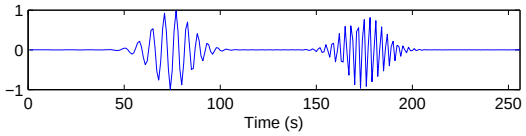

(a) Test signal

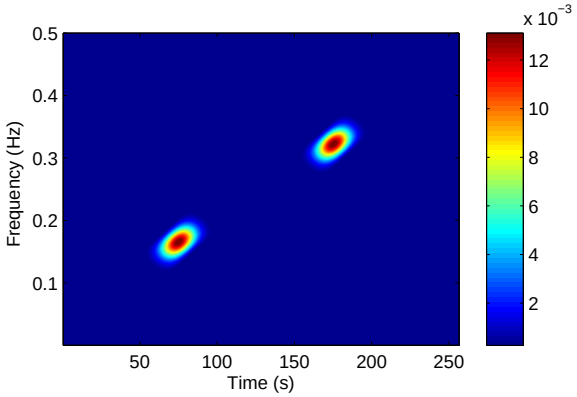

(b) $\sigma = 0.5$

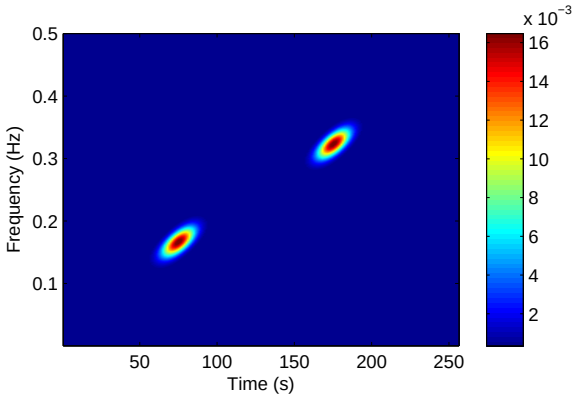

(c) $\sigma = 10$

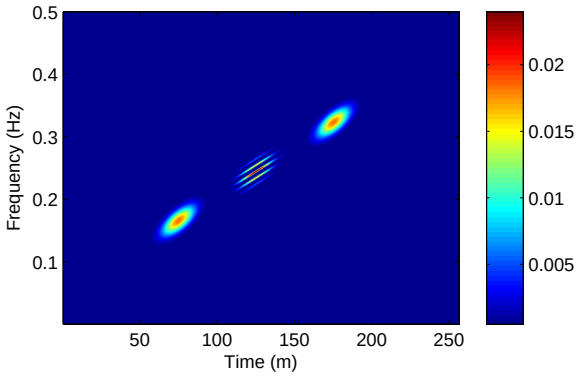

(d) $\sigma = 1000$

Figure 4.2: The Choi-Wailliams distributions, with different σ values, of a multi-component test signal

The design of the kernel is implemented such that the kernel is close to 1 around the origin of the ambiguity domain, and exponentially decreases toward zero in regions further away from the origin. As described in section 4.1.1, cross-term components lie further away from the origin than the auto-term. Multiplying the Choi-Williams distribution kernel with the ambiguity function before calculating the TFD ensures that the cross-term components have been attenuated. Notice that the kernel has a value of 1 along the lag and Doppler axes (i.e. when τ or ν equal to 0). This is designed to ensure time and frequency marginals [39].

4.2.2 Modified B-Distribution

The Modified B-distribution (MBD) [44] is a modified version of the B-distribution [43], proposed as a way to track the instantaneous frequency of a multi-component signal. The kernel is defined as follows:

$$G(t,\tau) = \frac{\cosh^{-2\beta} t}{\int_{-\infty}^{\infty} \cosh^{-2\beta} \xi d\xi} \tag{4.18}$$

where β is the parameter that controls the cross-term attenuation. The shape of the kernel, with different values of β, in the ambiguity domain can be seen in Figure 4.3.

Compared with the CWD, the MBD lacks the symmetry between the lag and Doppler dimensions. It is also independent from τ, the variable in the lag direction. In terms of its performance, the kernel favours frequency resolution, and the kernel shape in the time-frequency domain shows that the distribution smoothes the Wigner-Ville distribution in the time-direction. This is caused by the asymmetrical kernel shape. The Doppler-Lag kernel of the Modified B-distribution is uniform along the Lag direction (the inverse Fourier transform of the frequency domain) and the width of the peak in the Doppler direction (Fourier transform of the time domain) is controlled by the parameter β. As β decreases, the curve in the Doppler direction resembles the delta function. The shape of the Modified-B kernel in time-frequency domain will therefore be delta in the frequency direction (Fourier of unity is the delta function) and as β decreases, the curve of the kernel in the time direction approaches unity. As discussed in section 4.1.1, the time-frequency distribution can be seen as the convoluted intergral of the Wigner-Ville distribution and time-fequency kernel. With this in mind, one can see that the Modified B-distribution does not filter in the frequency direction, and will act as a low-pass filter in the time direction. The amount

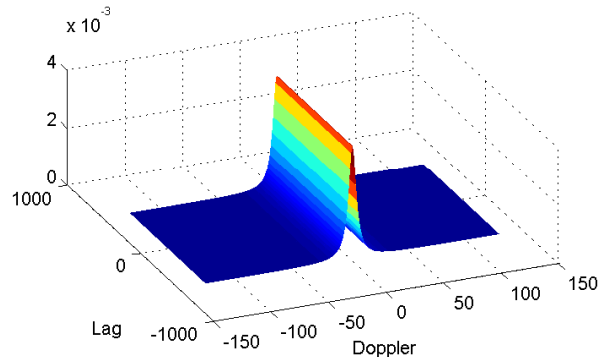

(a) $\beta = 0.05$

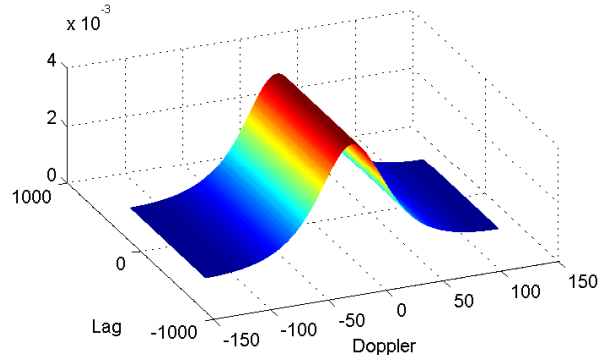

(b) $\beta = 0.5$

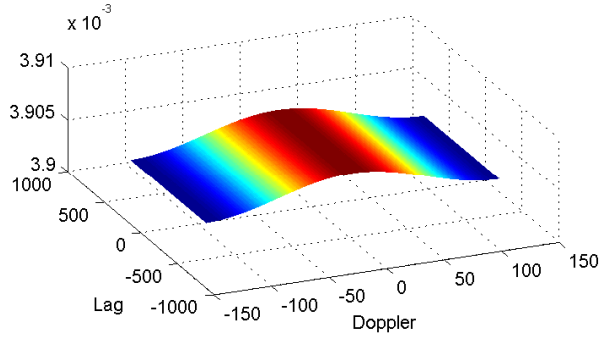

(c) $\beta = 10$

Figure 4.3: The Modified B-distribution kernel with different β values in the Doppler-Lag domain

of smoothing is controlled by the parameter β. The effect of β can be seen in Figure 4.4.

While this distribution may be inferior in time-resolution, it has the potential to show up frequency trends better because of its smoothing in the time direction.

4.2.3 Kernel Parameters and Smoothing Window Decisions

Various parameters were used with the two distributions to estimate the optimal value. These parameters were used to calculate the TFD of segments of neonatal EEG. In equation 5.6, the signal is expressed as a complex signal, $z(t)$. Since EEG signal in a real signal, the analytic signal is used. This is implemented by taking the Fourier transform of the EEG signal, and performing the inverse Fourier tranform on the positive portion of the signal spectrum.

The window length and kernel parameters discussed in this section only provide an estimate of the optimal value, since the time-frequency related feature has not yet been decided and therefore there is no concrete value of optimise the parameters against. The parameter values and the window length were chosen to ensure the TFD produced will show the time-frequency components of the EEG signal clearly, without excessive cross-term artifact. Further optimisation is possible once specific time-frequency features of the EEG are identified. The identification of time-frequency features is outside of the scope of this thesis.

Choi-Williams Distribution

Figure 4.5 shows the CWD of a segment of EEG with different values of σ. It was found that resolution is too low for $\sigma < 1$, and resolution improved little for $\sigma > 5$. For the rest of this section, the value of σ is set at 5, which filters out some of the cross-term artifacts without over-smoothing the distribution. There are still artifacts within the distribution, and these can be filtered by applying Hamming windows to the kernel. The Hamming window, $w(n)$, is defined as:

$$w(n) = 0.54 - 0.64\cos\left(2\pi\frac{n}{N}\right) \tag{4.19}$$

where N is the window length.

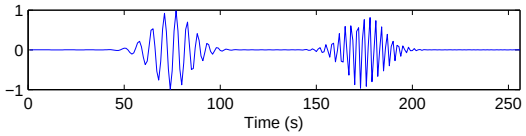

(a) Test signal

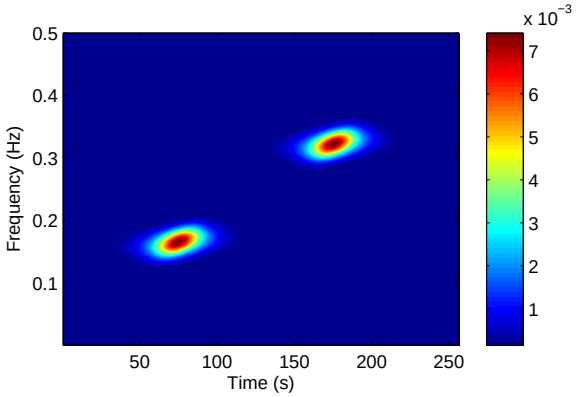

(b) $\beta = 0.05$

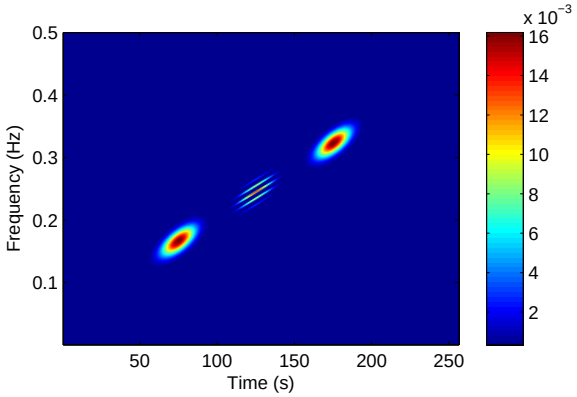

(c) $\beta = 0.5$

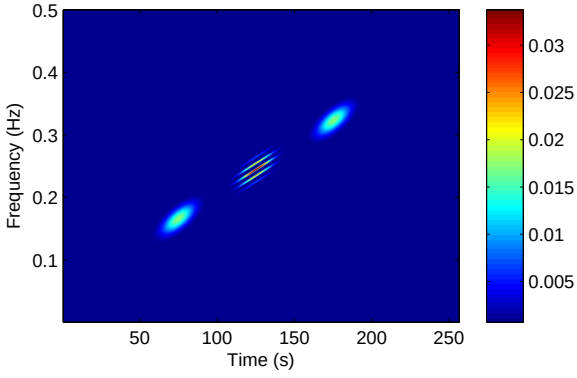

(d) $\beta = 10$

Figure 4.4: The Modified B-distributions, with different β values, of a multi-component signal

The effect of the Hamming window in the time domain is shown in Figure 4.6. The illustrations show that the artifacts perpendicular to the time axis are filtered out with Hamming windows applied in the time domain, and the 15 point Hamming window serves as a good compromise between resolution and cross-term attenuation. The same method of optimisation is also applied to Hamming windows in the lag direction to smooth the artifacts perpendicular to the frequency axis. The results are shown in Figure 4.7. It was decided that the estimated optimal length for the lag smoothing window is around 151.

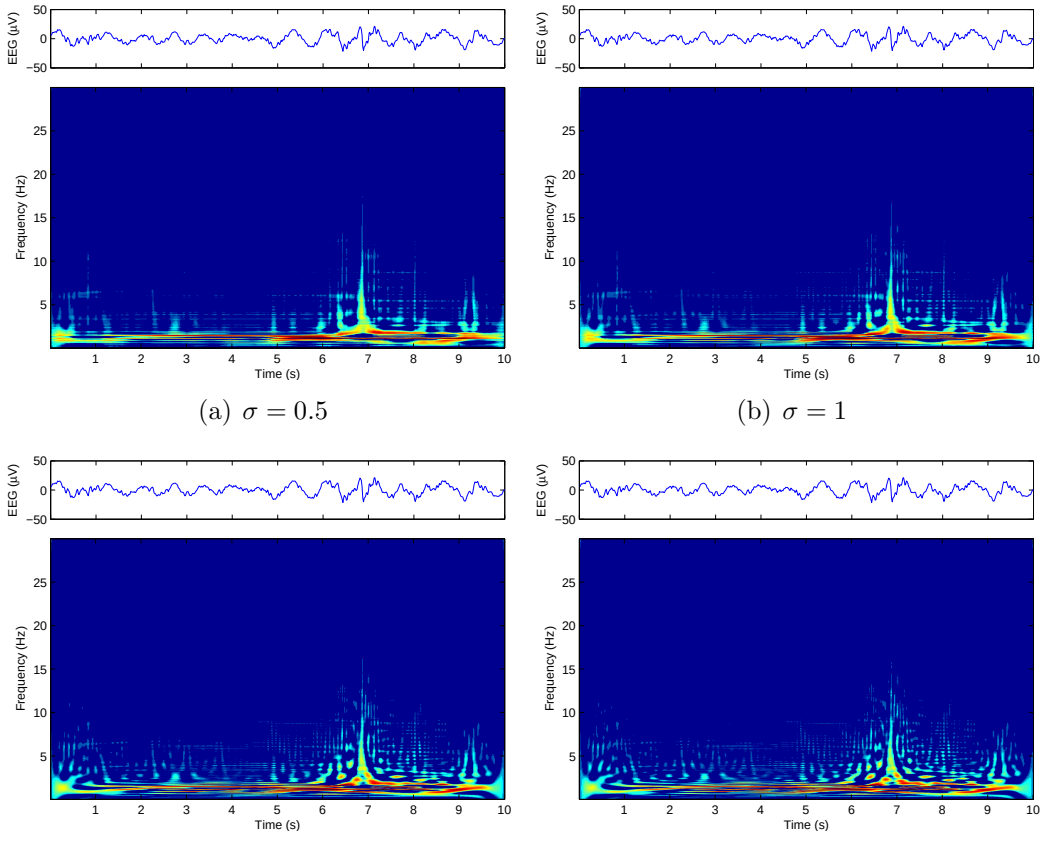

(a) $\sigma = 0.5$ (b) $\sigma = 1$

(c) $\sigma = 5$ (d) $\sigma = 10$

Figure 4.5: The Choi-Williams distributions of an EEG segment with different σ values

Modified B-Distribution

The parameter β in the MBD is optimised in the same way as σ in CWD described above. Figure 4.8 shows the MBD of the same EEG segment with varying values of β. It was found that a β value of 0.05 is a good compromise for resolution and cross-term

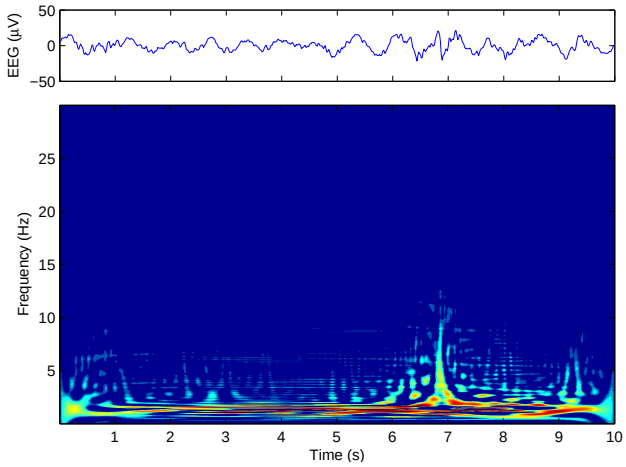

(a) 5 point Hamming window

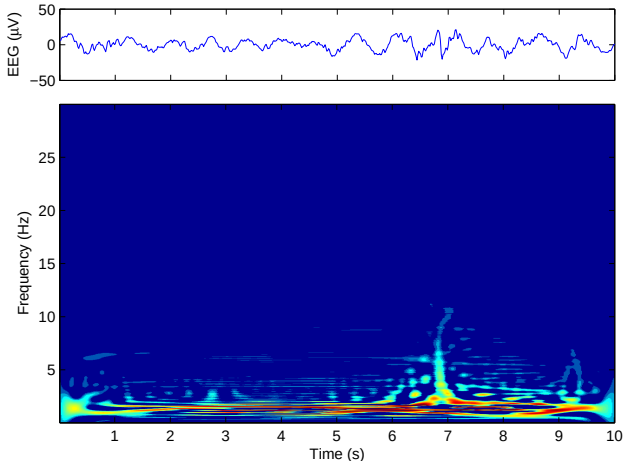

(b) 15 point Hamming window

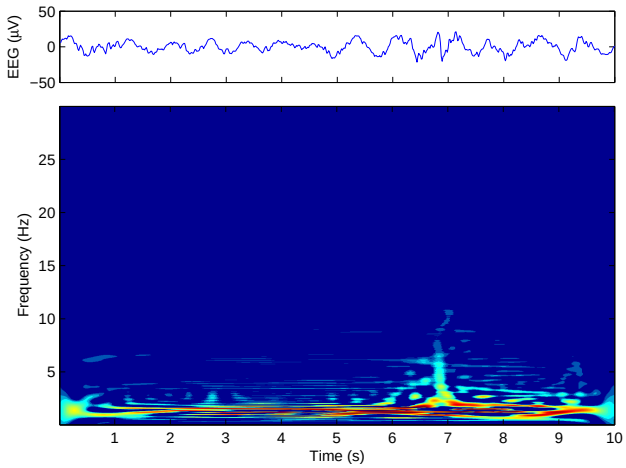

(c) 21 point Hamming window

Figure 4.6: The effect of applying Hamming windows of various lengths in the time domain to the Choi-Williams distribution of an EEG segment with $\sigma = 5$

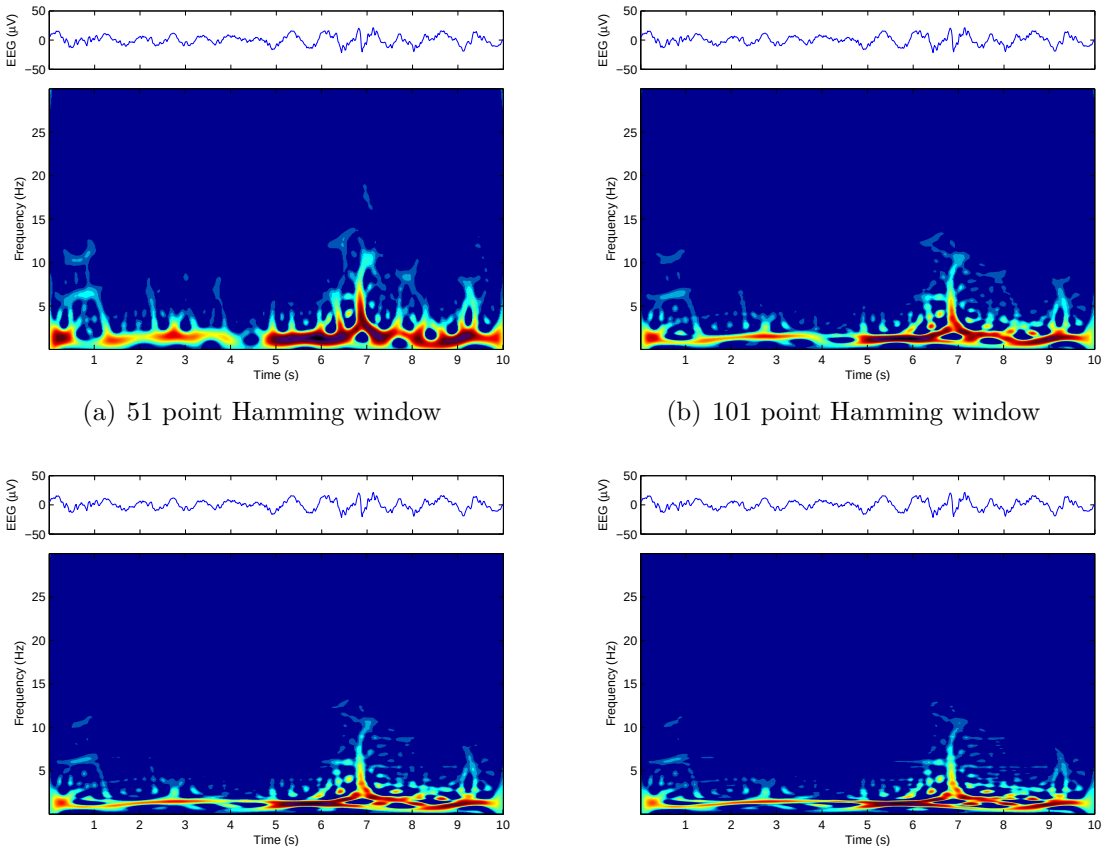

(a) 51 point Hamming window

(b) 101 point Hamming window

(c) 151 point Hamming window

(d) 201 point Hamming window

Figure 4.7: The effect of applying Hamming windows of various lengths in the τ domain to a Choi-Williams distribution with $\sigma = 5$ and 15 point Hamming windows in the time direction

attenuation. Note that the smoothing windows in the time and lag direction have not been optimised specifically for MBD. The same windows used for CWD are used for MBD for further smoothing.

4.2.4 Comparison Between the Two Kernels

The TFD of the EEG segments, as calculated using the smoothed version of the two kernels, are shown in Figure 4.9.

Although the CWD gives a better resolution, MBD performs better in the task of displaying the trend of the frequency components. This difference is highlighted in the case of seizure data, where the distribution of the signal is highly regular. The TFDs of a sample seizure signal segment using the two kernels are shown in Figure 4.10.

The high resolution CWD picks up the spikes in the seizure and displays them as vertical lines in the distribution. While this may be of interest, the spacing of the spikes can be estimated using the dominant frequency of the seizure signal, which appears as the bottom trend on the TFD. MBD smoothes the TFD in the time direction, which results in a better trend display.

Because we are interested in how the signal changes over time, the MBD is preferred over CWD at this stage. Because of its emphasis on the frequency resolution and smoothing ability in the time direction, it gives a better picture of the changes of the frequency component with respect to time. Further optimisation can be done after the features are finalised.

4.3 Time-Frequency Distributions of Different EEG Continuity Background States

Different signals were analysed by time-frequency analysis using the Modified B-distribution kernel described in the previous section. The signals were divided into different categories, and the distributions compared within each category and with distributions obtained from signals in other categories.

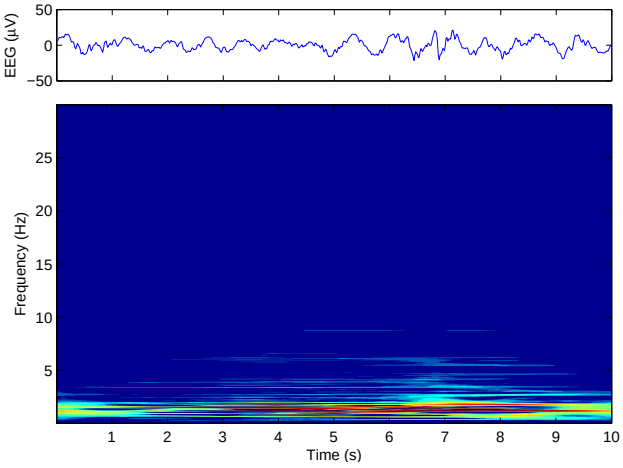

(a) $\beta = 0.01$

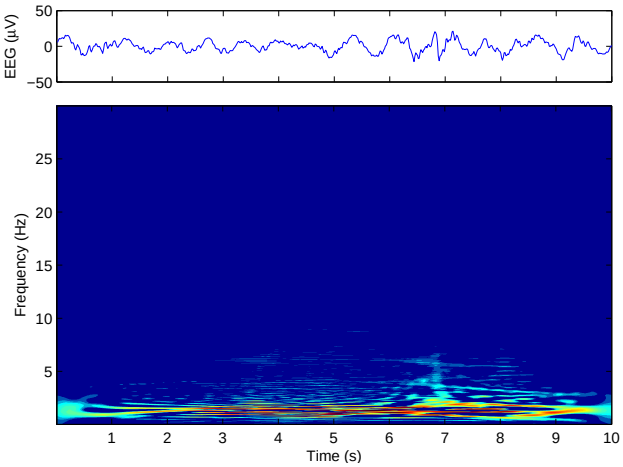

(b) $\beta = 0.05$

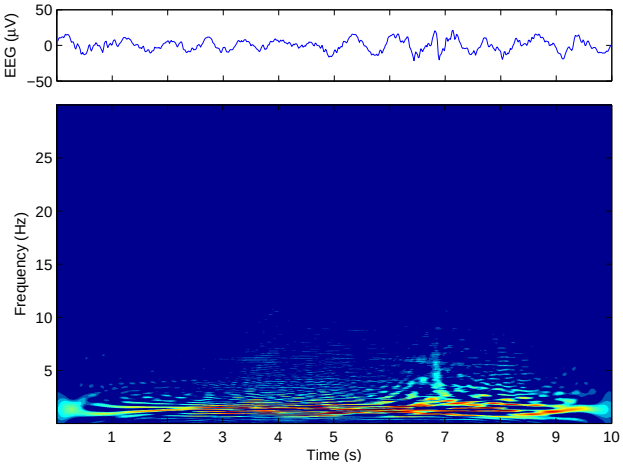

(c) $\beta = 0.1$

Figure 4.8: The Modified B-distribution of an EEG segment with different β values

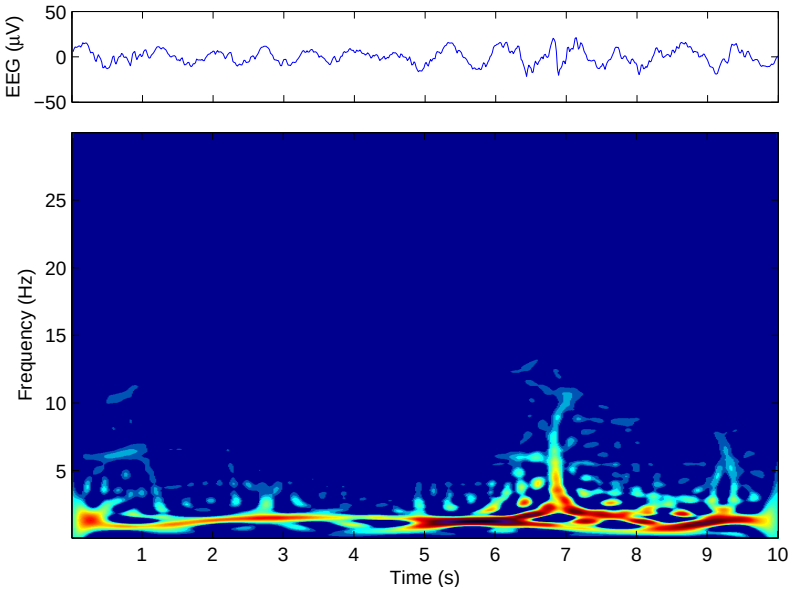

(a) Choi-Williams Distribution ($\sigma = 5$)

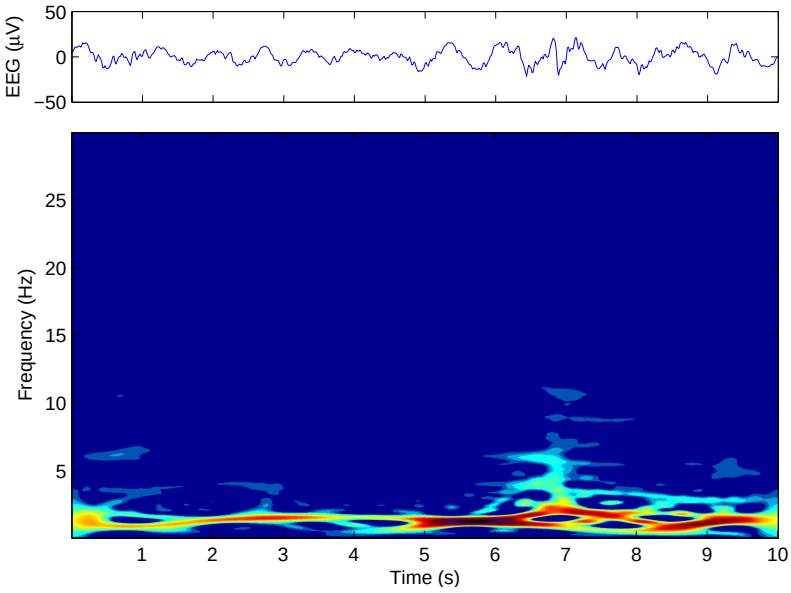

(b) Modified B-distribution ($\beta = 0.05$)

Figure 4.9: Comparison between Choi-Williams distribution and Modified B-distribution. Both distributions had a 15 point Hamming window applied in the time domain and a 151-point Hamming window applied in the lag domain.

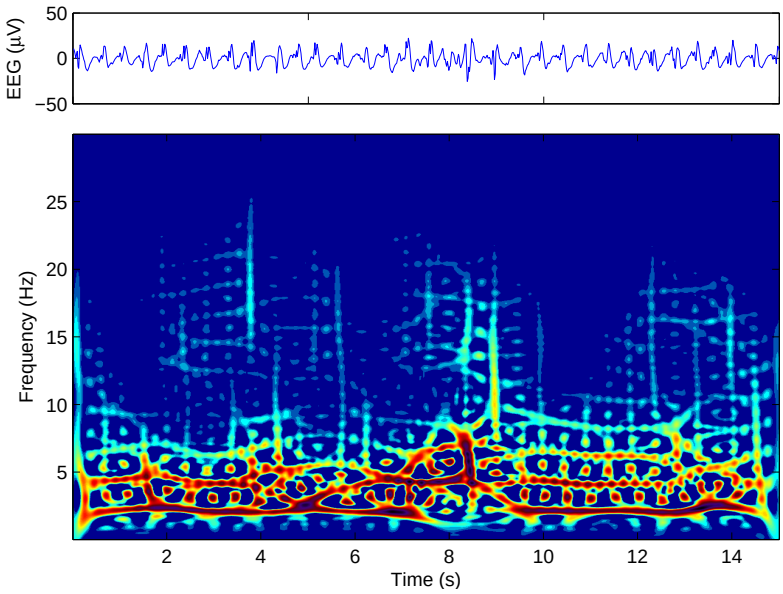

(a) Choi-Williams Distribution

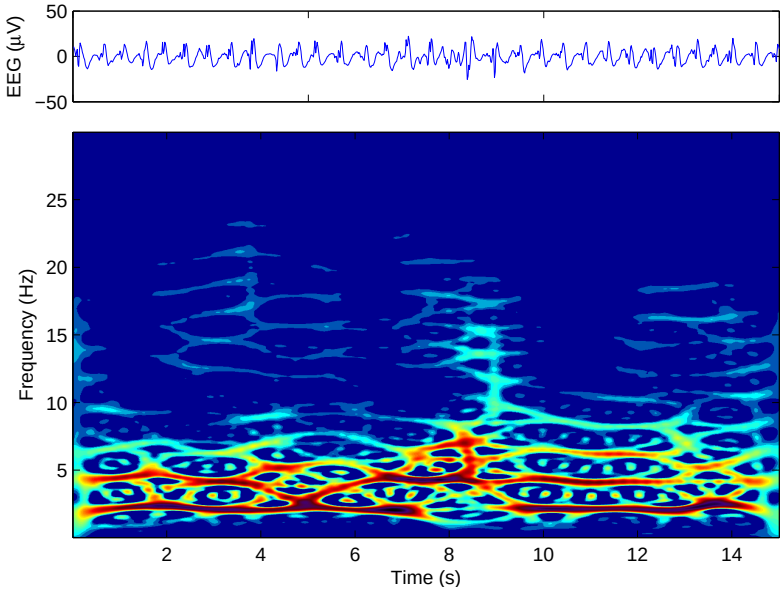

(b) Modified B-Distribution

Figure 4.10: The Choi-Williams distribution and Modified B-distribution of a seizure EEG sample

4.3.1 Normal Continuous Signals

A normal continuous signal is found in full-term or near full-term babies during wakeful periods. Characteristically, the normal continuous signal is a random looking signal within an amplitude of 25-50μV. The amplitude stays relatively constant, which is the definition of continuous signal. Figure 4.11 shows a typical time-frequency distribution (TFD) of a normal continuous signal.

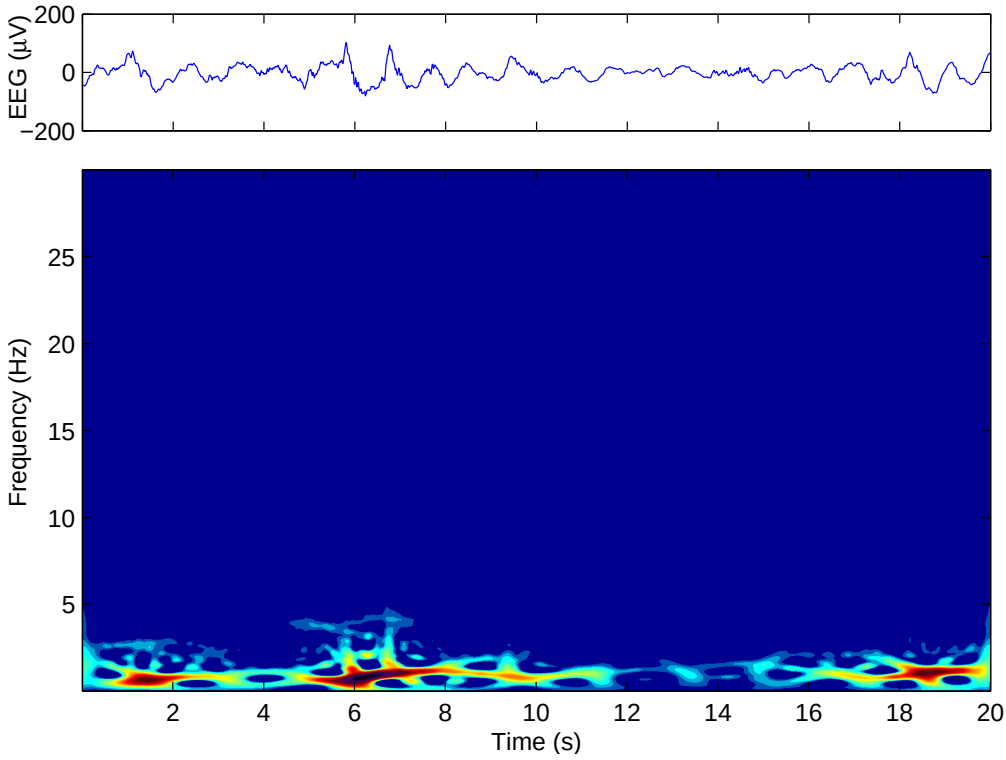

Figure 4.11: Time-frequency distribution of a typical normal continuous EEG signal

Note that the energy in this type of signal is concentrated below 5Hz, with subtle variations of intensity with respect to time. Since the skull acts as a low pass filter, any high frequency signal originating in the brain is attenuated. The major frequency components of the signal remain below 5Hz. There are no obvious patterns visible, which is predictable since the nature of a normal EEG recording is random.

4.3.2 Discontinuous Signals

Discontinuous signal is also known as *tracé alternant* in medical literature. The amplitude of discontinuous signals ranges from 25-150μV. The signals are characterised by

periods of high voltage activity, interrupted by relatively low voltage activity. These signals are common in premature infant recordings, and in recordings of full-term or near full-term infants during inactive sleep. The TFD of a typical discontinuous signal is shown in Figure 4.12.

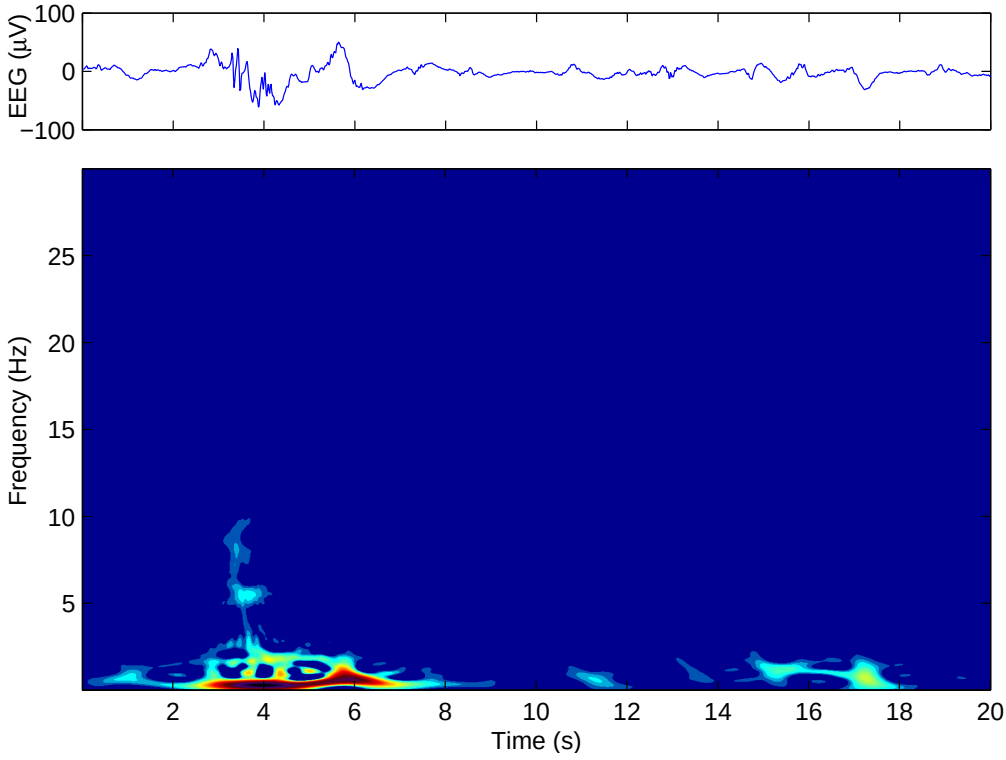

Figure 4.12: Time-frequency distribution of a typical discontinuous EEG signal

The distribution fails to show the energy contents of the areas with lower amplitudes, because most of the energy will be concentrated in the areas with a high amplitude. In the areas with high energy, the energy distribution is similar to that of the normal continuous signal.

4.3.3 Burst Suppression

Burst suppression refers to signal recordings that consist of bursts of high voltage activity, with prolonged periods of very low voltage, or inactivity, in-between bursts. It is known as *tracé discontinu* in very preterm infants and is considered normal when exhibited by these infants. This pattern should gradually evolve into discontinuous or *tracé alternant* patterns as the infant grows towards term. For term and older infants,

the presence of burst suppression is considered a sign of abnormality. A typical TFD of a burst suppression signal is shown in Figure 4.13.

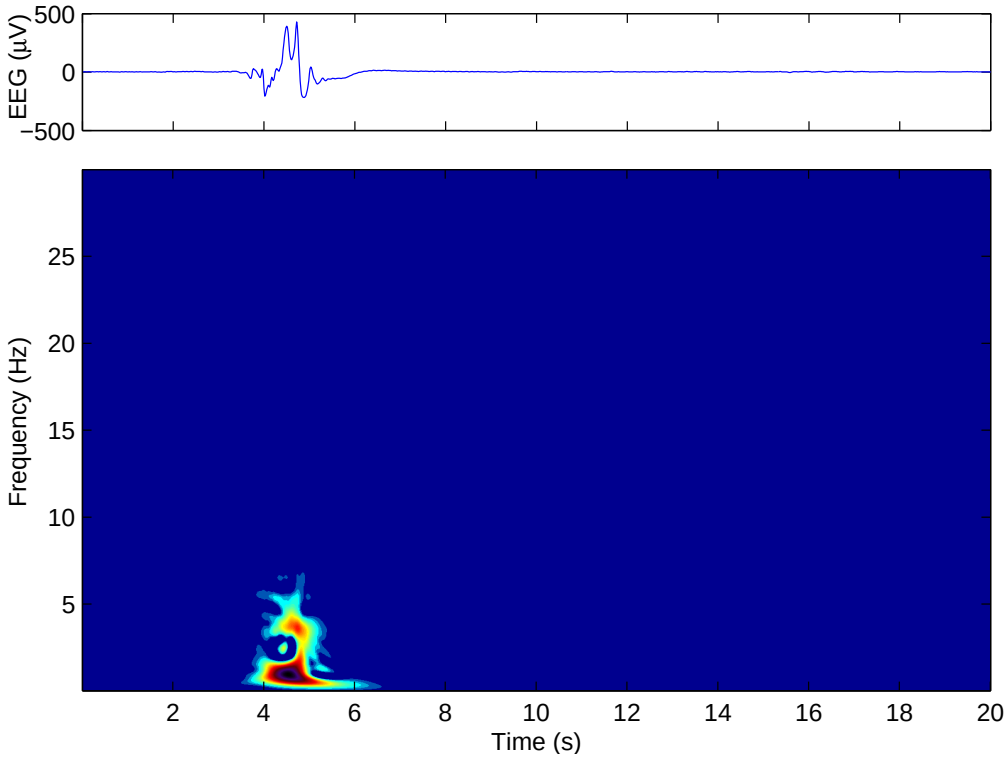

Figure 4.13: Time-frequency distribution of a typical burst suppression signal

Like the discontinuous data, energy is concentrated in the burst areas, where the amplitude is a lot higher than for the inter-burst periods. There is more energy in the higher frequencies compared with the discontinuous data, because of the spike-like nature of the burst signal.

4.3.4 Continuous Low Voltage

Continuous low voltage is an abnormal pattern where the voltage of the signal remains under 10μV. It is an indication of severe brain injury. As a rule, when the voltage of the EEG recording drops below 10μV consistently it is regarded as a sign of abnormality regardless of a lack of any other signs of abnormality. The TFD of a typical continuous low voltage signal is shown in Figure 4.14.

The characteristics of the continuous low voltage signal are similar to that of the continuous normal voltage, except with less intensity. Because the TFD is normalised,

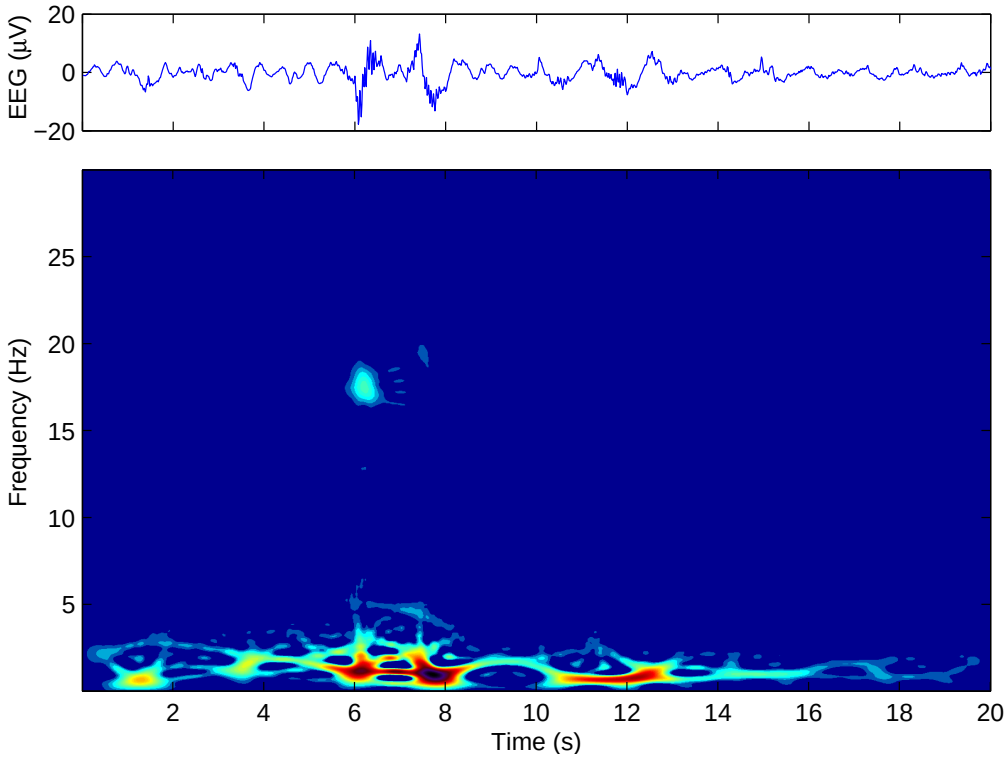

Figure 4.14: Time-frequency distribution of a typical continuous low voltage signal

it is difficult to distinguish between continuous normal voltage and continuous low voltage. However, since a constant low voltage is already a sign of abnormality on its own, a simple threshold approach can be used to detect continuous low voltage signals. Therefore, the similiarities between continuous normal voltage and continuous low voltage do not pose a serious problem in terms of the analysis of the signal. The high frequency muscle artifact is more prominent here compared with normal signals, since the signal to noise ratio is a lot lower in continuous low voltage signals.

4.3.5 Seizure

Seizure signals are any signals that exhibit a repetitive wave. The wave can take any shape or amplitude, and often evolves with changes in frequency and/or shape. Figure 4.15 shows the TFD of a seizure evolving over the duration of recording.

The TFDs of seizure signals, whether they evolve in the duration of recording or not, consist of lines which appear to be parallel in the frequency direction. This accounts for the periodic nature of seizure data. Notice that the parallel lines have

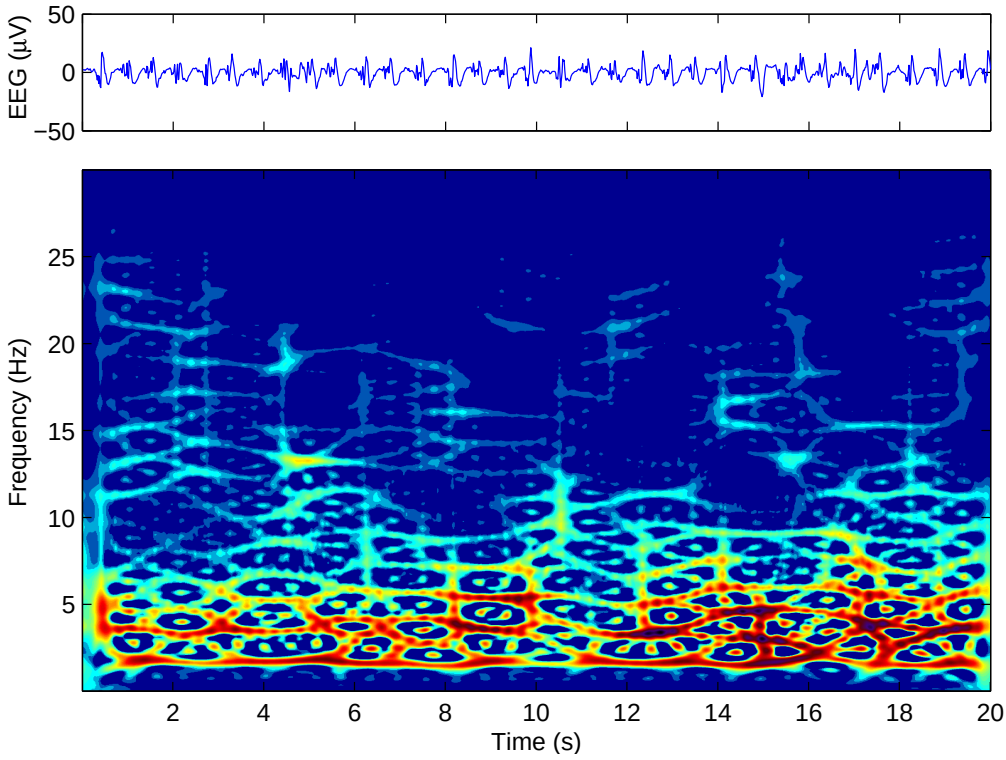

Figure 4.15: Time-frequency distribution of a seizure signal

even spacing, which is the approximate harmonic of the dominant frequency (the bottom trend). It is worth noting that during periods where seizures evolve, instead of a dramatic change in the TFD, the distribution seems to be adding more harmonics to the existing signal. Since seizure signals can contain spikes that occur in every period of the signal (see Figure 4.15 for examples of the spikes within the seizure signal), the time-frequency distribution will often include vertical lines where the spikes occur.

4.3.6 Summary of Observations

For premature infant EEG, most energy contents are concentrated below 5Hz, with occasional high frequency components. The high frequency components in the TFD are less intense than the low frequency components. Since the infants' skulls act as low pass filters, high amplitude components, if any, are attentuated. Since the energy of the signal is related to the amplitude, the TFD shows the changing intensity of a signal during different periods of the recording. This is most obvious in the case of discontinuous and burst suppression signals. Normal continuous data appear to be

random, with no obvious pattern visible. Most of the energy is concentrated below 5 Hz. Discontinuous data show a similar pattern to normal continuous data, with an emphasis on periods with high amplitudes. Energy presence in the lower amplitude region does not appear clearly in the TFD since the distribution of the energy is concentrated in the higher amplitude area, and the distribution of the energy in the lower amplitude region is around the same intensity as the cross term caused by the energy clustered around the high amplitude region. Burst suppression data shows more energy in higher frequencies compared to normal continuous data and discontinuous data, since bursts have spike-like components which contribute to high frequency contents. Seizure signals, because of their periodic nature, produce significantly different TFDs from other signal types. This type of signal is characterised by the parallel lines in the TFD in the frequency direction.

4.4 Discussion

The dominant patterns exhibited by preterm infants are burst suppression and normal discontinuous signals, with continuous signals gradually appearing more frequently as the infants grow towards term. The main problem with analysing discontinuous signals using TFD is that, since it is a distribution of energy, the area within the signal with relatively lower amplitude will show less variation compared with the high amplitude area. The logarithm of the TFD can be used to emphasise the energy distribution in the lower amplitude area, but it will also amplify any cross term artifact present in the high amplitude area. A solution to this problem would be to divide the signal into segments based on amplitude and calculate the segment TFDs individually. This can ensure that frequency components in both high and low amplitude periods are visible in the time-frequency distribution.

Looking at the actual TFD of the signal, it is interesting to note that besides seizures, normal continuous and discontinuous signals show no obvious underlying patterns. This means that the presence of a pattern in the TFD can help indicate certain markers. In terms of seizure detection, Boashash *et al.* has published a series of studies on the topic of using TFD for detection [47, 48, 49, 50]. Since seizure detection is outside the scope of this study, it will not be discussed further in this thesis.

4.5 Summary

In this chapter, time-frequency analysis and its applications in neonatal EEG analysis are discussed. The Cohen's class distributions offer time-frequency analysis techniques to represent the energy contents of a signal in the joint time-frequency domain. They are based on the Wigner-Ville distribution, a distribution designed to represent the instantaneous frequencies of a signal, with a kernel that controls the trade-off between time-frequency resolution and cross-term artifact. Serveral kernels have been proposed for various applications, with the Choi-Williams distribution and the Modified B-distribution being most suited to EEG analysis. The Choi-Williams distribution focuses on reducing the cross-term interference, while the Modified B-distribution is designed to optimise the instantaneous frequency tracking of signals with multiple frequency components. The Modified B-distribution is preferred over the Choi-Williams distribution for its smoothing in the time domain, which helps with the identification of frequency components. Different background states show different time-frequency compositions. Besides seizure patterns, normal EEG background states show no obvious pattern in the time-frequency domain. The time-frequency distribution cannot show the energy distribution of lower amplitude regions when there is a high amplitude region close by. Neonatal EEG is discontinuous in nature, therefore segmentation is needed to separate regions with different amplitude levels.

Chapter 5

EEG Segmentation

EEG signals are non-stationary in nature. Segmentation can divide a signal into psuedo-stationary segments, which can be further analysed using existing methods for examining stationary signals. Dividing signals into segments also mimics the way human experts mentally organise EEG signals into series of blocks of different characteristics or amplitudes, and isolating areas of interest such as burst segments during burst suppression. This chapter looks at three existing segmentation algorithms and presents an evaluation of them in the context of neonatal EEG analysis.

5.1 Segmentation Methods Overview

5.1.1 Spectral Error Measurement (SEM)

Spectral Error Measurement (SEM) is a way to segment a non-stationary signal, proposed by Bodenstein and Praetorius [51, 52]. The idea is to estimate the spectral difference between the signals in the test window and the reference window using autoregressive modelling.

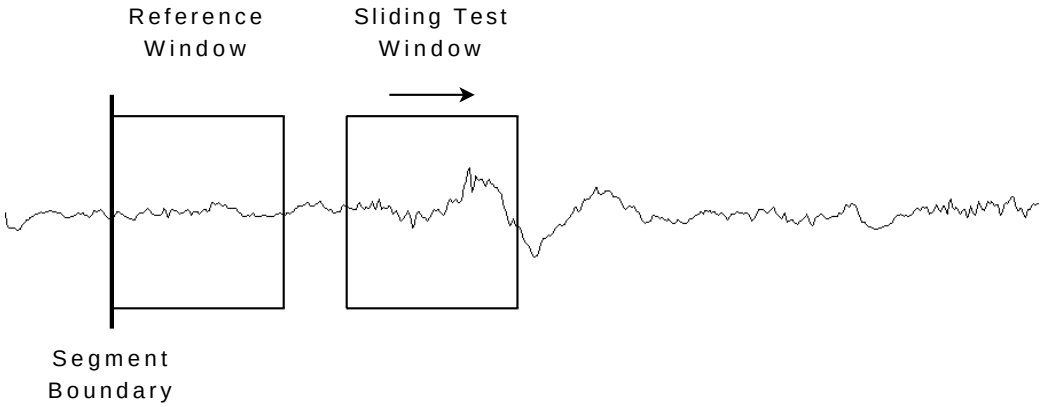

Figure 5.1: Definition of the fixed reference window and the sliding test window in the SEM algorithm

A fixed window is defined at the start of every segment. A set of linear prediction coefficients is obtained using the fixed window. An initial sliding test window is defined as a window of the same length as the reference window and starts immediately after the reference window. Figure 5.1 shows graphically how these windows relate to the segment boundaries [51] .

The predictive error of the testing window is analysed and its power spectral density is used as an estimation of spectral error between the reference window and the testing window. The segmentation criterion is defined in (5.1), where $r(n, m)$ is the autocorrelation function of the predictive error from the sliding test window, at time = n and lag = m.

$$\text{SEM}_n = \left(\frac{r(0,0)}{r(n,0)} - 1\right)^2 + 2\sum_{m=1}^{M}\left(\frac{r(n,m)}{r(n,0)}\right)^2 \tag{5.1}$$

where M is the maximum lag, which equals to the testing windows length.

The SEM_n value is compared with a predefined threshold. When SEM_n exceeds this threshold, a boundary is placed at time n. The algorithm will then restart, using the boundary as the start of the signal.

5.1.2 Nonlinear Energy Operator (NLEO)

A segmentation algorithm involving a non-linear energy operator (NLEO) was proposed by Agarwal and Gotman [34, 35]. NLEO is a way to represent energy in a frequency weighted operator, and is defined as follows.

$$\Psi(n) = x(n-1)x(n-2) - x(n)x(n-3) \tag{5.2}$$

where x is the signal in interest.

The NLEO value is a frequency weighted representation of the localised energy. For signals with the similar energy contents but different frequencies, the signal with a higher frequency will produce a higher NLEO value. For signals with similar frequency contents, the NLEO values will be proportional to the energy contents of the signal. This means that any discontinuity in the NLEO measurement indicates a change in the signal composition, either in energy or frequency contents, and therefore a possible segment boundary.

To detect a sudden change in the NLEO, a variable used for segmentation criterion is defined as follows:

$$G_{nleo}(n) = \sum_{m=n-N+1}^{n} \Psi(m) - \sum_{m=n+1}^{n+N} \Psi(m) \tag{5.3}$$

Where $2N$ is the window size. $G_{nleo}(n)$ reaches a peak when the $\Psi(n)$ is discontinuous. Figure 5.2 shows the effective windows used in the algorithm. The boundaries can therefore be found by detecting peaks from $G_{nleo}(n)$.

Because $\Psi(n)$ is highly localised, a windowed version was also examined, defined as follows:

$$\Psi_w(n) = \sum_{i=n-1-M}^{n-1+M} x(i) \sum_{j=n-2-M}^{n-2+M} x(j) - \sum_{k=n-M}^{n+M} x(k) \sum_{l=n-3-M}^{n-3+M} x(l) \tag{5.4}$$

where $2M$ is the window length used. The windowed NLEO (w-NLEO) algorithm uses the same criterion as the non-windowed version to detect the segment boundaries.

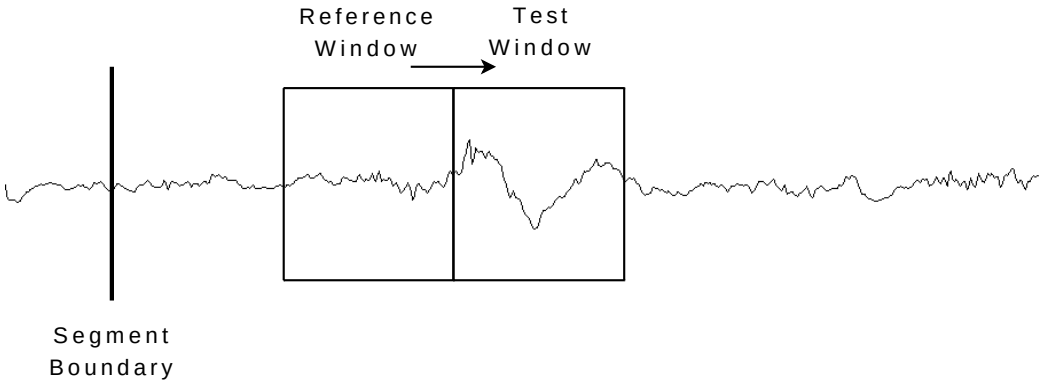

Figure 5.2: The effective reference test windows in the NLEO algorithm

5.1.3 Generalized Likelihood Ratio (GLR)

Generalized Likelihood Ratio is a method for time series proposed by Appel and Brandt [53, 52]. Like the SEM algorithm described above, it uses autoregressive models. The reference window is defined as the signal from the start of the current segment to the start of the sliding test window. As the sliding test window moves during each iteration, the reference window grows to cover all points in the signal that belong to the current segment. The sliding test window is of fixed length, and slides to the next point as the algorithm analyses the signal. The window formed by combining the reference and test windows is referred to as the pooled window. Figure 5.3 shows how the various windows relate to one another and how they relate to the signal [53]. The idea of GLR is to analyse the predictive error in the three windows mentioned above, to determine the amount of predictive error that will be generated should the sliding test window be regarded as part of the current segment.

The Generalised Likelihood Ratio, $d(n)$, is defined as (5.5), where n is the start of the test window, L is the length of the test window, and $e(n)$ is the predictive error at time n.

$$\begin{aligned} d(n) &= (n+L)\ln\left(\frac{\sum_{k=1}^{n+L} e(k)^2}{n+L}\right) - \\ &\quad [(n-1)\ln\left(\frac{\sum_{k=1}^{n-1} e(k)^2}{n-1}\right) + \\ &\quad L\ln\left(\frac{\sum_{k=n}^{n+L} e(k)^2}{L}\right)] \end{aligned} \tag{5.5}$$

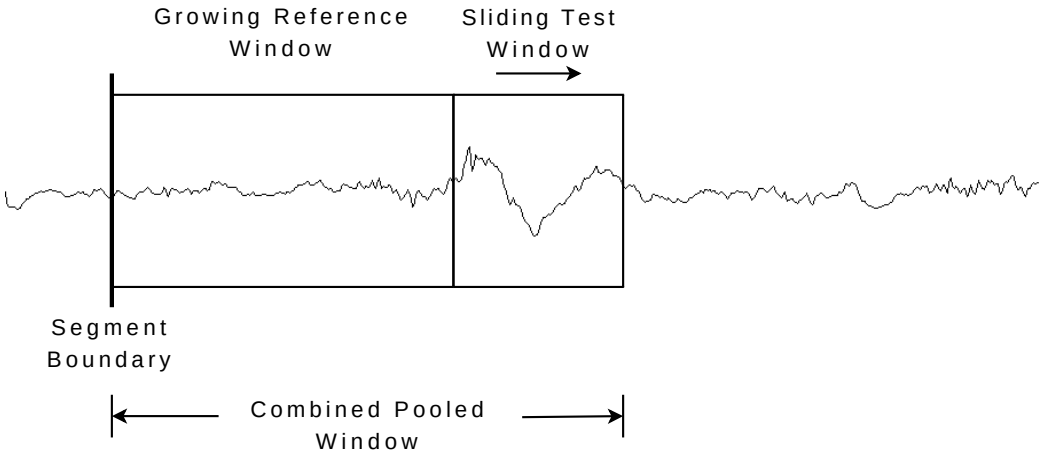

Figure 5.3: Definition of the growing reference window and the sliding test window in the GLR algorithm

The variable n continues to increment until $d(n)$ exceeds a certain predefined threshold. Once the threshold is exceeded the algorithm will go into the next phase to optimise the segment boundary position. A segment boundary is assumed to be present within the test window. The boundary position is optimised by applying a similar algorithm, as shown in equation 5.5. This time the combined pooled window length is fixed, and moves the boundary between reference and test windows instead. Figure 5.4 shows the windowing of the boundary optimisation phase [53].

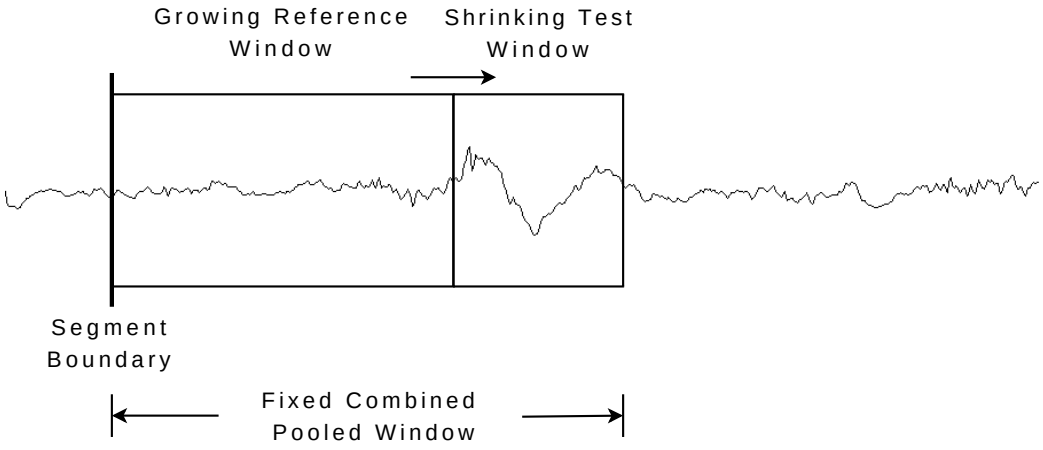

Figure 5.4: Definition of the time windows in the GLR algorithm during boundary optimisation phase

The boundary between the reference and test windows continues to increment in

an attempt to locate the boundary position that maximises $d(n)$. Once the boundary position is optimised, the algorithm will start again, using the segment boundary as the beginning of the signal.

5.2 Evaluation of the Methods

5.2.1 Method

A randomly selected pool of neonatal EEG recordings (12 2-channel recordings, 10 of which are 1-minute segments, the remainder 20 segments) were used in the evaluation phase [54], where each channel was processed separately with all four methods. To evaluate the appropriateness of the segment boundaries, the results from the segmentation algorithms were compared with the time-frequency distribution of the original signal. Recall from Chapter 4, the distribution is defined as follows:

$$\rho(t,f)_z = \int\int G(t-u,\tau)z(u+\frac{\tau}{2})z^*(u-\frac{\tau}{2})e^{-j2\pi f\tau}dud\tau \tag{5.6}$$

where $z(t)$ is the analytical signal derived from the EEG signal (as described in section 4.2.3), and $G(t,\tau)$ is an arbitrary kernel that controls the trade-off between the time-frequency resolution of the distribution, and the crossterm artifact arising from the bilinear nature of the distribution [41].

To compare the different segmentation methods, the original signal is displayed with the different segmentation boundaries superimposed on it. Along with this, the TFD is displayed for comparative purposes. The kernel used in the TFD for this task is the Modified B Distribution [44], as described in (5.7), with $\beta = 0.05$. The Modified B Distribution was chosen because it provides smoothing in the time direction to give smoother trends without too much loss in time-frequency resolution. The method described in this chapter is independent of the time-frequency distribution used, and therefore other distributions can also be used. The Modified B-distribution kernel is defined in the time-lag domain as follows:

$$G(t,\tau) = \frac{\cosh^{-2\beta} t}{\int_{-\infty}^{\infty} \cosh^{-2\beta} \xi d\xi} \tag{5.7}$$

Because the goal of the segmentation is to divide the non-stationary signal into

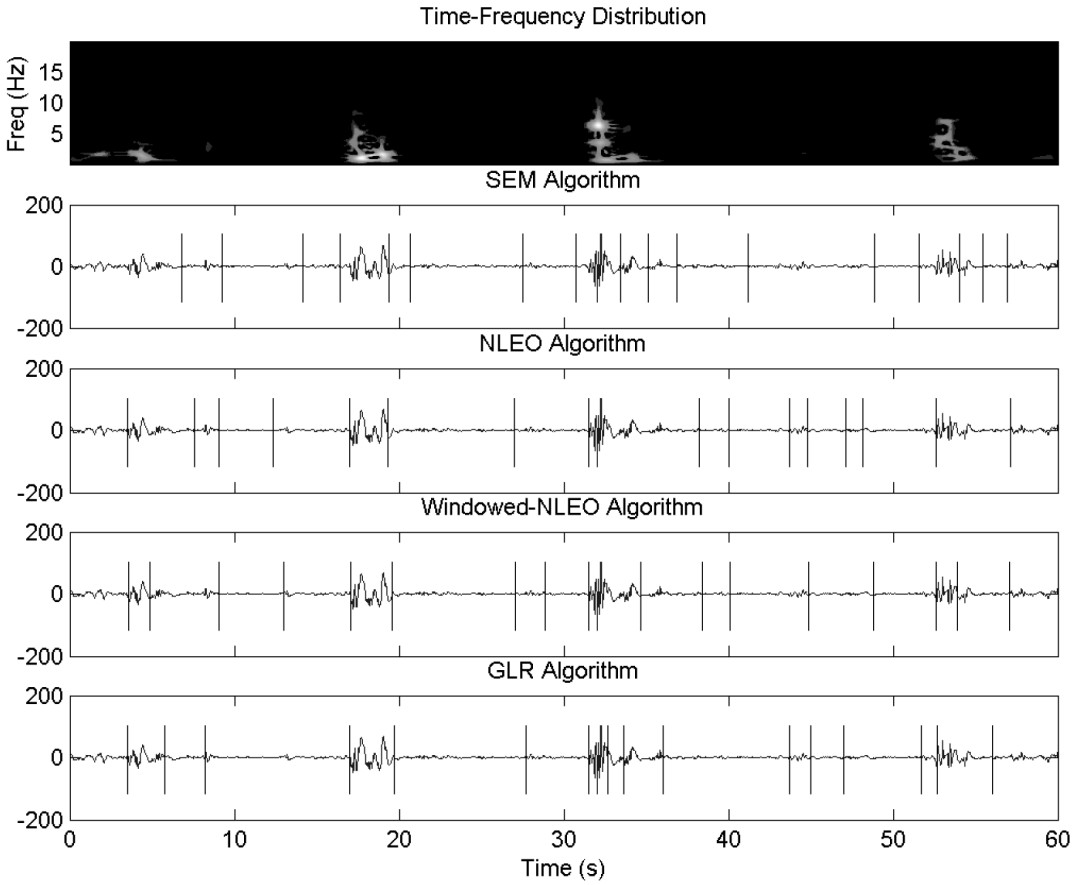

Figure 5.5: Screen shot of the segmentation method evaluation screen

pseudo-stationary blocks, one can see how appropriate the segment boundary is by comparing the segmentation boundaries with the energy contents of the signal in the time-frequency domain. A good segment should not include any abrupt changes in TFD. The frequency components of a good segment should be easily describable in terms of the time-frequency distribution. Figure 5.5 shows a sample of the comparison.

Figure 5.6 shows a close-up of a potential segment. In this example, all algorithms have detected a segment boundary around the area where the discontinuity occurs in the TFD, with the segment boundary found by the GLR most accurately locating the point of discontinuity. The boundary segments were manually evaluated, and classified as one of the following four categories. A segment boundary was classified as "correct" when it reflects the actual position where a discontinuity occurs in the

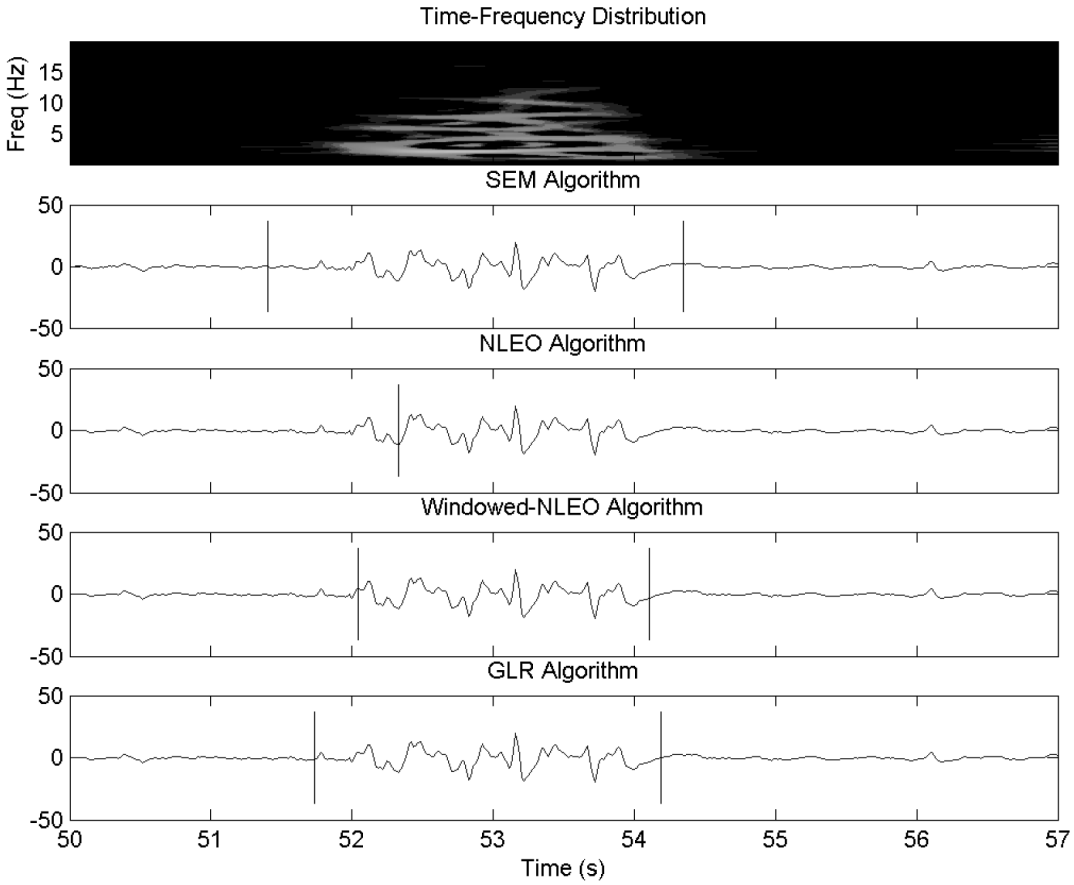

Figure 5.6: Close-up showing the segmentation resulting from the different algorithms

TFD. A boundary was classified as "inaccurate" when the boundary is located in the general proximity of the discontinuity, but not accurately enough to be classified as "correct". A boundary is classified as "incorrect" when no discontinuity in the TFD was found around the boundary, and a failure to detect any boundary for a discontinuity in the TFD was considered a "missed" boundary. Figure 5.7 shows how the boundary categories relate to the discontinuity of the TFD.

In the example shown in Figure 5.7, a boundary detected in the region labelled as "a" is categorised as "correct". A boundary detected in region "b" is categorised as "inaccurate". A boundary detected outside these regions is categorised as "incorrect". If no boundaries are detected in either region "a" or "b", a "missed" boundary is recorded. Table 5.1 shows the performance of the different algorithms. A total

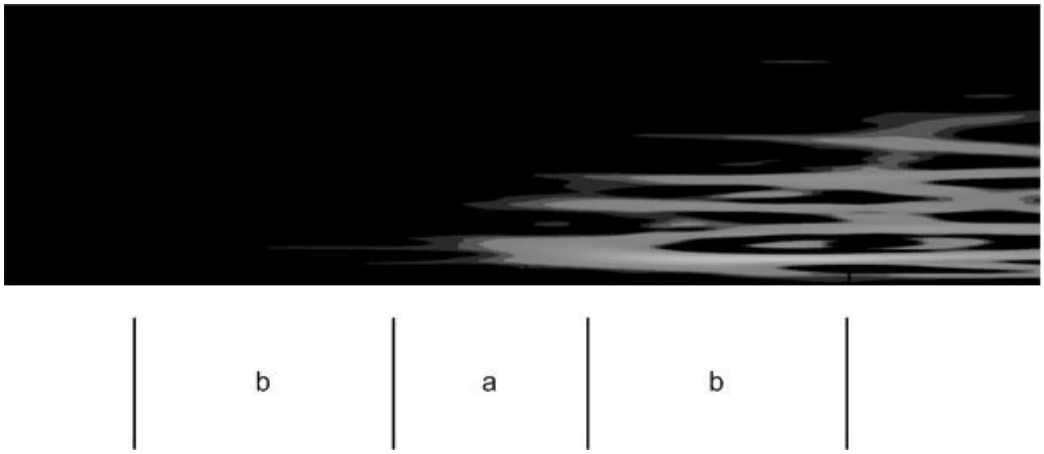

Figure 5.7: The boundary category in relation to the TFD.

Segmentation Method	Correct	Inaccurate	Incorrect	Missed
SEM	129	80	174	22
NLEO	153	18	239	60
w-NLEO	158	10	249	63
GLR	169	19	103	43

Table 5.1: Results from evaluation of the segmentation methods. (No. of manually detected reference boundaries = 231)

number of 231 segment boundaries were manually detected from the sample signals.

5.2.2 Discussion

From table 5.1, the GLR method performed the best in terms of the number of correct segmentation boundaries, and the lowest number of incorrect boundaries detected. It also had a relatively low number of inaccurate boundary detections, and was second in terms of not missing boundaries that should have been detected.

NLEO is the simplest method evaluated, since the method does not involve any autoregressive modelling. As table 5.1 demonstrates, NLEO and the windowed NLEO method performed almost as well as the GLR method, with the windowed version performing slightly better than the non-windowed version, in terms of boundaries correctly detected, and a lower number of inaccurate boundaries. However, the windowed version missed slightly more boundaries than the non-windowed version, and

incorrectly detected more boundaries than the non-windowed version. Depending on the application of the segmentation, this may or may not be a significant defect.

SEM performed the worst in terms of correctly detected boundaries, while missing the lowest number of boundaries out of the different methods. This was caused by the high number of boundaries that were detected, but placed inaccurately. More investigation can be undertaken to improve the accuracy of the algorithm by changing the windowing methods used.

It can be concluded that the GLR algorithm would be most suitable for the proposed task of EEG segmentation. This is possibly due to the windowing method of the algorithm, which takes into account the whole segment rather than just the initial region, or the region immediately prior to the testing window. This results in the highest amount of information to determine whether the segment in the testing window belongs to the current segment. In [53], an online segmentation form of the algorithm was presented that used iterative updates of various variables used in the algorithm. This reduces the computational cost of the algorithm and makes it more suitable for real-time applications.

Also worth noting is the fact that some segments in a relatively low amplitude area of the signal are not visible in the time-frequency domain. While the TFD can be a good tool to determine where abrupt changes occur, since it is a distribution of energy, areas where the amplitude is relatively lower than the rest of the signal will not be shown as clearly. This may account for the seemingly high number of incorrect segment boundaries. Some of these boundaries may be boundaries for segments that are present in the relatively low amplitude regions of the signals.

In terms of the application being considered, namely to provide pseudo-stationary segments for continuity analysis, the goal is to provide segments where the amplitude remains relatively even, such that the segments mimic the way a human EEG expert would mentally divide the signal into blocks. The boundary grading system explained in the previous section helps to identify where these boundaries might be placed without the need for an expert. Since human experts perform the task mentally, it is not always possible for a human expert to mark precisely how they divide up the signal when considering the continuity of a signal. The two clear choices are NLEO and GLR. The decision was made to use GLR as the preferred segmentation algorithm, since its accuracy is slightly better than the NLEO algorithm. GLR is more computationally expensive, however at this exploratory stage of research, accuracy is

more important than computational speed. NLEO also has the disadvantage of being too sensitive to localised transient noise and variation. These localised transients are not usually considered when continuity is concerned, since the continuity of an EEG signal is defined for a window of time in the minute scale, not seconds. Although the windowed version helps to desensitise the algorithm to such local variations, it is at the expense of accuracy. This higher accuracy is why GLR is preferred over the NLEO algorithms.

5.3 Autoregressive Model Order Optimisation

Once the decision is made to use GLR, which uses auto-regressive modelling, the order of the model needs to be determined. To optimise the order of the model used for EEG, the total square error (TSE) is used. The segments are modelled individually, and the predictive error is squared and summed for the duration of the test signal. For this task, only burst suppression and discontinuous signals are tested, as the purpose of the segmentation for this research is to divide the signal into segments where the amplitude stays constant. The order is optimum when the addition of parameters does not improve the TSE significantly.

It is assumed that the segments are stationary, and that they can be modelled using the same number of parameters, hence have the same autoregressive order. There is bias here, since errors in the high amplitude segments would be larger than errors in the low amplitude segments. Since only one threshold is defined in the GLR algorithm, this bias would mean that a segment boundary between two relatively low amplitude segments may not be detected, despite a difference in the AR model. However, since the goal of the segmentation in this application is focused on dividing the signal into segments of various amplitude levels, the bias is not neccessarily significant.

Four burst suppression patterns and five discontinuous patterns have been chosen for this task. The signals were chosen for their lack of artifact interference. Orders from 6 to 32 were tested, and the total square error, as described previously, is plotted against the number of order in Figure 5.8.

As shown on the graph, the rate of decrease slows at around P = 14, where P is the order of the autoregressive model. Although the addition of further parameters can minimise the error even further, the difference is not significant, and adding

more parameters will also increase the computational cost. Therefore, the order for the autoregressive modelling in the GLR algorithm is set as 14. This ensures that the model can sufficiently represent the signal without compromising on the computational cost of the algorithm.

5.4 Summary

In this chapter, three existing segmentation methods were analysed and compared against each other. Spectral Error Measurement (SEM) and Generalised Likelihood Ratio (GLR) methods both use the predictive error in an autoregressive model. SEM uses the predictive error to estimate the spectral differences between the reference and test windows, while GLR uses the predictive error of the reference and testing windows to estimate the likelihood of the test windows being in the same segment as the reference windows. Nonlinear Energy Operator (NLEO) is another method proposed to produce an energy operator that is related to both the energy of the signal and the frequency of the signal. The idea was to place a segment boundary where an abrupt change occurs in the NLEO of the signal. The three segmentation algorithms were compared with the time-frequency distribution of the EEG signal. It was found that GLR performed the best in terms of accuracy.

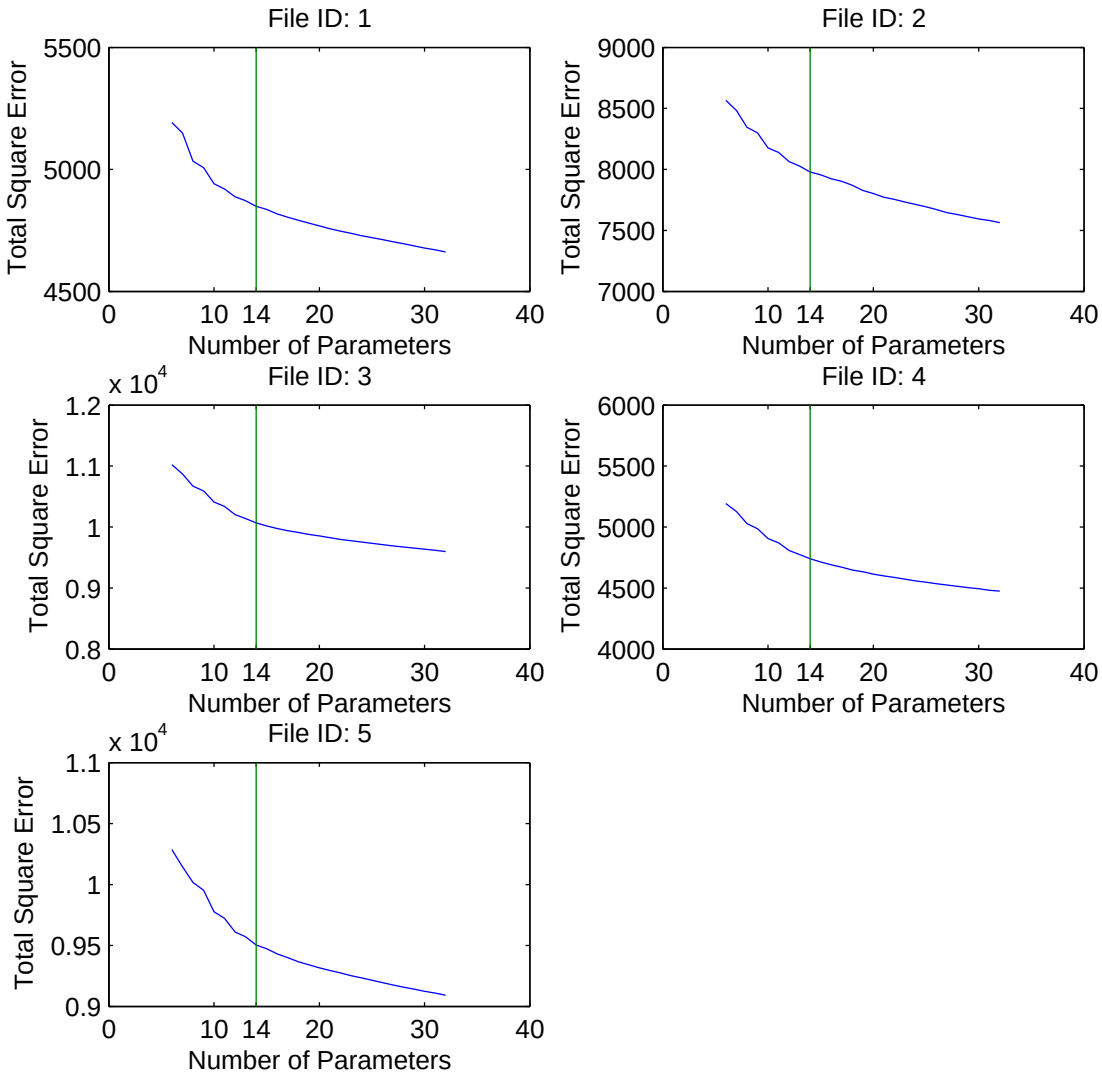

Figure 5.8: Results of order optimisation testing

Chapter 6

Quantifying Continuity

Continuity refers to the amplitude variation of the EEG signal. It is one of the qualities clinicians monitor in order to determine the maturation progress of neonatal EEG. It also provides a context for further EEG analysis, especially since the variation within a single recording can be dramatic, compared with the differences between segments in similar continuity states in EEG recordings from different ages. This chapter examines the quantification of continuity by looking at the amplitude distribution of EEG signals.

6.1 The Need to Quantify Continuity

Continuity, as explained in Chapter 2, is one of the main qualities that clinicians look for when diagnosing an infant. Continuity is related to maturation: preterm infants with a CA of less than 29 weeks show one consistent background state (*tracé discontinu*, or burst suppression) and gradually mature to display a more continuous signal with different sleep stages that exhibit different background states. Clinical research shows that various continuity related features, such as inter-burst intervals, are closely related to the prognosis of an infant. While the task of defining the background state of an EEG signal may be an intuitive assessment for a human, it

remains a qualitative measurement. In order to automate the process, and provide the context for further data analysis, a quantified measure of continuity is needed.

There are several benefits to quantifying the continuity measurement. Firstly, it gives clinicians another value that can be correlated with clinicial outcomes. At present, some continuity-related values are available for clinical researchers to perform statistical analyses on, but not a numerical value for continuity *per se*. A quantified continuity measurement could be used alongside current features to provide extra information. An objective way to define continuity would also assist medical researchers to better compare research outcomes. In terms of clinical value, a quantified continuity could help caregivers without EEG expertise gain insight into the state of an infant, since continuity is related to sleep cycles. In terms of engineering research, a quantified measurement of continuity provides a context, which is needed for further analysis. Since EEG signal is non-stationary, and the characteristics can vary dramatically within a recording, without the context of continuity, the subtle difference between two recordings during periods with similar background states is far less than the variation within the recording itself. With a quantified measurement of continuity, areas with similar continuity states can be identified and isolated for further processing, meaning that other information unrelated to continuity (e.g. frequencies, entropy) can be compared.

6.2 Current Methods of Continuity Measurement

The most common approach to determine the background states of a recording is visually scanning the EEG. Clinicians do so by quickly scrolling through the recording to get an idea of how the amplitude varies over the time of the recording. Only qualitative labels are assigned for continuity using this method.

Amplitude-integrated EEG (aEEG) [55, 9], which is a compressed display of an estimated signal envelope, is another way to determine continuity. This approach uses the general guidelines shown in Figure 6.1 to determine the background states [9]. Currently, the definitions of continuity labels, such as “continuous” and “discontinuous”, are qualitative. Some guidelines have been established to estimate the maximum and minimum points of the aEEG, which are in turn used to determine the background state. The label assigned is based on the estimated lower band (the bottom edge of the aEEG signal as it appears on an aEEG plot) and the upper band

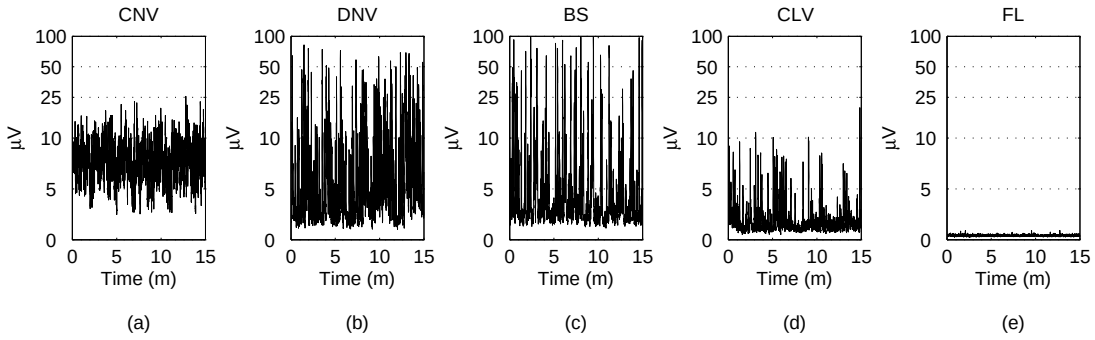

Figure 6.1: Guidelines for background state classification by aEEG. The two horizontonal lines indicate $5\mu V$ and $10\mu V$. From left to right: (a) Continuous normal voltage; (b) Discontinuous normal voltage; (c) Burst Suppression; (d) Continuous low voltage; (e) Flatline. The estimated maximum and minimum of the aEEG is compared with the $5\mu V$ and $10\mu V$ thresholds as classification criteria. Note the conventional semi-log scale used for aEEG (i.e. linear from 0 to 10 μV and in log scale from 10 - 100 μV)

(the top edge of the aEEG signal as it appears on an aEEG plot), and where these bands are relative to 5μV and 10μV. The label "continuous" is given to signals with an aEEG where the lower band is between 5μV and 10μV, and the upper band above 10μV. "Discontinuous" is assigned to signals with aEEG lower band below 5μV and upper band above 10μV. "Burst Suppression" refers to signals with a lower band close to 0μ, an upper band below 5μV , and with "spikes" above 10μV appearing in the aEEG plot. "Continuous low voltage", an abnormal EEG pattern, is assigned to aEEG patterns where an upper band is not obvious, and the plot has a lower band very close to 0μV , with occasional spikes reaching between 5μV and 10μV . However, these are estimates and are largely subjective.

Medical researchers use measurements such as interburst interval (IBI) and burst amplitude as indications of continuity. As discussed in the literature review in Chapter 3, measurements such as IBI and burst amplitude correlate with maturation, and can be used for prognostic purposes. These measurements usually cover a wide range, so in general only the maximum/minimum or the median is used for analysis [4, 17]. Therefore, while the values give an indication of the continuity of the signal, they serve as an indication of the global continuity of the recording, rather than showing the changes of continuity with respect to the recording period.

Another way medical researchers measure continuity quantitatively is to use intensity measurements. Certain thresholds, usually 10μV, 25μV, 50μV, and 100μV,

are used, and the proportions of the signal that lie under the thresholds are recorded. Usually, the values corresponding to every threshold are used, so each recording gives four different values that indicate the proportions of time during which the signal is lower than the different threshold levels. Like other burst suppression related features such as IBI, the intensity measurement only provides a global indication of the signal.

The algorithm presented by Navakatikyan et al [38] trys to improve the objectivity of the guidelines by using rEEG [37] and sets a standard to determine the maximum and minimum amplitudes using the rEEG. This provides a more objective way to classify the signal, but rEEG, like aEEG, is still very sensitive to local signal variations, and can also be prone to the interference of noise such as muscle artifact. The algorithm presented in [38] focuses on term infants, who have more developed brains and generally clearer distinctions between continuous and discontinuous signals.

6.3 The Continuity Feature

To quantify continuity, one has to understand what the clinical definition of "continuity" is, and how it can be translated into a mathematical quantity. Continuity, in the clinical sense, is the inverse measure of the variability of the amplitude of the EEG signal in a certain time span. However, the current labels are only a guideline, and the classification process is highly subjective. This is discussed in detail in Chapter 2. In terms of a quantified version of continuity measurement, the focus should be put on the statistical distribution of the amplitude, which provides information about the variability of the amplitude.

6.3.1 Amplitude Vector

In order to determine the continuity feature, the term "amplitude" needs to be defined. In Chapter 5, the concept of segmenting the EEG signal into psuedo-stationary sections is discussed, and this can form a basis for defining the amplitude of the signal. The advantage of using the psuedo-stationary segments, rather than looking at the envelope of the signal, is that it mimics the way clinicians define continuity, as discussed in Section 2.1.1.

To produce a vector to represent the amplitude of the signal, the GLR segmentation method is first used to segment the EEG signal. The signal is first processed

using a high pass filter with a stop band of 0.01 Hz and a pass band of 0.05 Hz to filter out any DC component. A low pass filter with a pass band of 10 Hz and stop band of 15 Hz is applied in an attempt to filter out any high frequency noise such as muscle artifact. The cut-off frequency is set as 10 Hz, as most energy in neonatal EEG is present at frequencies lower than 10 Hz, as shown in the initial study with time-frequency analysis. Both filters are Butterworth filters designed using the MATLAB built in Filter Order Determination Function. Butterworth filters were used because of the flat frequency response in the passband, and a relatively linear phase response. Each filtered segment in the EEG recording is then rectified, and its mean value is used to represent the amplitude of that segment. For segment i, the averaged absolute, v_i, is defined as:

$$v_i = \frac{1}{N_i} \sum_{n \in k_i} |s_n| \tag{6.1}$$

where k_i is a set of time indices in the EEG signal that belong to segment i, and N_i is the number of samples in segment i.

The mean value is preferred over the power because power of a signal is proportional to the amplitude squared. The use of the signal power would therefore make the algorithm more sensitive to a difference in amplitude between two high amplitude segments than the difference between two low amplitude segments. However, the medical emphasis on the analysis of EEG amplitude is placed on the difference in amplitude between two low amplitude segments. A change in amplitude in low amplitude segments is more significant for prognosis and diagnosis, and is one of the deciding factors of labelling continuity states such as "discontinuous" and "burst suppression". While artifacts such as electrode pop would affect the averaged absolute voltage, the effect would not be as noticeable. Although noise rejection algorithms can help eliminate the effects of artifacts, such an approach is beyond the scope of this thesis.

The mean absolute voltage is used to represent the segment. The resulting amplitude vector, $\mathbf{a}$, is a vector with the same length as the original EEG signal, where:

$$a_n = v_i|_{n \in k_i} \tag{6.2}$$

An example of the contruction of this amplitude vector is shown in Figure 6.2.

Note that the mean absolute value of each segment is repeated for the duration of the segment. This takes into account the length of the segment as well as the amplitude, and the resultant amplitude vector is therefore an estimate of the signal envelope. The average of the segment is used instead of the extrema values such as peak-to-peak amplitude to minimise the effect of occasional transients in EEG signal, such as spikes or sharp waves.

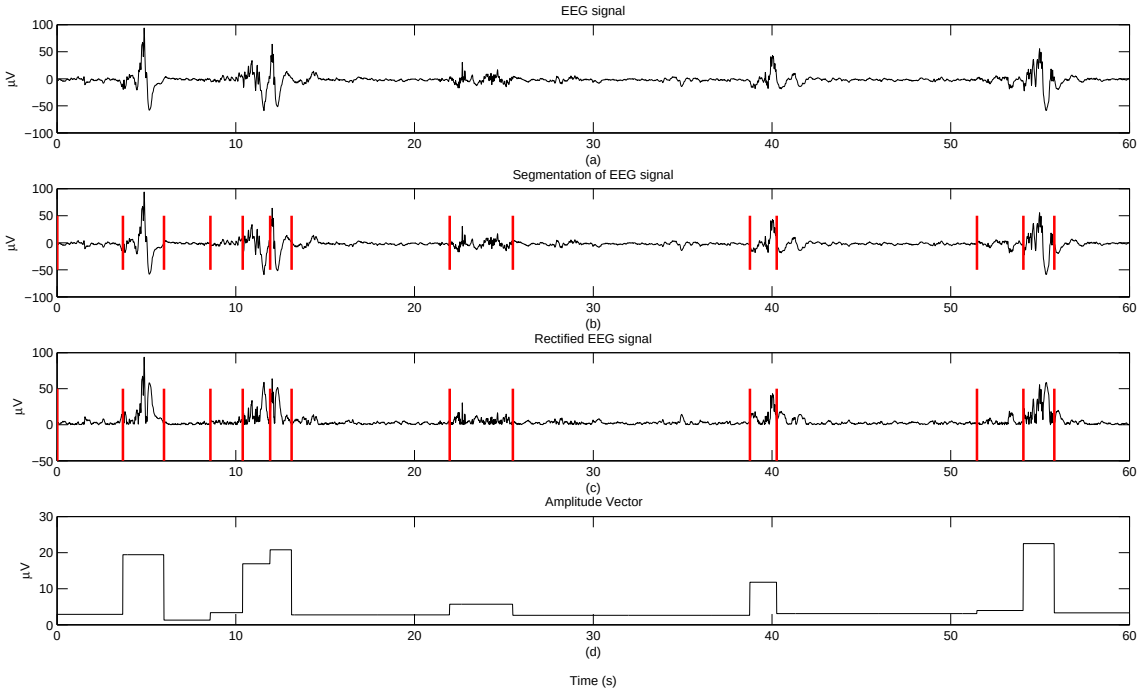

Figure 6.2: Construction of an amplitude vector. (a) The original EEG signal (b) Segmentation results from GLR algorithm (c) Absolute value of the EEG signal (d) Resultant amplitude vector, with the mean value of each segment used to represent the estimated amplitude for the duration of the segment

6.3.2 Modelling Amplitude Distribution

After defining the amplitude in terms of continuity criteria, the distribution of the amplitude vector, **a**, is examined. Since continuity is a measurement that describes the EEG signal during a certain window, a moving window of a predefined length should be used. Figure 6.3 shows a histogram of the amplitude vector from a discontinuous EEG segment, using various window lengths.

The trade-off for the length of the window is between resolution and providing

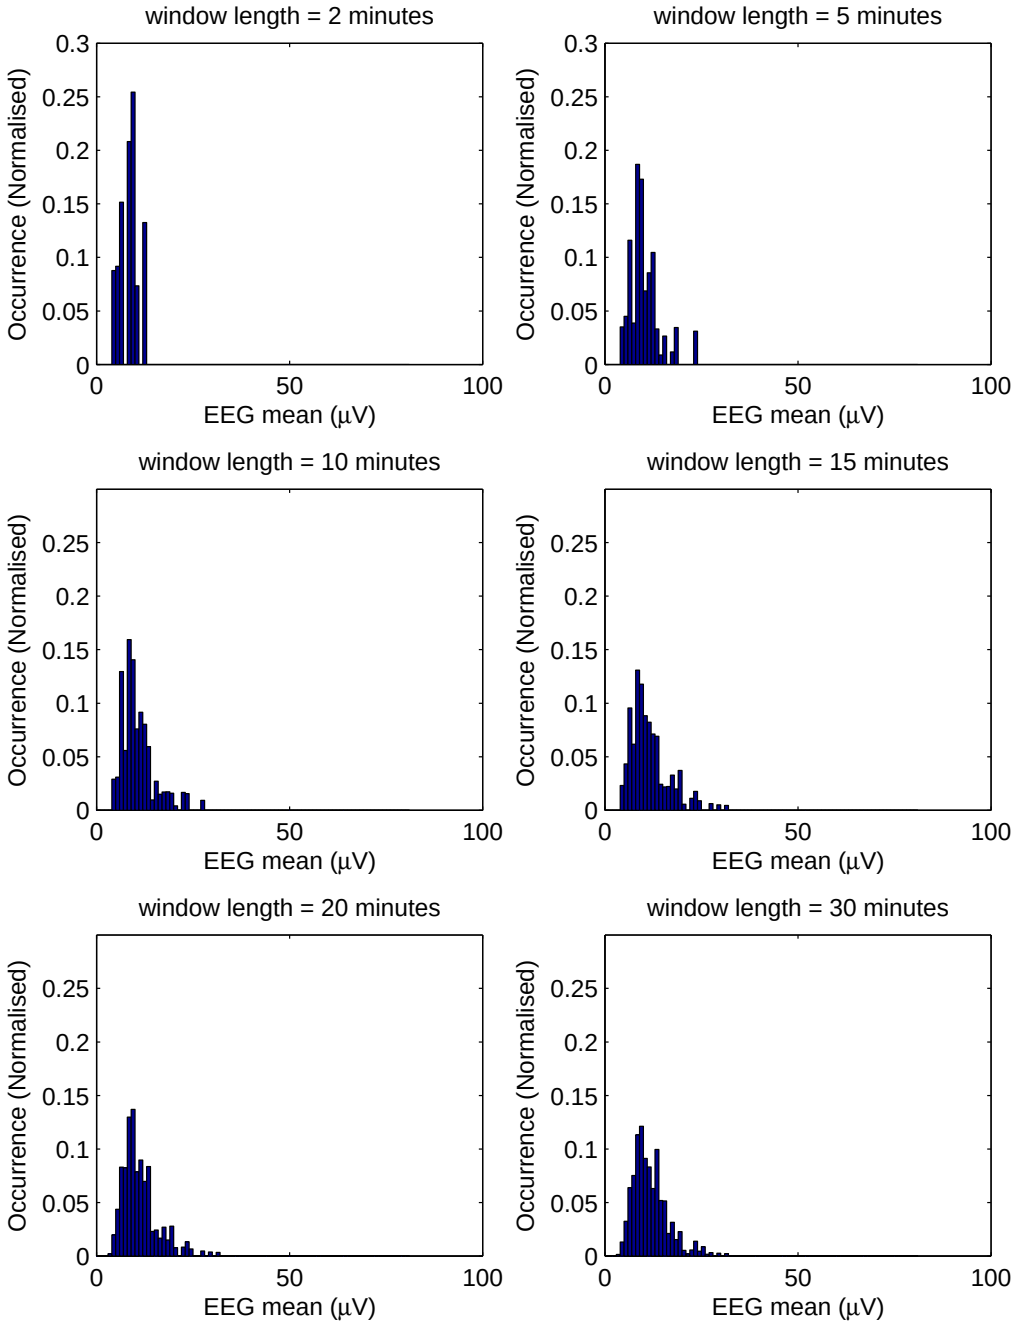

Figure 6.3: Histogram of amplitude vector with various window lengths

enough data to make the window carry enough information about the signal. In theory, as the length of the window increases, the amplitude distribution should approach a smooth curve, representing the amplitude distribution of the underlying EEG mechanism. However, during the change of states (e.g. from continuous to discontinuous EEG, which happens during sleep), a long window will blur the boundary between the two states. A window length of 10-minutes was chosen for the system since it provides a good amount of data for the system to model the ideal distribution, without being excessively long and compromising the resolution of the feature to be extracted from the distribution.

With this information, the distribution to model the amplitude can be chosen. The log-normal distribution was selected to model the data because of the skewed nature of the distribution, and its suitability for modelling biological distributions [56]. The probability density function of the log-normal distribution is defined as:

$$f(x;\mu,\sigma) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x-\mu)^2}{2\sigma^2}} \tag{6.3}$$

where μ and σ are the theoretical mean and standard deviation for the logarithm of the variable x. In this instance, x will be the windowed amplitude vector, $\mathbf{a}$, as described in (6.2). To estimate μ and σ, the maximum likelihood estimation of the parameters is used. Thus:

$$\hat{\mu} = \frac{\sum_k \ln x_k}{n} \tag{6.4}$$

$$\hat{\sigma}^2 = \frac{\sum_k (\ln x_k - \hat{\mu})^2}{n} \tag{6.5}$$

Examples of the distribution histogram and the estimated probability density function are shown in Figure 6.4.

Although the estimated distribution does not give the exact error free estimation of the amplitude distribution, it has the advantage that the two parameters controlling the shape of the probability density function give a good indication of the peak of the distribution ($\hat{\mu}$) and the spread of the distribution ($\hat{\sigma}$). The skewness of the distribution is also related to the parameter $\hat{\sigma}$. Another way to consider the distribution estimate is by looking at the histogram of the logorithm of the amplitude vector, since a log-normal distribution assumes that the logorithm of the data

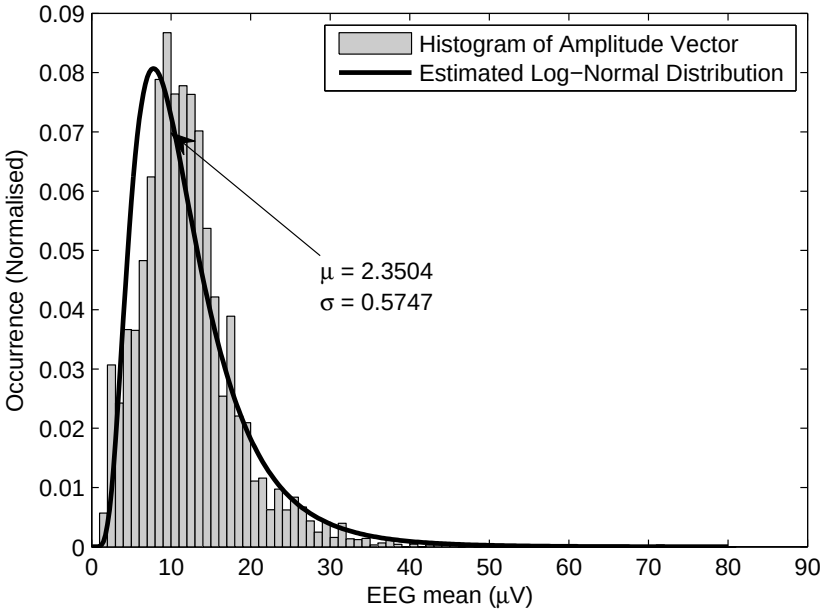

Figure 6.4: Example of a histogram of the amplitude vector (normalised to unity area)

follows the normal distribution. Figure 6.5 shows the histogram of the logorithm of the amplitude vectors shown in Figure 6.4.

6.3.3 Distribution Parameters as a Continuity Feature

After modelling the amplitude distribution with the log-normal distribution, the parameters of the log-normal distribution can be plotted and compared between segments with different continuity states [57]. Segments with 10-minute length and manually determined continuity states were processed and their estimated distribution parameters plotted in a scattered plot in Figure 6.6.

The plot shows a correlation between the continuity and the parameters of the distributions. Furthermore, the age of the infant at the time of the EEG recording is also plotted against the distribution parameters, shown in Figure 6.7. It shows a correlation between the distribution parameters and the age of the infant. Because the dataset used in producing the graphs includes different infants, the correlation might be affected by the slightly different rate of maturation between the different infants. The effect of maturation on the distribution parameters will be further investigated in Chapter 7.

6.4 Displaying Continuity

Each recording can be processed to produce a continuity feature corresponding to the EEG signal for the duration of the recording. A 10-minute sliding window with step size of 1 minute is used on each channel to produce a continuity feature with the effective resolution of one set of features per minute. The feature is plotted in various ways to highlight different aspects of the continuity progression.

6.4.1 Line Plot of the Continuity Feature

The continuity feature can be plotted against time and produce a graph similar to aEEG. The estimated mean ($\hat{\mu}$) is plotted against time, and two lines are plotted for one standard deviation ($\hat{\sigma}$) above and below the estimated mean. An example is shown in Figure 6.8 with the aEEG of the same signal. The aEEG is used for comparison rather than EEG because aEEG is more compact and can be plotted on a similar time scale as the continuity feature.

Using this plotting convention, the continuity feature plot displays similar characteristics to the aEEG display. Continuous signal resulted in a narrow band (low

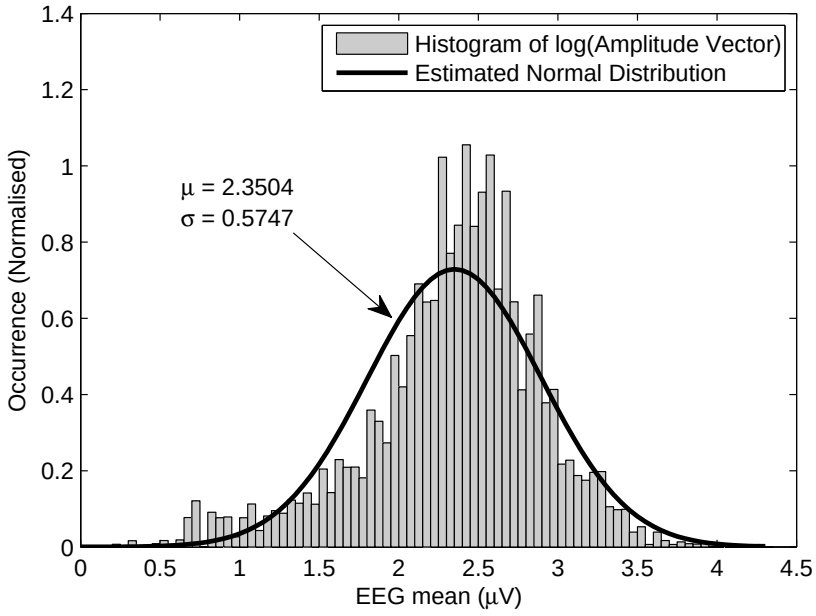

Figure 6.5: Example of a histogram of the logorithm of the amplitude vector (normalised to unity area)

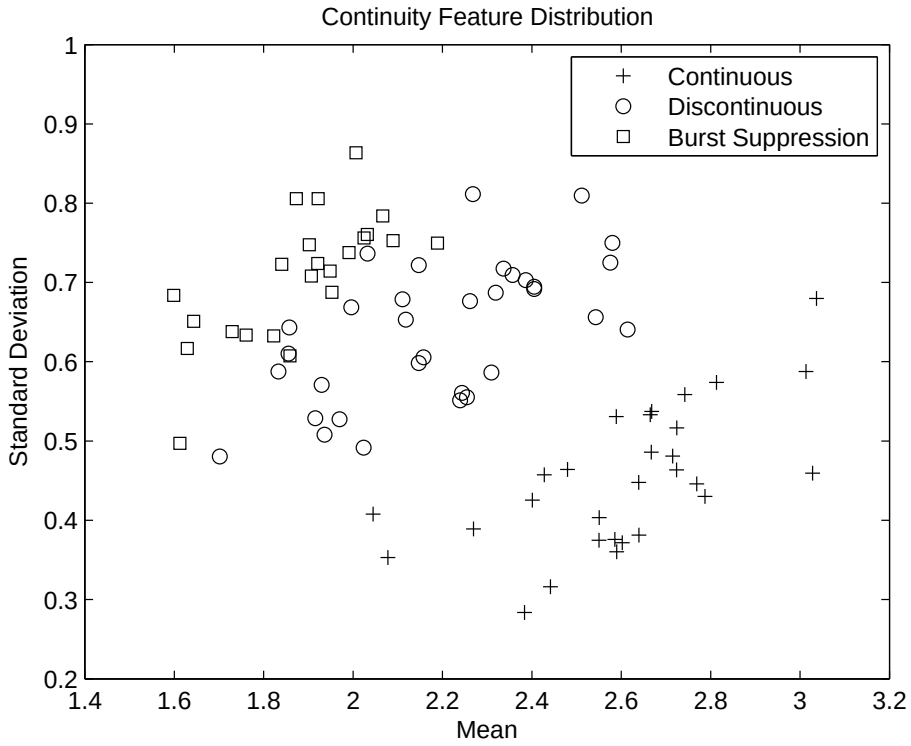

Figure 6.6: Scatter plot of amplitude distribution parameters for 10 minute segments of different background continuity states, taken from infants with healthy clinical followup between ages of 25 to 35 weeks CA.

standard deviation) higher on the graph (high mean). As continuity decreases, the lower band of aEEG and the lower standard deviation line in the feature plot both decrease in value, and the "bandwidth" of the aEEG plot and the distance between the two standard deviation feature plots increases. This is predictable since, assuming the amplitude of segments follows the log-normal amplitude, the minimum amplitude (represented by the "lower band" of the aEEG plot) and the estimated mean are related. Although the mean of the distribution is, in practice, the mean of the logarithm of amplitude, the fact that aEEG is displayed by a "semi-log" scale (linear from 0 - 10 μV and log scale for amplitude $> 10\mu$V) means that visually the minimal voltage displayed in aEEG and the mean voltage in the feature plots are in the same order of magnitude. The "bandwidth" of the aEEG, which is the thickness of the graph, is in effect the range of the amplitude, and is related to the standard deviation of the distribution. The use of standard deviation over range means that a transient event such as a spike has a lesser effect. The resulting feature produces a graph much smoother than an aEEG. Although the established guidelines for aEEG cannot be

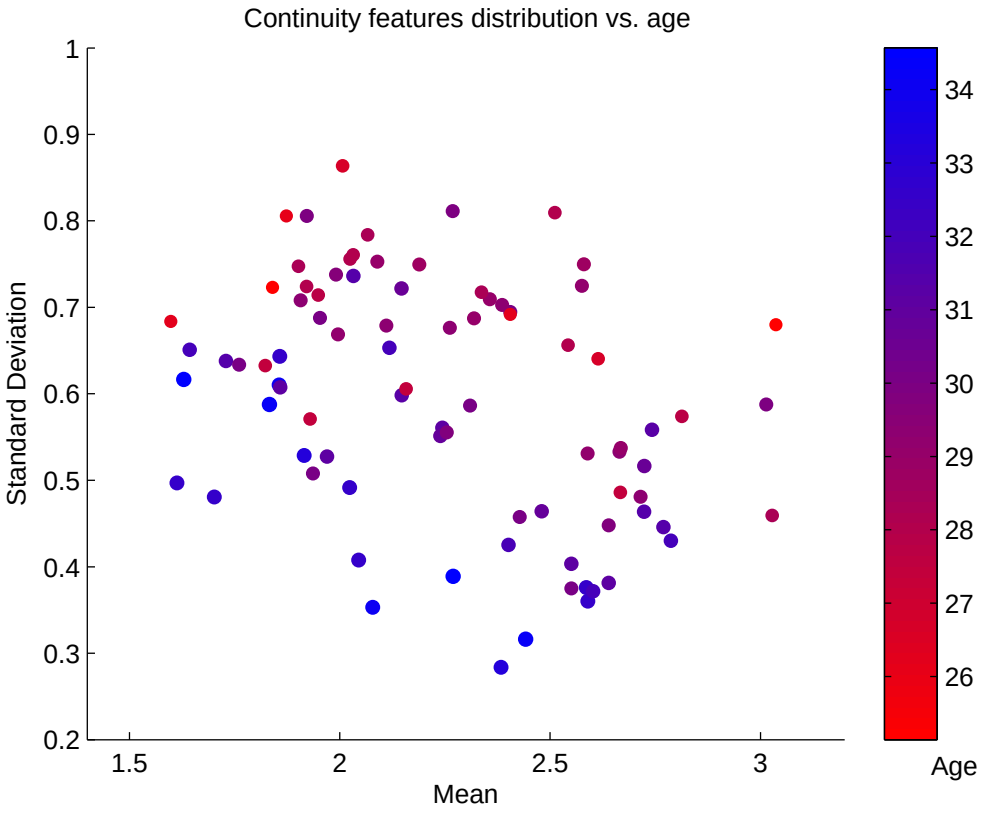

Figure 6.7: Scatter plot of amplitude distribution parameters for infants of varying CA

directly applied to the feature plot, the general trend is similar enough that a clinician can quickly adapt the guidelines to work with this new plot.

6.4.2 Scatter Plot of the Continuity Feature

Another way to display the feature is to plot it in a scatter plot, without the time reference. This display is best for comparison between different recordings, to show the general continuity distributions. A line can be drawn to show the passage of time if needed. This also allows a visual guide as to whether any intrinsic clusters are present, to form a better definition of a background state. Figure 6.9 shows a scatter plot of the feature in one single recording. The graph shows that the data clusters at the one end of the major axis of the dataset, being the end corresponding to continuous background activity. The line in the plot shows the passage of time, and together with the data points they show that the change between the two states

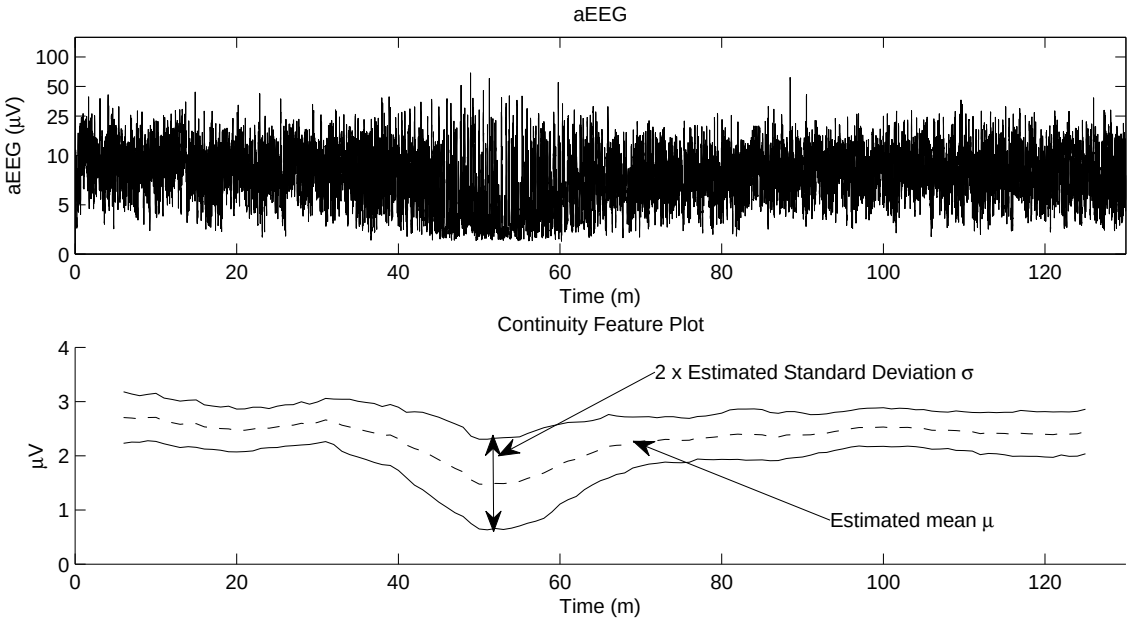

Figure 6.8: Continuity feature plot with aEEG plot as reference

is gradual and the change from one state to another is a smooth transition. Once the signal has reached a stable state the continuity stays relatively stable until another change in state occurs. These changes in states are explored further in Chapter 8.

Figure 6.10 shows the scatter plot of the features for four different recordings at different times from the same preterm infant.

The changes in the feature are similar to those shown in the plots of other infants. As the infant grows towards term, the trend observed is that both the mean and standard deviation decrease. This maturation trend is more obvious at the continuous end of the datasets than at the discontinuous end of the datasets.

6.5 Quantifying Continuity Using Principal Component Analysis

Once the correlation between continuity in the clinical sense and the continuity feature extracted from the distribution has been established, a quantified approach to continuity can be devised.

The most intuitive way to extract continuity information out of the feature described in this chapter is to use Principal Component Analysis (PCA) to transform

the feature into another domain. PCA is used to determine a new coordinate system which will result in components with the highest variance as the first domain, and subsequent components as orthogonal domains in order of the variance in data. Looking at Figure 6.6, the major axis of the data distribution is also the orientation best representing the continuity of the data. This means that after using PCA on the feature, the first component will be a good indicator of continuity. The problem with this approach is that, while the first principal component does indeed give a good representation of continuity, it is record specific. It also assumes the dataset being transformed provides various different continuity background states. Should a recording only contain data in one particular state, the principal component may not reflect the continuity of the data. Another problem is that this continuity measure is only relative to the data used to generate the PCA transform. This means that, should the continuity of two different recordings be analysed separately, the principal components cannot be compared between the recordings, since there is no guarantee that the PCA transformation is the same for the two datasets.

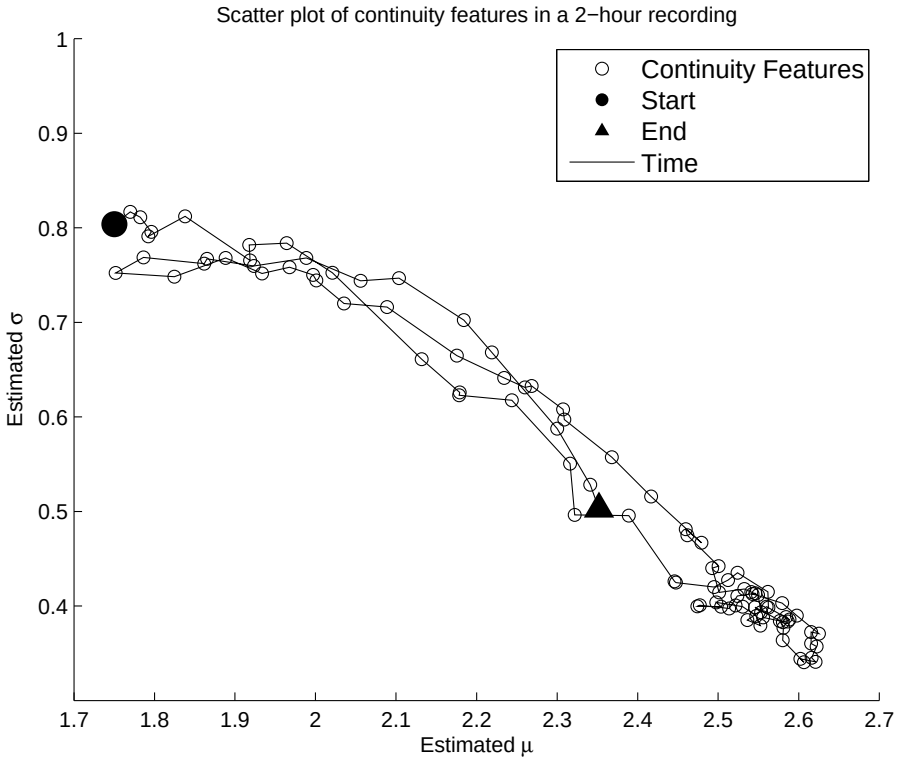

Figure 6.9: The continuity feature in a 2 hour recording.

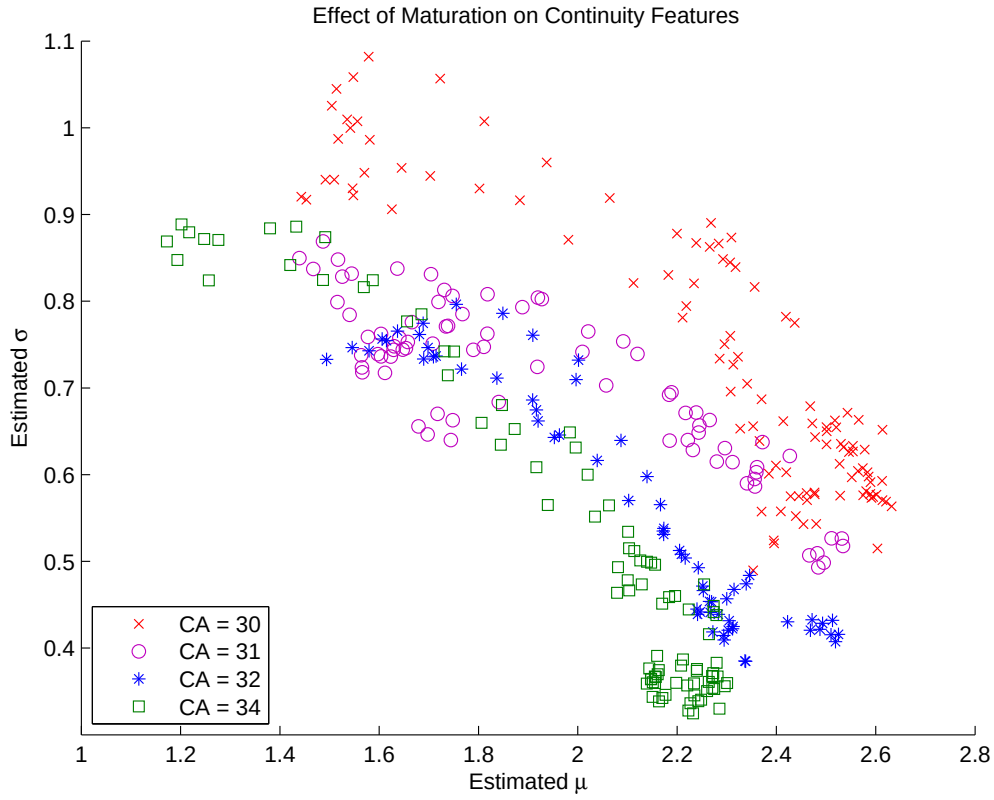

Figure 6.10: Scatter plot showing changes in continuity feature for one healthy infant during CA 30 to 34 weeks.

One way to overcome this problem is to use features from recordings of a number of infants to come up with a transformation that will be used for every file. With a standardised transform, the continuity index (in this case, the first principal component) can be compared across different recordings. With this in mind, the dataset used in Section 6.3.3 was used to form the basis of the PCA transformation used for continuity quantification. The dataset contains data from infants with CA from 25 to 35 weeks, which is approximately the range of CA available in the database of preterm EEG used in the study. From each recording, 10 minute segments of continuous, discontinuous and burst suppression patterns (where available) were used. The continuity feature as described in Section 6.3.3 was calculated for each segment, and PCA was performed on the set of features from all of the segments. The training dataset was normalised to ensure that both dimensions of the feature ($\hat{\mu}$ and $\hat{\sigma}$) have a mean of 0 and standard deviation of 1. Figure 6.11 shows the resulting new coordinate system that forms the PCA transformation of the dataset.

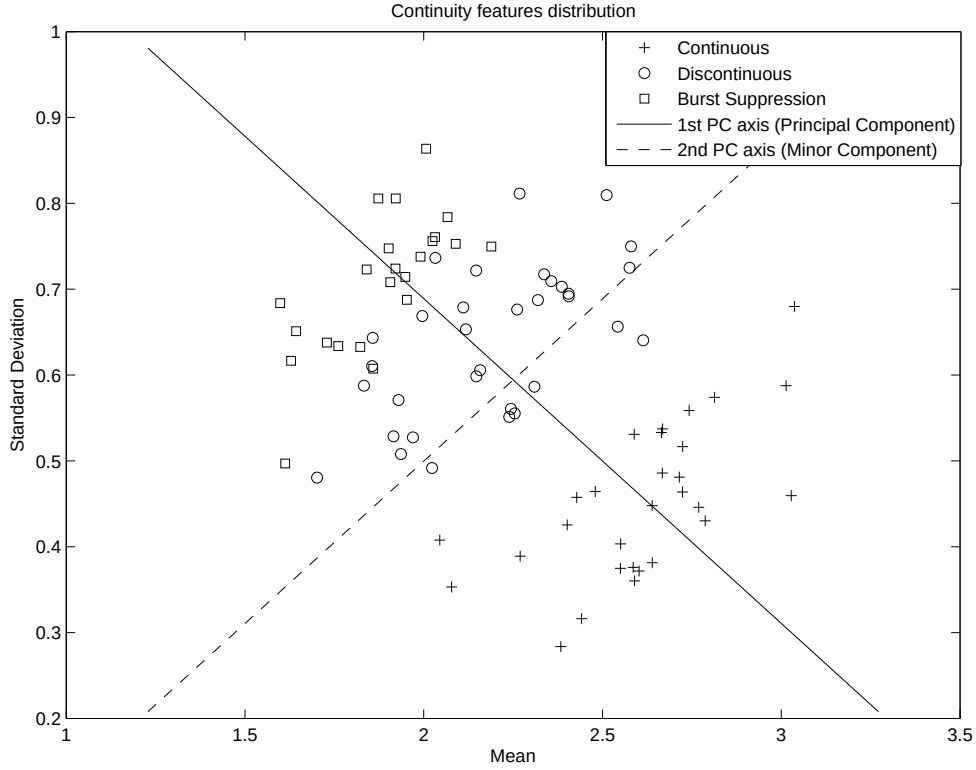

Figure 6.11: Principal components of the training dataset.

The principal component calculated from this dataset represents the orientation that gives the largest variance. This also coincides with the direction that indicates continuity. Using the transformation, the feature extracted from a recording can be transformed to give a quantified measurement that indicates continuity. This gives continuity information in a more localised fashion, as opposed to the global measurement used for maturation studies discussed in the medical literature review in Chapter 3. Figure 6.12 shows a recording of an infant at 34 weeks CA and the principal component against the aEEG and the continuity grading from Navakatikyan's classification algorithm of the same recording. The PCA transformation of the continuity feature, as discussed in Section 6.3.3, is normalised using the same mean and standard deviation as the training dataset. Both the principal and minor component are displayed for discussion purposes.

Lacking a reference for a quantified continuity measure, visual scanning of aEEG remained the best way to analyse the continuity of the signal. The Navakatikyan rEEG related grading was designed and trained using term infant EEG recordings.

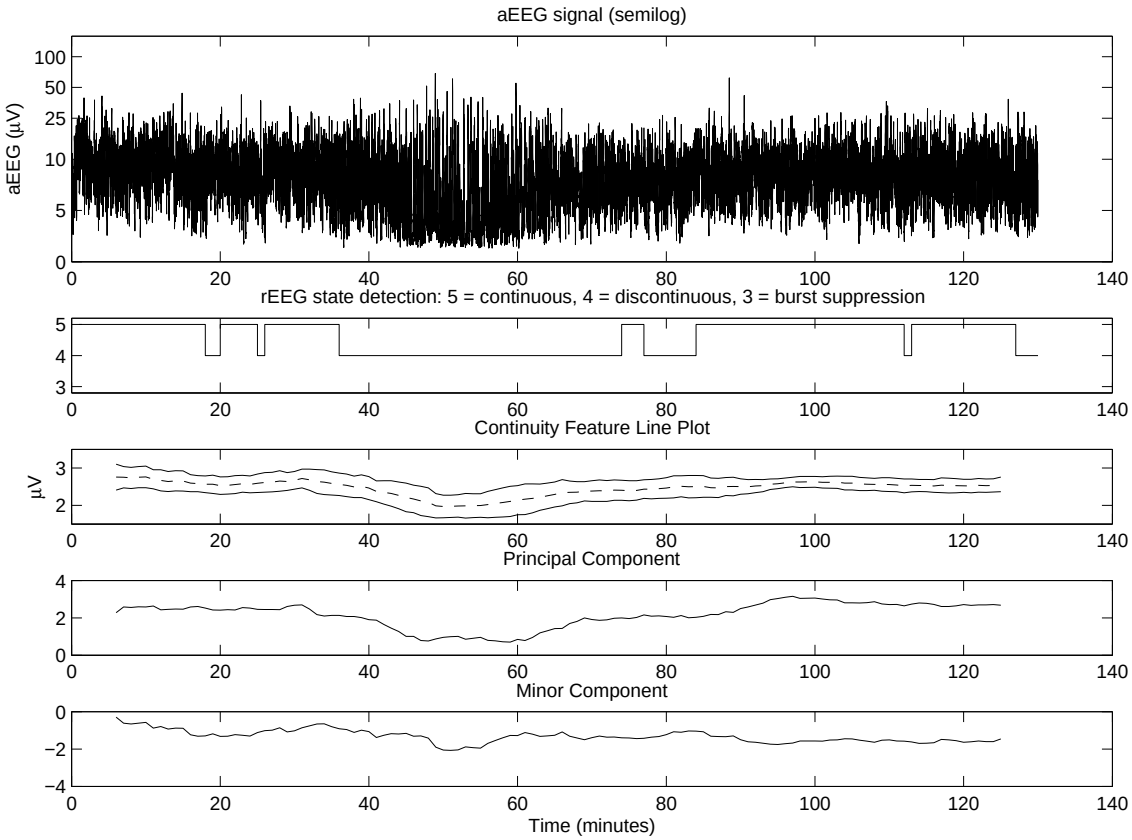

Figure 6.12: Principal component transformation of the continuity feature compared with aEEG and rEEG classification

While not a perfect representation of the background states for preterm infants, it provides a comparison of how the quantified continuity measure relates to a current automatic labelling method. As shown in Figure 6.12, the rEEG method demonstrates some uncertainty around regions where the EEG is in transition from one state to another. The continuity measure, however, gives a smoother transition, and provides a better indication as to the rate at which the states are changing, which is not shown in a qualitative approach. The continuity measurement is higher in value in continuous areas and lower in value in discontinuous areas. Figure 6.13 shows a comparison between the rEEG state labelling and the continuity index as described. Ten files from infants with healthy clinical follow-ups, aged between 32 to 38 weeks CA were compared, with recordings lasting approximately 180 minutes. As shown in the plot, although there are some overlaps between the state labels as assigned by the rEEG algorithm, the continuous states have generally higher values in the continuity index then the discontinuous states. Recall from Figure 6.12, where a

plot showing the state assignment is displayed, the rEEG labels are very sensitive to local variation and when performed on preterm infants the algorithm can mislabel the signal. Because the standard deviation is used in the continuity index, instead of the maximum and minimum values of the EEG signal, the effect of outliers and local transients is minimised and the transition between states is shown as a smoother transition with less fluctuation. This may explain the overlaps in the continuity index in the states assigned by the rEEG method.

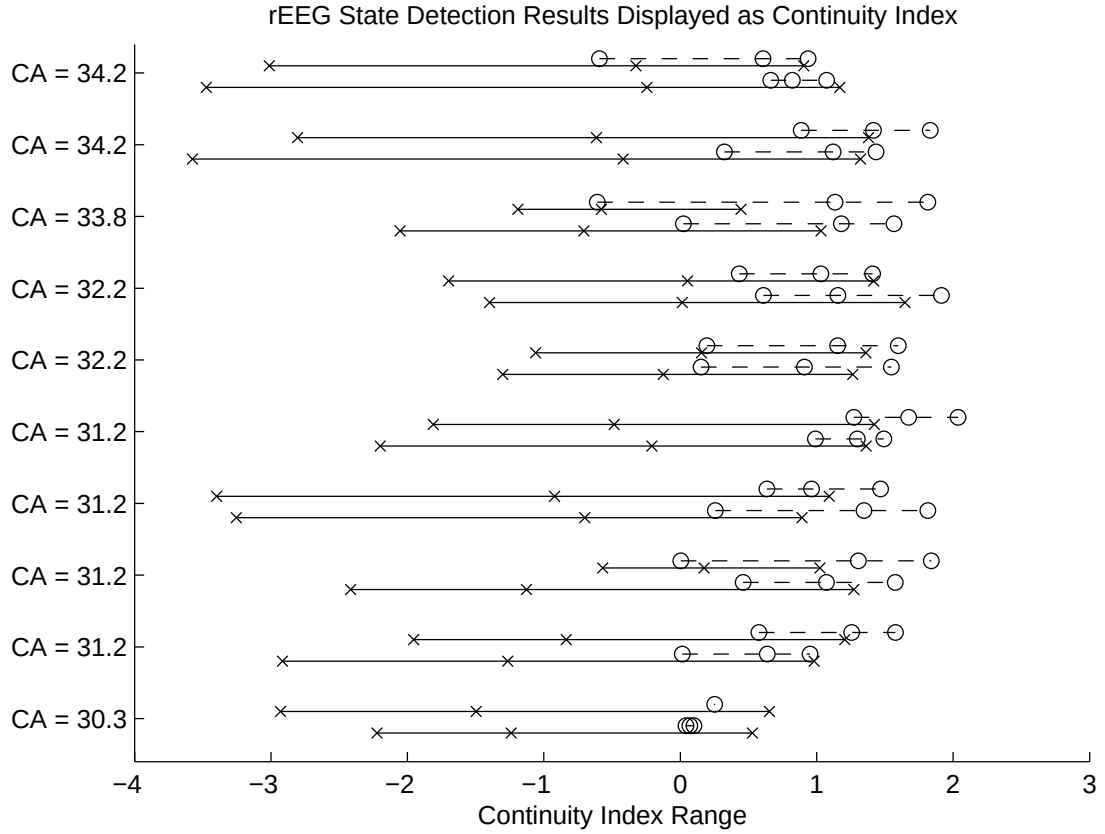

Figure 6.13: Plot showing the continuity index range for each background state assigned by Navakatikyan's classification algorithm. Solid lines represent discontinuous states and dashed lines represent continuous states. The minimum, average and maximum values are shown as plotted values on the lines. Two pairs of lines are plotted for each file. Each pair (solid and dashed) represent one channel of the EEG (left or right hemisphere).

The minor component does not reflect any continuity information, since the minor component is orthogonal to the principal component, which corresponds to the degree of continuity. The variation of the minor component during the recording is less than that of the principal component, which is expected since the principal component is

the projection of the feature along the major axis. As discussed briefly in Section 6.3.3, the minor component is related to maturation, therefore the minor component should not be changing dramatically during the recording. The relationship between the minor component and maturation is further discussed in Chapter 7.

6.6 Summary

Continuity is an inverse measure of amplitude variance within a recording. Currently, it is a largely qualitative measurement, and the only quantitative measurements available are simple measurements obtained by using a few conventional threshold levels to represent the proportion of the signal that lies beneath each threshold. This chapter proposed a quantified continuity measure using statistical distribution parameters of the amplitude as the continuity feature. An amplitude vector is defined by segmenting the signal into psuedo-stationary segments, and the mean absolute value of each segment is used as a representation of the amplitude of the segment. This amplitude vector is modelled using the log-normal distribution, and the estimated parameters, namely the mean and standard deviation of the logarithm of the amplitude vector, are used as the continuity feature. A 10-minute sliding window at a 1-minute rate is used to ensure enough information is available for modelling. Segments with known qualitative continuity states are modelled and their continuity features plotted in a scatter plot in order to establish the relationship between the feature and the continuity of the segment. It is observed that continuous segments have relatively higher means and lower standard deviations, and that the continuity of the segment can be modelled by the first principal component of the feature dataset. The raw feature and the principal component can be used to represent continuity and can be displayed in different ways to help clinicians analyse EEG signals.

Chapter 7

The Continuity Feature and Maturation

As discussed in the literature review, continuity is one of the markers for maturation in preterm EEG recordings. With the continuity features extracted from the EEG, the continuity of the signal can be quantified and correlated with the age of the infants. The principal component from the features provides a way to quantify continuity relative to the recording, but stripped of age information. This age related information, however, can be retrieved from the minor component, as shown in the initial investigation of the previous chapter. In this chapter, the second PCA coefficient is used as a maturation indicator, and the correlation between this coefficient and maturation is examined. The effect of brain injuries such as white matter injury on maturation is also explored by examining how this maturation index changes for infants affected by such injuries.

7.1 Maturation and Continuity

Continuity is one of the major markers of maturation in neonatal EEG. As described in Chapter 2, by looking at the continuity of the EEG signal, a human expert can

estimate the CA of the infant, and thus determine whether the infant is developing neurologically at a normal rate. EEGs of very preterm infants are highly discontinuous, and gradually become more continuous as the infants grow towards term. Sleep-wake cycles, which consist of alternating periods of continuous and discontinuous EEG, also start appearing around 30 weeks CA. The presence of sleep-wake cycles is also important when determining the maturation of the EEG signal.

Although there are clinical guidelines for EEG interpretation, these guidelines are largely qualitative, and require years of experience in EEG interpretation to accurately determine the CA of the infants. A lot of the interpretation is subjective and relies on, to a certain degree, intuition on the clinician's part. Human brains are very good at pattern recognition and therefore experts can be trained to recognise certain patterns relating to the maturation process during different stages.

Using the feature described in Chapter 6, a parallel can be drawn between the way clinical experts view EEG, and the mathematical attempt to define how maturation affects the continuity of the EEG signal.

7.2 Changes in the Continuity Feature Throughout Maturation

In order to define how maturation can affect the continuity of the EEG, one needs to ensure that the changes detected are those resulting from normal maturation. This is difficult in preterm research because preterm infants were born before the normal gestational age (GA) is reached, and thus by definition are not normal in the clinical sense. The goal is, therefore, to determine which infants have no brain injury and show no adverse effects in a two-year neurological followup. The data used in this study was collected by A/Prof Terrie Inder for the Victorian Infant Brain Study (VIBeS) in Melbourne, Australia. Infants enrolled in the study had four EEG recordings taken within about a month of birth, and MRI scanning at term (40 weeks CA). Clinical follow-ups were conducted at two years to measure their motor and speech development. A list of the recordings in the database, along with the clinical outcomes of the infants, is available in Appendix A.

The infants used in this section include infants with no abnormalities during the MRI scanning at term, and demonstrating healthy motor and speech development

at the two year follow-up. The time between EEG recordings and birth varies with the infants. The infant CAs are recorded in the graphs shown in this chapter for comparative purposes.

Figure 6.10 from Chapter 6 shows the progression of the continuity feature as an infant matures. The graph shows that, in general, the standard deviation of the amplitude distribution decreases as the infant grows towards term. This reflects known medical guidelines for EEG interpretation. As preterm infants grow towards term, their very discontinuous signal generally becomes more continuous. This includes an increase in amplitude in the low amplitude periods of the signal, as well as a decrease in amplitude in the high amplitude periods (or the "bursts"). In the continuity feature domain, this corresponds to a decrease in the standard deviation. The mean parameter increases at the continuous end of the spectrum, and decreases at the discontinuous end. Although both continuous and discontinuous signals mature in similar ways as described earlier, because of the different proportions of high and low amplitude segments present, maturation affects the mean parameters differently. Continuous periods have a more even distribution of relatively high and low amplitude segments, and the difference between the amplitude values in these segments is a lot smaller than for discontinuous periods. Discontinuous segments, on the other hand, consist of more low amplitude segments than high amplitude ones. Because the feature used is the distribution of the log of the amplitude, an increase in the amplitude of the lower amplitude segments will result in a bigger difference in the log-normal distribution mean than the difference made by a decrease of the same amount in the higher amplitude segments. The change due to maturation in the amplitude of the high amplitude segments, especially in the burst segments in burst suppression (*tracé discontinu*), is more dramatic than the change occuring in the low amplitude segments [10]. Using the log-normal distribution mean, the bias towards the difference in high amplitude segments is minimised, and a change in amplitude in the interburst period will affect the continuity feature.

In Chapter 6, the idea of using PCA to extract continuity information was examined. The minor component from the PCA transformation appears to coincide with the way maturation affects the continuity feature. To test this theory, the mean value of second PCA coefficients for the continuity feature was calculated for each recording, and plotted against the CA of the infant at the time of recording. Recordings with obvious artifacts or excessively high impedence values were excluded from this graph.

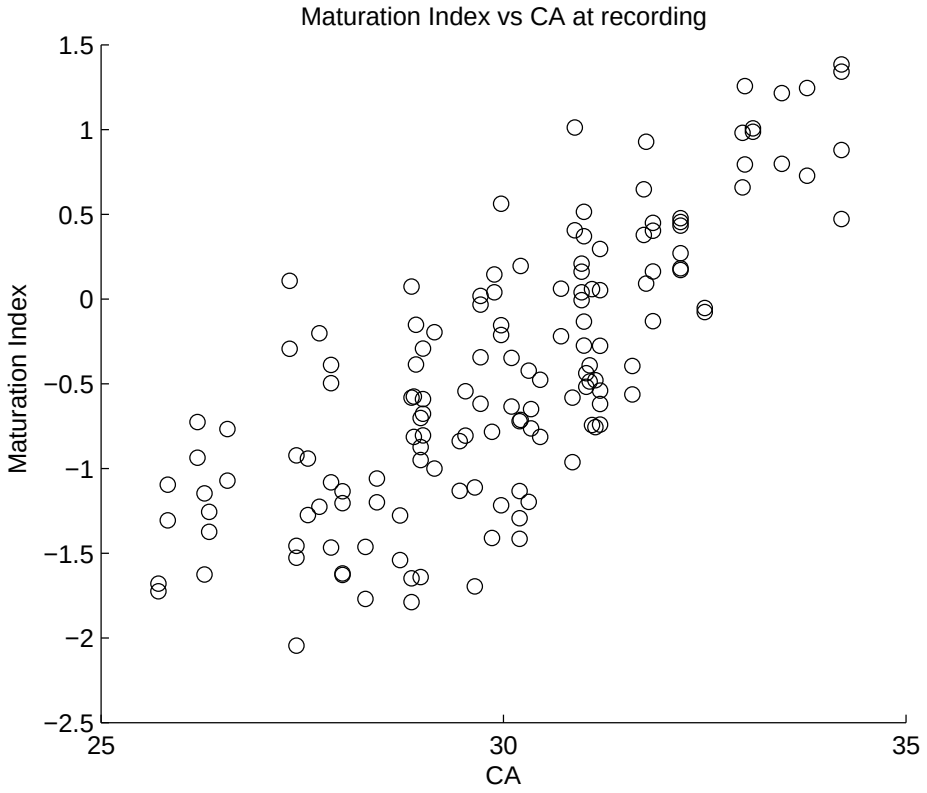

Figure 7.1: Scatter plot of minor component in PCA (maturation index) vs CA of infant at time of recording

From Figure 7.1, the correlation between CA and the mean value is used since the second PCA coefficient is othorgonal to the principal component, which is related to the continuity of the signal, and therefore remains relatively stable during the course of the recording. Further research can be applied to enhancing the maturation index with specific continuity states, to further correlate the maturation index with the age of the infant.

7.3 Effects of Brain Injury on Maturation

Figure 7.1 shows the maturation progress in infants with no known brain injury and a healthy two year clinical outcome. Although this study targets the maturation of healthy infants, it is worth looking at how infants with brain injury differ from infants with a healthy clinical follow-up. Figure 7.2 shows a scatter plot of an infant with grade 3 white matter injury (grade 1 being no injury and grade 4 being most severe).

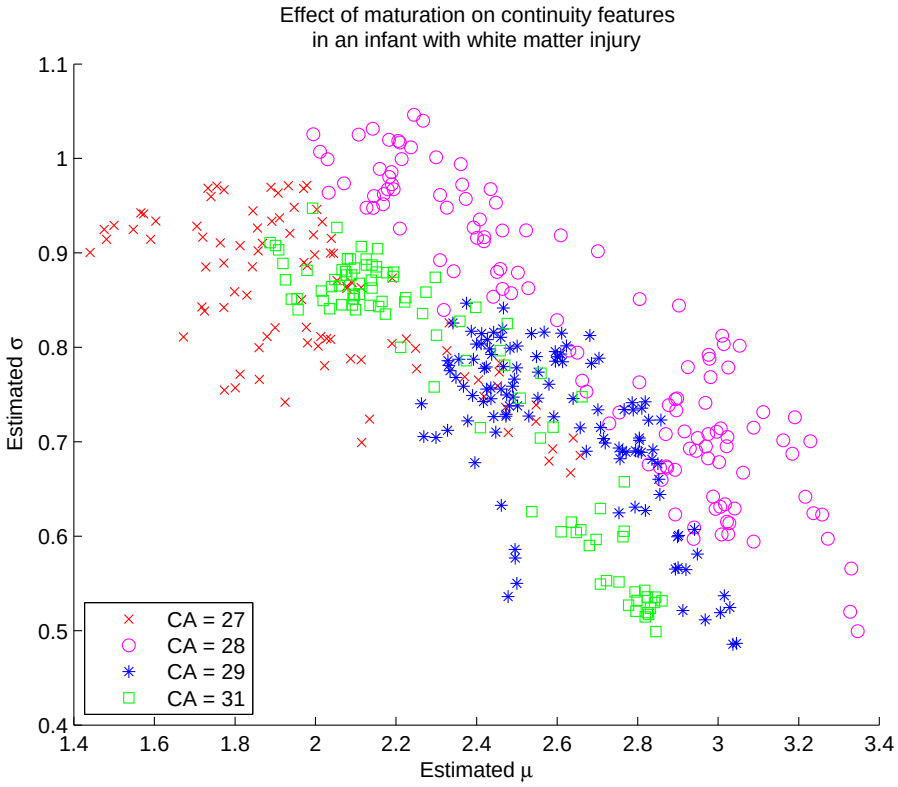

Figure 7.2: The continuity feature of recordings from an infant with grade 3 white matter injury

Compared with Figure 6.10, which depicts the continuity feature from an infant with no brain injury, the features shown in Figure 7.2 lack the progression which is shown in the feature as a decrease in both $\hat{\mu}$ and $\hat{\sigma}$. It is important to note that some infants with brain injury demonstrate a similar pattern to infants with healthy brain scan results, instead of the pattern exhibited in Figure 7.2. The investigation of possible causes of these outlying results would require additional medical investigation, and is outside the scope of this engineering research.

When the mean of the maturation index is plotted for recordings of infants with brain injury, the data points remain within the range of the healthy infants. However, if the progressive recordings of the infants are plotted and joined, the rate of increase shows a small but noticeable difference. Figure 7.3 shows the maturation index averages in sucessive recordings for both infants with a healthy two year follow-up and infants with grade 3 or 4 white matter injury with mental retardation at the two year follow-up. Only recordings with an impedence less than 10 kΩ were used, and only infants with two or more recordings are displayed. Infants with a healthy follow-up are

those who show average or better neurological test results in both motor and language skill tests at the two year follow-up, with no signs of white matter injury in the MRI scan undertaken at term. Infants with white matter injury were chosen for abnormal scores in both motor and language skill tests at the two year follow-up, with white matter injury scores of 3 (major) or 4 (severe) for the MRI performed at term. While the blue lines, representing the infants with a healthy two year follow-up, generally increase at roughly the same gradient, with a few infants being the exceptions, the infants with white matter injury, represented by the red lines, show either a slower rate of increase, or, in some cases, a decreasing maturation index. Figure 7.4 shows a box and whisker plot, displaying the gradients of the lines of best fit in Figure 7.3 for each of the infants. The plot shows that, in general, sick infants have a lower gradient of maturation changes. Some healthy infants also exhibit a low gradient, and further research can help to identify other features that can differentiate the two groups of infants further.

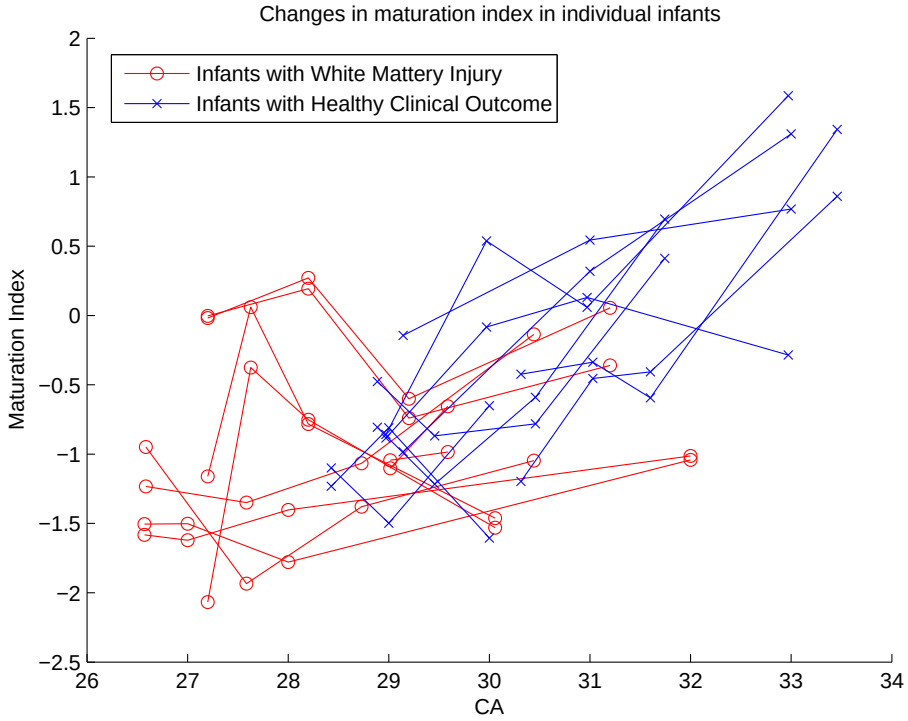

Figure 7.3: Maturation progress for infants with a healthy follow-up and infants with white matter injury. Red lines represent infants with white matter injury and blue lines represent infants with a healthy clinical follow-up. Each line represents the progress of one brain hemisphere of one infant.

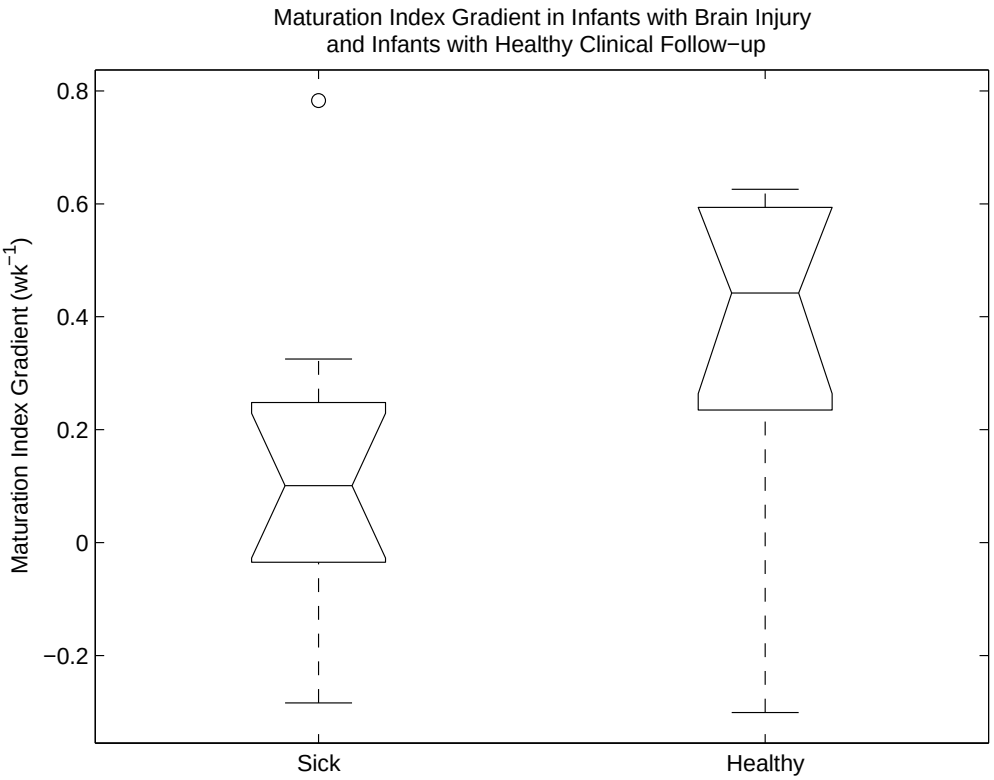

Figure 7.4: Box and whiskers plot of the estimated gradient of the maturation index vs. age of infants. The graph is plotted using the boxplot function in MATLAB, and boxes with notched areas that do not overlap indicate a difference in median value at the 5% significance level.

7.4 Discussion

These differences in the rate of change of the maturation index suggest that the maturation rate of the EEG signal, judged by the the continuity alone, can be used to identify infants with white matter injury. However, at this stage it is still only a rough estimate, and the index can benefit from other information from the EEG, such as symmetry and synchrony (see Chapter 2). At present, the maturation measure is still prone to interference from artifact such as that caused by movement. While a basic noise rejection algorithm, such as rejecting recordings with high impedence, can be used to attempt to exclude EEG recordings containing some form of noise, some files used may still be corrupted by noise such as muscle artifact. Manual artifact detection is very time consuming, and is based on experience. Without a reliable noise rejection algorithm, it is difficult to determine whether an unexpected change

in the maturation index for one infant is caused by a neurological problem or is due to noise in one of the recordings. At this stage, the maturation index can be used as an experimental measurement for the EEG recording. Performed on recordings with very little or no noise, it can help to track the maturation of the infant, and provide another quantitative measurement enabling clinicians to correlate the EEG signal with the clinical outcome of the infant.

7.5 Summary

EEG changes in behaviour as preterm infants grow towards term. In terms of continuity, EEG signals start exhibiting burst suppression patterns for very young preterm infants, and become more continuous as the infants grow. This chapter presented a way to quantify the maturation of EEG according to the quantitative continuity feature. As shown in Chapter 6, the continuity feature changes in the direction corresponding to the minor component of the PCA transformation. The mean of the minor components in the recording is used as the maturation index for the recording. A scatter plot of this maturation index against the conceptional age of healthy infants at the time of recording shows a linear correlation between the age of the infant and the maturation index. Furthermore, a plot of the changes in the maturation index throughout several recordings shows that this index changes at a faster rate in healthy infants than in infants with severe white matter injury. However, this index is still prone to artifact interference, which affects the correlation between the age of the infant and the index. Also, continuity is only one aspect of maturation, and other elements of the EEG signal should be taken into account for a more comprehensive description of maturation. The maturation index at this stage serves as an experimental value that helps track the way continuity changes as infants mature, and the progress of these changes through time.

Chapter 8

Background State Classification

Besides a quantified continuity measurement, the continuity feature described in previous chapters can also be used as input for a classification system resulting in labels familiar to clinicians. Having the qualitative labels can also assist in tasks like sleep-wake cycle detection which require recognition of a series of changing continuities. This chapter will discuss a few different approaches to the task of classification, and how they compare with one another.

8.1 Overview of Continuity Classification

Although a quantified continuity measurement addresses the lack of a continuity index correlating age and clinicial outcomes, there are certain tasks where a qualitative label is preferred. Tasks where the background state is used as the context of the analysis, such as burst detection during burst suppression, require the algorithm to detect the background states of the signal in a qualitative manner. Therefore, a classifier for continuity is still important in EEG analysis.

There are different approaches to this classification problem. The two main categories are unsupervised classification and supervised classification. In unsupervised classification, the learning data are not labelled, and the number of classes is either

determined by an optimisation algorithm or set to a predetermined number. In supervised classification, the learning data are labelled, and the resulting classifier will assign as output the labels present in the learning dataset.

The advantage of the unsupervised classifier is that the resulting labels will not be limited to the labels currently used by clinicians. Since the number of classes can be determined by an optimisation algorithm, the classes determined by a unsupervised classifier may be more suited for the dataset than the conventional labels given by clinicians in terms of reflecting any intrinsic clusters present in the dataset. However, an unsupervised classifier does not guarantee that labels generated reflect the continuity states, should any unrelated pattern affect the feature. Also, if no obvious clusters are present in the dataset, the classes determined by an unsupervised classifier may not be meaningful at all.

Supervised learning requires labelled data, which are not always available. However, if the dataset contains a good representation of the classes desired for the classification system, supervised learning can be an effective method to classify the desired output classes. In the application of continuity classification, since current continuity determination is a subjective exercise on raw EEG or aEEG signal, the labels used for learning data are manually assigned by a visual assessment of the EEG and aEEG recording.

8.2 Evaluation of Classification Systems

The evaluation of the continuity classifiers poses an interesting problem. Since continuity is a qualitative feature, and one that is evaluated with subjective bias, the statistical significance of the classifiers cannot be evaluated in the traditional sense. In this research, the classifiers will be compared with a qualitative classifier previously shown to classify term infant EEG data [38]. Readers are reminded that term and preterm infant data have different EEG continuity distributions, and what is considered discontinuous for a normal term infant may be considered continuous in a preterm infant, as preterm infant data has generally fewer continuous EEG patterns. The aEEG recording is also shown as a visual guide to show examples of classification results. aEEG interpretation is explained in section 6.2. The classification results were also compared with the continuity index described in Chapter 6. Given that this continuity index reflects the degree of continuity in an EEG signal, this can

also provide a good comparative measure of how good a classification system is.

8.3 Linear Discriminant Analysis

Linear Discriminant Analysis (LDA) is a supervised classifier. LDA looks at the data in the different classes to determine a linear mapping that increases the between-class variance and minimises the within-class variance [58]. The within-class variance matrix is defined as:

$$S_w = \sum_{j=1}^{C} p_j \times (cov_j) \tag{8.1}$$

where C is the number of classes, cov_j is the covariance matrix of class j, and p_j is the priori probability of class j. The within-class matrix, S_w, is therefore a $n \times n$ matrix, where n is the number of dimension of the data to be classified. The $n \times n$ between-class variance matrix is defined as:

$$S_b = \sum_{j=1}^{C} (\mu_j - \mu) \times (\mu_j - \mu)^T \tag{8.2}$$

where μ is the global mean and μ_j is mean of the class j. The projection matrix is defined as the eigenvectors of $S_w^{-1} \times S_b$. The transformation is optimised to ensure the ratio $\det|S_b|/\det|S_w|$ is maximised. The transformed data are then used for classification purposes, using the Euclidean distance between the testing point and the centre of the data in each class of the testing dataset.

Both crisp and soft classifications were performed for comparative purposes. Crisp classification refers to classification systems where a class label is assigned to each data point to be classified, so each data point is said to be belong to one and only one class. Soft classification, on the other hand, gives the probabilities of the data point belonging to each of the classes. Soft classification is particularly useful for the grey area between discontinuous and burst suppression data. The implementation of the LDA classifier is from an open source MATLAB toolbox called PRTools, available at [59]. The LDA classifier from the toolbox does not take into account the prior probabilities of the classes. This does not affect the problem at hand, as the prior probabilities of the classes are considered equiprobable. The linear discriminant is estimated by minimising the error in the least square sense.

From a database of preterm EEG recordings, 10-minute segments of EEG were

selected as the training set after examining the aEEG and raw EEG data to ensure the segments were good representations of continuous or discontinuous signals. The continuity feature was extracted from the training set as described in Chapter 6, and an LDA classifier as described previously was trained using this set of data. Figure 8.1 shows the distribution of the feature in the training set.

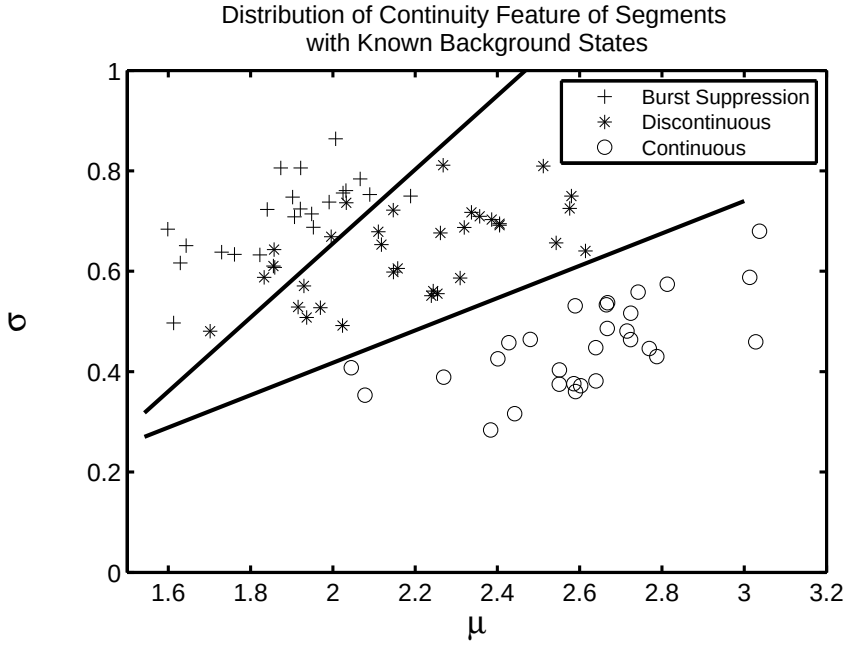

Figure 8.1: Distribution of the $\hat{\mu}$ and $\hat{\sigma}$ in the training set.

From the same database, 60 recordings, approximately 2 hours in length each, were selected as a testing dataset. Selections were based on the quality of the recording, and signals without seizures or significant mechanical artifacts were selected as testing signals. The EEG recordings from the testing dataset were processed using the algorithm described in Chapter 6, and the results were visually compared with the aEEG and an existing algorithm for term infant background detection. The first and last 5 minutes of each 2-hour recording were not classified. No attempt was made to reject artifacts of any nature. Both the crisp and soft classifications were graphed against the background state detected using the algorithm described in [38] and the aEEG signal, the latter being an established way for clinicians to determine the background continuity.

8.3.1 Results and Discussion

Figure 8.2 shows an example of classification results from the proposed method and the rEEG [38]. Note that the rEEG method is designed for full term infants (with a CA of 40 weeks), while the EEG recording used for Figure 8.2 is from an infant with a CA of 34 weeks. The distribution of the feature for the same example is shown in Figure 8.3, plotted with the decision boundaries of the LDA classifier. From the Figure, the aEEG shows the signal is generally continuous, with a period of discontinuous signal between 40 to 60 minutes. The rEEG classification results, as well as the crisp and soft labelling results from the LDA classification system described in Section 8.3, are shown.

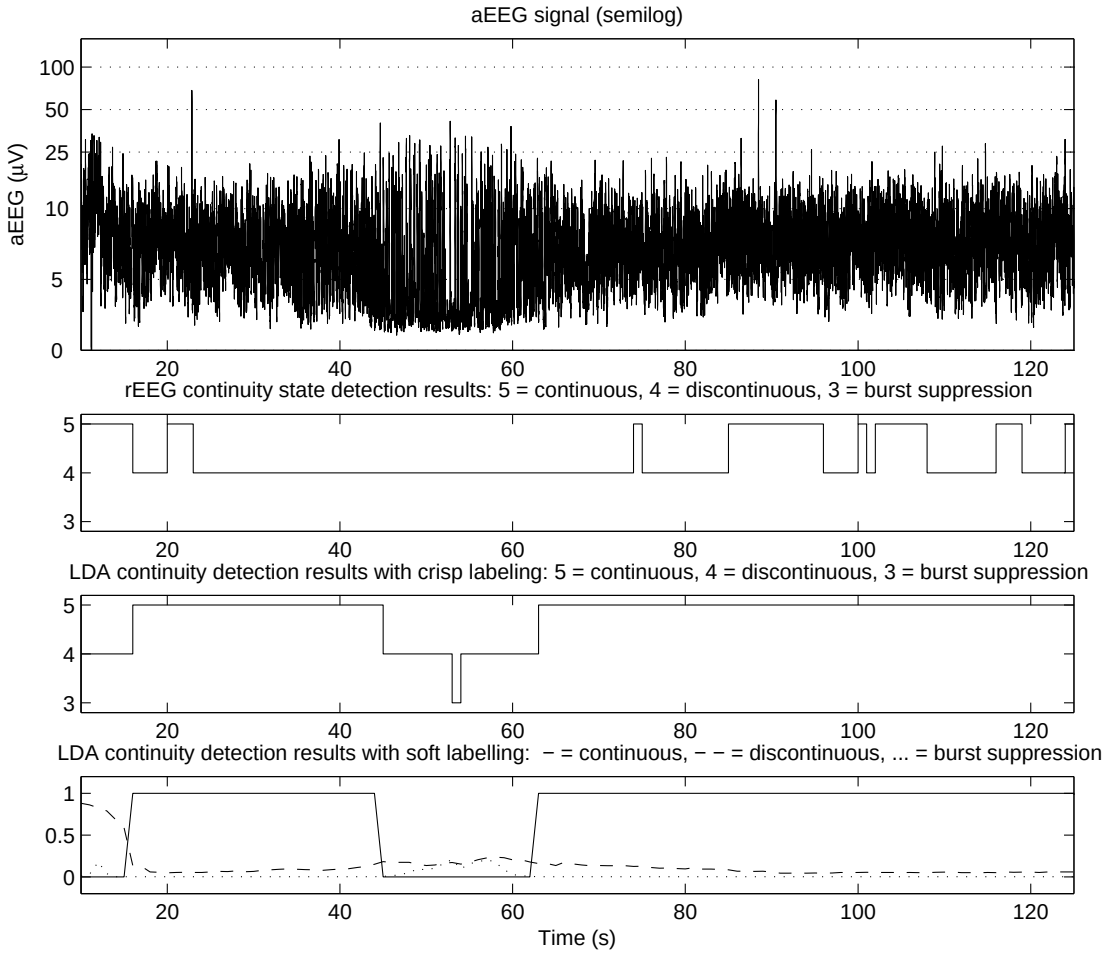

Figure 8.2: Classification results as compared with rEEG based algorithm for an EEG recording of a healthy 34 week CA infant.

In the absence of an existing automatic system to classify preterm continuity, a system designed for term infants was used as well as aEEG as visual references. The 10-minute window used for the feature extraction stage ensures that the features extracted from the window (i.e. $\hat{\mu}$ and $\hat{\sigma}$) take into account a long enough period of the EEG to determine the background state. This means that the resulting classification system provides states that are more stable and less prone to noise interference and less sensitive to local variation of the signal, as shown in Figure 8.2. The fact that the proposed classifier was trained using preterm data also increases the accuracy to make it better suited to preterm infants. From the example shown in Figure 8.2, the term infant algorithm identifies some of the continuous regions of the EEG. However, since the continuity threshold is different from that of a preterm infant, even though the aEEG shows a change in states around 45 to 60 minutes into the recording, the term infant algorithm has identified a larger area as being discontinuous. It is also worth noting that, beside more accurately describing the state changes thoughout the widest band of the aEEG, the proposed system also accurately identified the brief period of burst suppression (characterised by a very low value for the lower edge and a "spikey" appearance in the aEEG) around the 50 minute mark in the recording. The mislabelling of the term infant algorithm is more obvious for younger infants. Figure 8.4 shows the EEG continuity classification results of a recording from an infant with a CA of 31 weeks. Compared with the file from an infant of 34 weeks CA (Figure

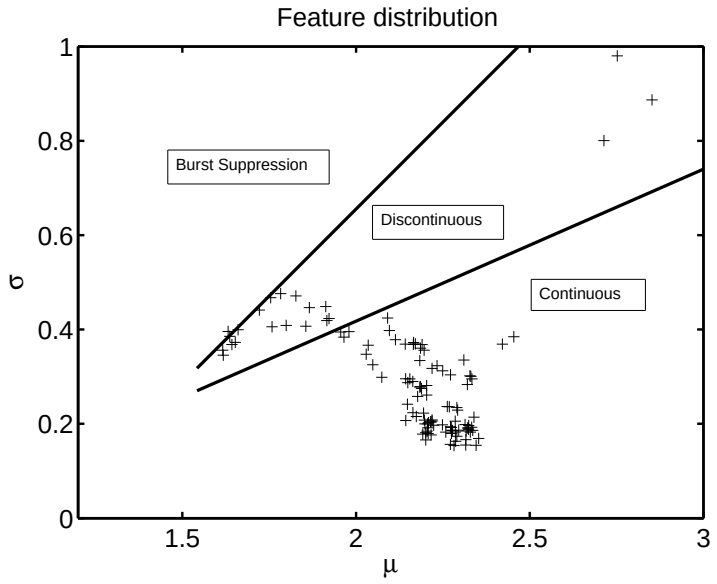

Figure 8.3: Distribution of continuity feature from the example shown in Figure 8.2

8.2), the rEEG algorithm, designed for term infants, identified several short segments as continuous, and others are labelled as discontinuous. The LDA identified a certain period of the signal as discontinuous, which corresponded to the change in behaviour in the aEEG of the signal (a subtle drop in the lower band and a drop in the upper band, with a spikey appearance).

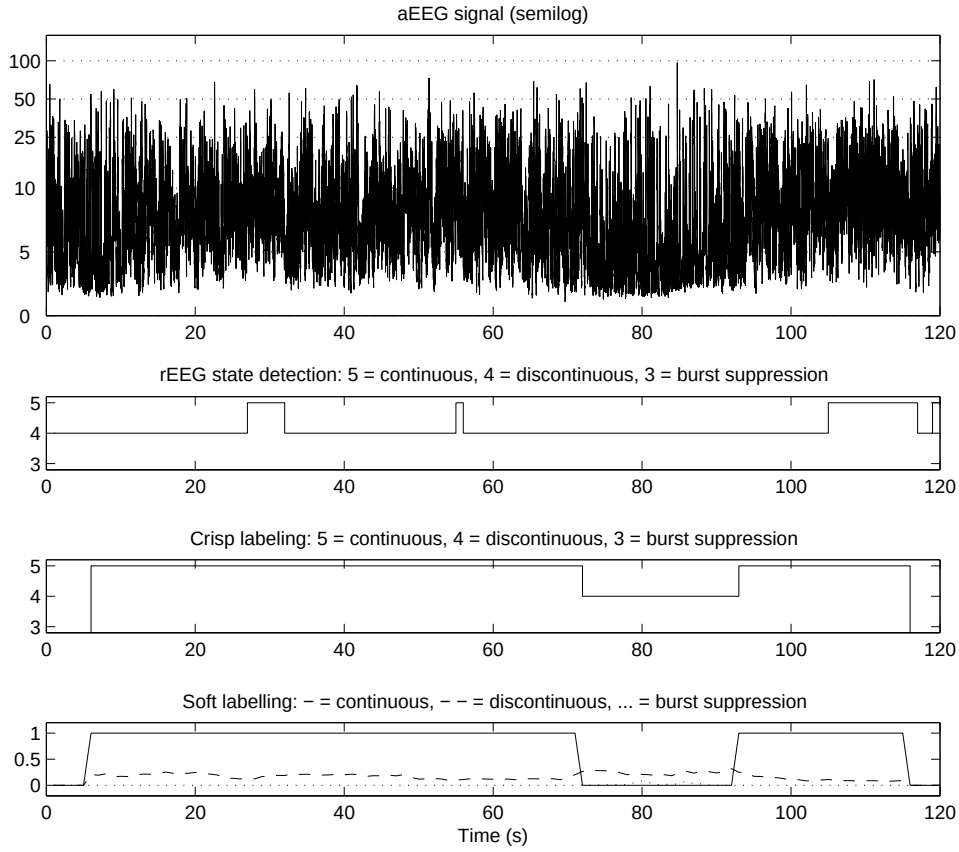

Figure 8.4: Classification results from an EEG recording of a 31 weeks CA infant.

One problem with the background state detection is the fact that the changes between one state and another do not occur instantly but rather, from the aEEG graph, gradually change from one state to another. Using the soft classification, each window is classified with a probability of belonging to the three classes, instead of being assigned a label as in a crisp classifier. The soft label can be easily converted to crisp labelling, while providing more information about the background state of the EEG. It can also assist future work in EEG analysis by defining the area of the signal where no state changes occur. Looking at the training data distribution in

Figure 8.1, data points in the discontinuous and burst suppression classes overlap in the feature space. Compared with the boundary between the continuous and discontinuous classes, where the data points do not overlap and have an obvious gap between the two classes, the boundary between discontinuous and burst suppression presents a larger degree of uncertainty. A better defined training set can help reduce the uncertainty between the two classes.

8.4 Self Organising Map

Self Organising Map (SOM) is an artificial neural network normally used for unsupervised clustering. A two-dimensional network of neurons is used to represent a mapping of the features used for clustering. Each neuron is represented as a feature vector with the same length as each of the feature vectors used for training and clustering, known as the "weight" of the neuron. In the training phase, the neuron most similar to the training feature vectors (defined as the neuron with the least Euclidean distance from the training vector) is identified, and known as the best matched unit (BMU). Its neighbouring neurons are subsequently updated. The updating function is:

$$W_i(t+1) = W_i(t) + \Theta(i, v, t)\alpha(t)(D(t) - W_i(t)) \tag{8.3}$$

Where $W_i(t)$ is the weight of the neuron i, $\Theta(i, v, t)$ is the neighbourhood function given v, the BMU for the input, $D(t)$. The neighbourhood function can be as simple as 1 for immediate neighbours and 0 otherwise, or graduated in various degrees depending on the distance on the SOM between the neuron i and the BMU v. The learning parameter, $\alpha(t)$, is usually time dependent so the initial learning phase will alter the weights more dramatically than the later fine-tuning phase. The result from this training phase is a mapping representing the training data, where neurons with similar weights are close together and neurons with dissimilar weights are further away from each other. Any cluster boundary intrinsic to the dataset will also appear as a large Euclidean distance between neighbouring neurons.

Figure 8.5 shows a graphic example of an SOM mapping along with the clustering results. Figure 8.5(a) shows the U-matrix of the resultant SOM. The U-matrix is a way to show the Euclidean distances for the normalised feature vector between

neighbouring neurons. For each neuron in the SOM, the Euclidean distances between its weight and each of its immediate neighbours' weights are calculated. These differences are shown in the U-matrix by an additional cell (not represented in the SOM map) between the two cells. A high value in the U-matrix represents a possible cluster boundary in the dataset. Figures 8.5(b) and 8.5(c) show the values of $\hat{\mu}$ and $\hat{\sigma}$, respectively, for each neuron, represented as hexagons, in the SOM mapping. Figure 8.5(d) shows the k-mean clustering results, such that each colour represent one of the two classes. Comparing Figure 8.5(a) and 8.5(d), the boundary between the two clusters corresponds to high values in the U-matrix. Once clustering is performed on the SOM, data can be classified by finding the BMU, and identifying the cluster the BMU belongs to.

The K-means clustering described above uses the weights of the neurons as the dataset. The cluster centres were first initialised with random value. Each neurons is associated with the cluster with the closet cluster centre. Each cluster be assigned a new centre equals to the mean of its members. The dataset is re-clustered using the new centres, and new centres are found using the new clusters. This process continues until the classification label of the data does not change for two iterations.

Using unsupervised clustering in this application has an advantage, because of the loose definition of the different continuity labels and the lack of a definitive continuity classification protocol. An unsupervised clustering method can pick up any intrinsic clusters in the dataset. A heuristic formula can be used to determine the optimal number of clusters in the dataset. The disadvantage of an unsupervised method is that the clusters found may not reflect the classification desired. Also, if no obvious clusters are present in the dataset, the clusters determined by the algorithm may have no physical meaning at all. Because the feature used is related to continuity, as shown in Chapter 6, finding clusters not reflecting the desired class definitions is not likely be become an issue. However, should the dataset include continuity information of areas of the signal that are heavily affected by noise (e.g. movement artifact), clusters with continuity information may not be as distinguishable as clusters showing areas with and without noise. This may be useful in a noise detection application, but since the primary concern in this project is to investigate continuity, areas heavily affected by artifact will be excluded.

There are different approaches to using SOM for the task of continuity classification. The continuity feature of a pool of segments with different continuity back-

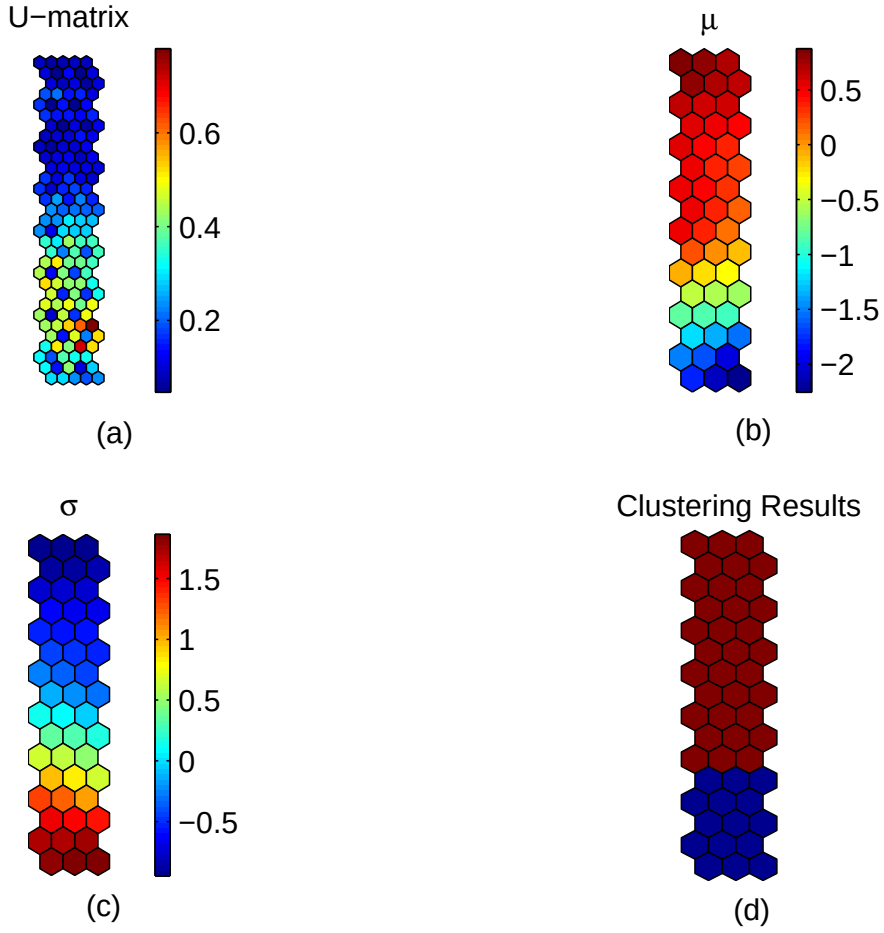

Figure 8.5: Example of a 16-by-3 SOM mapping. (a) U-matrix showing the Euclidean distances between each neuron. (b) and (c) Resultant mapping for the dimensions $\hat{\mu}$ and $\hat{\sigma}$, respectively. (d) The clustering results from k-mean clustering.

ground states, as determined by human experts, can be used as training data to come up with a definitive mapping for all EEG recordings. This approach could produce a classifier that will assign labels that are related to known continuity states, and the classification would be consistent across all recordings. In contrast, each recording can be used as a training set to come up with an SOM unique to the recording. This approach would focus on identifying hidden patterns, and divide the segments into states that are native to the recording. Since each file seems have to stable states that are different from each other, this approach can help identify stable states unique to each file.

The SOM implementation used here is an open-source MATLAB toolbox called SOM Toolbox [60]. A built in function allows k-mean clustering to be performed on

a trained SOM. Since the trained SOM is a mapping of the training set where the weight vectors of the neurons are prototypes of the training data, a k-mean clustering performed on the trained map would be similar to using k-means clustering on the training set itself. The number of clusters is set at two since the EEG feature patterns show that, in 2-hour recordings with more than one state, the signal appears to oscillate between two states.

8.4.1 Results and Discussion

SOM provides a way to organise and display a set of multidimensional data. The advantage of using SOM over other clustering methods is the fact that a mapping is produced on completion to represent the training data, and it depicts the multidimensional data as a two-dimensional map that can be easily visualised to show the relationships between the different feature dimensions. In this case, however, since there are only two dimensions in the data to be classified, the advantage of SOM is not very obvious, since two-dimensional data can be easily plotted in a scatter plot to show any correlation. Figure 8.6 shows an example of the classification resulting from the SOM.

The results are very similar to those of the LDA based classifier, which can be shown by Figure 8.7, where the LDA boundaries are depicted with the feature in the same plot. The SOM classification results cluster the data points into two groups, where one group corresponds to the continuous state of the EEG signal, and another group corresponds to the discontinuous and burst suppression states. One drawback of claffication using SOM is the fact that no soft labels can be assigned, since the features are classified by comparison with the mapping, which consists of prototypes with labels assigned using k-mean clustering. Apart from the lack of soft labelling, the performance of SOM is good: the algorithm correctly located the change of states and marked the portion of the signal that was discontinuous in Figure 8.6. The centres of k-mean clustering were examined to determine the state each of the clusters represented, on the assumption that one cluster represents discontinuous activities and one represents continuous activities. Because only two clusters are found, if the recording include both discontinuous and burst suppression EEG, as well as continuous EEG, the algorithm will only produce two labels, and therefore one will not be able to differentiate between discontinuous and burst suppression EEG. In most cases, EEG signals that contains more than one background continuity states can be

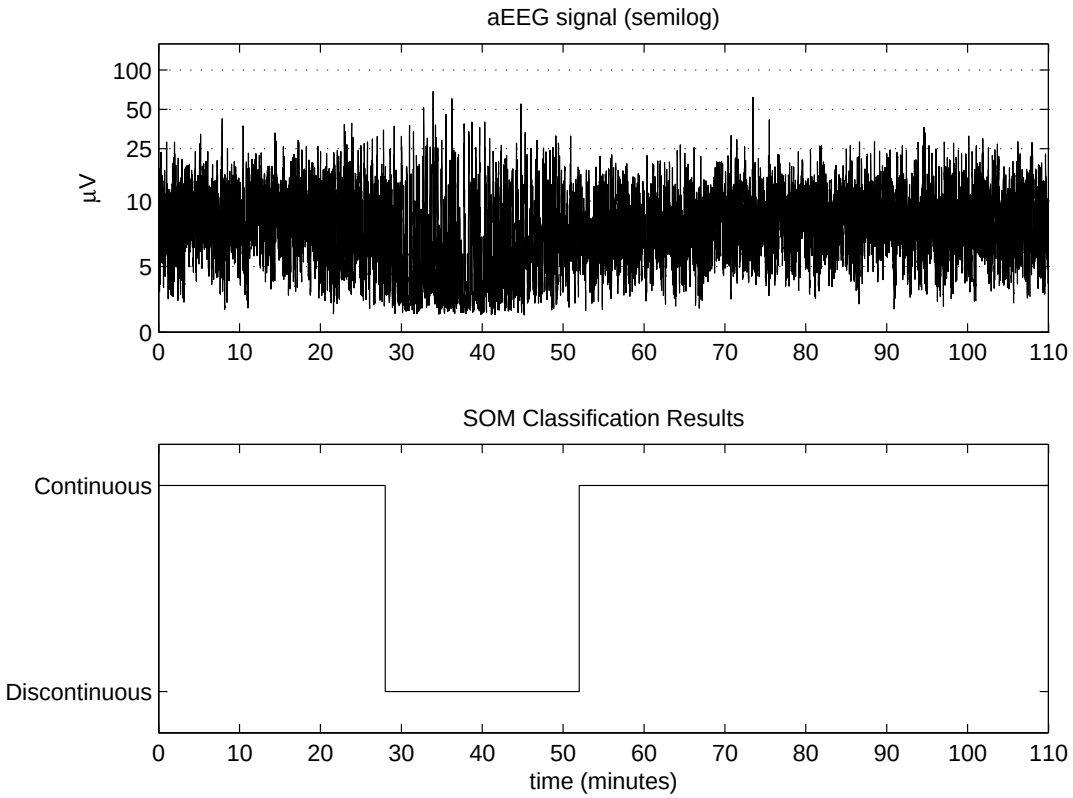

Figure 8.6: Classification results of SOM classification compared with aEEG

described as having one continuous state (e.g. awake) and one discontinuous state (e.g. asleep), which can be either discontinuous or burst suppression, depending on the age of the infant. This means that a two state classifier should be sufficient to differentiate the two main states. Further work on the optimal number of clusters can be further investigated.

8.5 Gaussian Mixture Model

Gaussian Mixture Model (GMM) is a way to model a set of unlabelled data. GMM is based on the assumption that each data point is a member of a number of Gaussian distributions. In this application, each Gaussian distribution represents a state. Using the continuity feature of the recording, the most likely Gaussian distributions that represent the underlying state are calculated. The algorithm used is as implemented in the open source MATLAB toolbox called PRTools, available at [59]. The algorithm

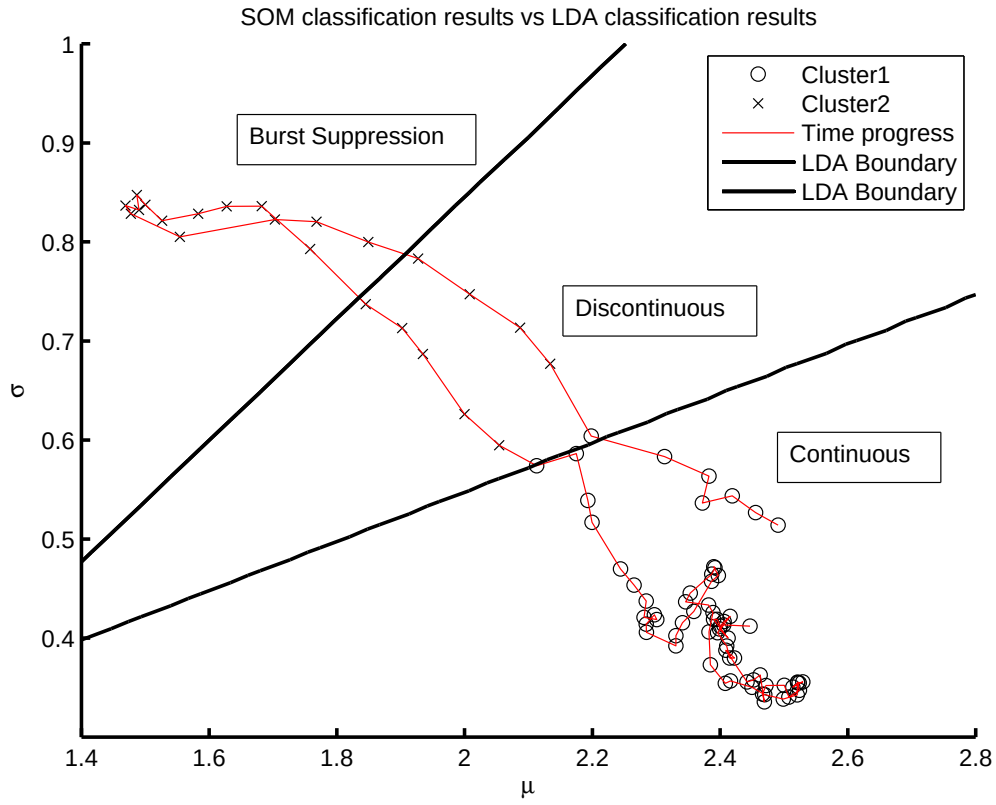

Figure 8.7: Continuity feature clusters in the signal shown in Figure 8.6 with LDA decision boundaries

uses the K-centres algorithm to cluster the the data. Using the K-centres algorithm, a predefined number of centres are determined, such that the maximum distance between a centre and any data point is minimised. The K-centres algorithm is similar to the K-means algorithm described in section 8.4, but instead of using the mean of the cluster member as the new centres in the next interation, K-centres uses the cluster member with the minimal distance from other members as the new centre for the clusters. These centres will be used as the centres of the Gaussian distributions. Each data point is assigned to the closest cluster centre. The covariance matrices of the clusters are calculated using the data in the cluster. Each feature vector (the $\hat{\mu}$ and $\hat{\sigma}$ pair) is analysed and the probablity of the data point belonging to each cluster is calculated according to the mean and covariance matrix of the clusters. The probability of a feature vector $\mathbf{x}$ belonging to a cluster with centre at μ and the covariance matrix $\mathbf{\Sigma}$ is defined as:

$$\mathbf{P}(\mathbf{x}) = \frac{\mathbf{1}}{(\mathbf{2}\pi)^{(\mathbf{N/2})}\sqrt{|\mathbf{\Sigma}|}} \exp\left(-\frac{\mathbf{1}}{\mathbf{2}}(\mathbf{x}-\mu)^{\mathbf{T}}\mathbf{\Sigma}^{\mathbf{-1}}(\mathbf{x}-\mu)\right) \tag{8.4}$$

where $|\Sigma|$ is the determinant of the covariance matrix, Σ.

Gaussian mixture model assumes each data point is a member of a number of Gaussian distributions with different probabilities of membership. In this instance, the probability of the feature vector belonging to each of the Gaussian distributed clusters is normalised so the sum of the probabilities is 1. It is assumed that no major movement artifact is present in the recording, and two clusters are present in each recording. The reason only two clusters are used rather than three is because, judging from the scatter plots of the feature in recordings which consist of more than one state, the EEG appears to change between two states: one towards the continuous end and one at the burst suppression end of the feature distribution. The centres of the clusters are used to determine the state of each cluster: with the highest $\hat{\mu}$ being the continuous state, and the lowest value representing the discontinuous or burst suppression state.

8.5.1 Results and Discussion

Figure 8.8 shows an example of the classification results compared with aEEG. The classification results are similar to those resulting from the LDA and SOM. The main difference is in the way the class labels are determined: in the LDA method, labelled data were used to train the classifier and the mapping was identical for all recordings, where in the GMM method, the labels were assigned according to the clusters found in the data itself. Figure 8.9 shows the feature extracted from the signal used in Figure 8.8, along with the decision boundaries for the LDA classifier for comparative purposes. As shown in the plot, the LDA decision boundaries divide the plot without regard for the distribution of the feature within the signal. In contrast, the GMM method located the two intrinsic clusters: one for the continuous state and another state at the discontinuous end of the spectrum, which reflects the two intrinsic states of the EEG recording more accurately.

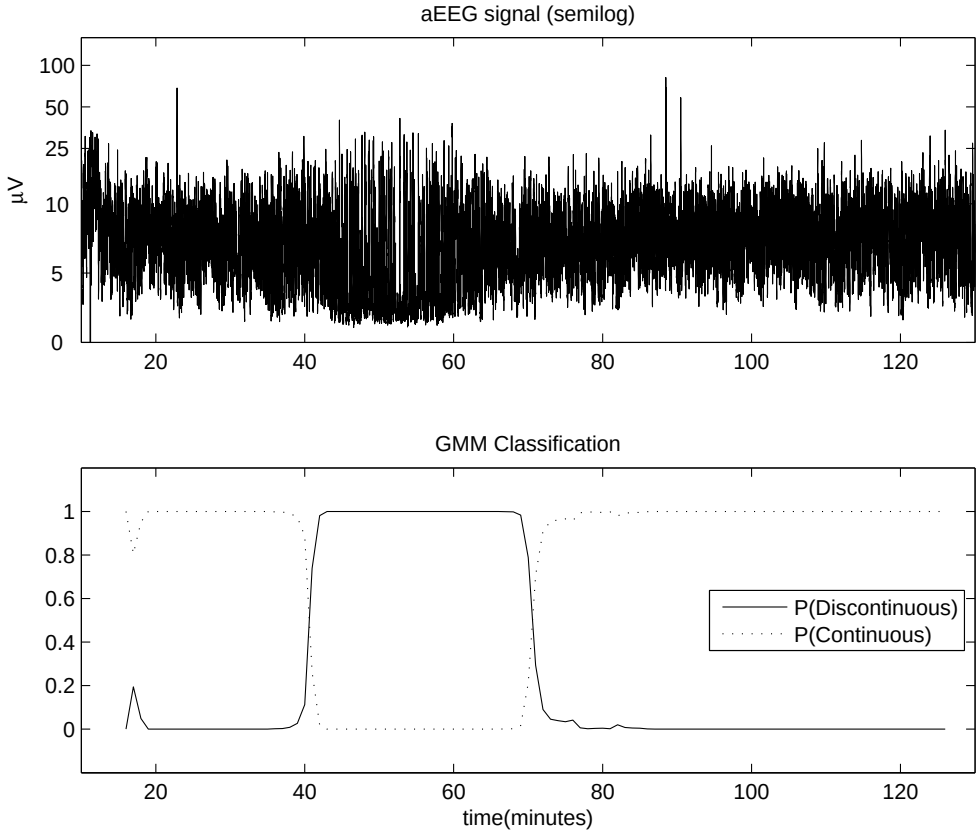

Figure 8.8: Classification results of GMM classification compared with aEEG

8.6 Comparison of Algorithm Performance

Both supervised and unsupervised classification systems were used to compare these two common types of classification method. Figure 8.10 shows the classification results from the three classification methods. As shown in the graph, the performance of the three algorithms is very similar. The LDA algorithm uses three states, where the two unsupervised methods (SOM and GMM) tried to identify two clusters in the recording: one for the continuous state and another for the discontinuous state. The discontinuous state regions identified by the unsupervised methods are represented in the LDA classifier as "discontinuous" or "burst suppression", depending on the degree of continuity of the signal. The continuity index of the signal (the principal component of the continuity feature, as stated in Chapter 6) is also plotted in the Figure, for comparative purposes. Table 8.1 shows the continuity index statistics of the different states for the three algorithms used in 10 EEG recordings of healthy infants.

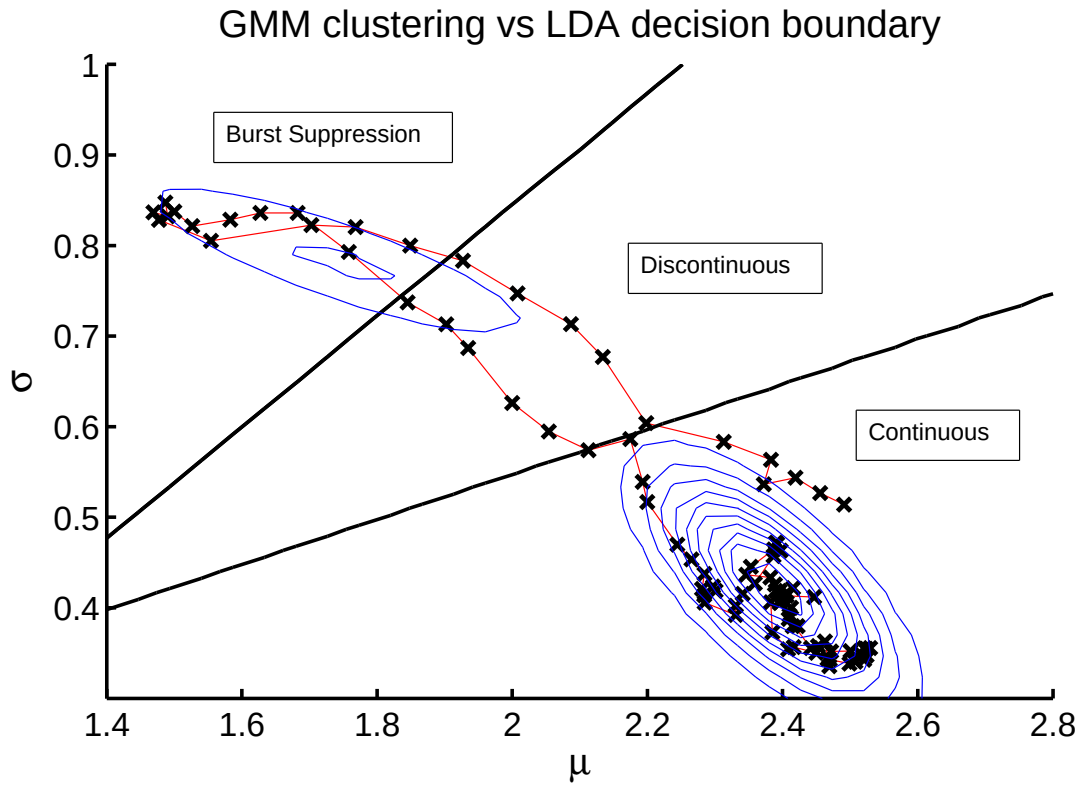

Figure 8.9: Continuity feature clusters in the signal shown in Figure 8.8 with LDA decision boundaries displayed for comparison purposes

Given that the "burst suppression" and "discontinuous" labels used in the supervised algorithm correspond to the "discontinuous" label in the unsupervised algorithms, the maximum continuity values of the "discontinuous" label and the minimum continuity values of the "continuous" label can be compared across the three algorithms. These values give an indication of where the decision boundary is for the two main continuity states of the EEG signal. It is interesting to note that, for the supervised algorithm (i.e. LDA), the maximum continuity value for the "discontinuous" label ranged from -0.52 to +0.10, while the unsupervised algorithms showed a larger range (-1.23 to +0.26 for SOM and -1.23 to +0.20 for LDA). Since the LDA algorithm uses the same mapping for all recordings, the classification is more consistent. With the unsupervised methods, the classification takes into account the distribution of the continuity feature within the recording, and therefore is different from recording to recording.

Although classifications can be obtained from EEG by using human experts to label data, as mentioned by Navakatikyan et al in [38], the subjectiveness of the

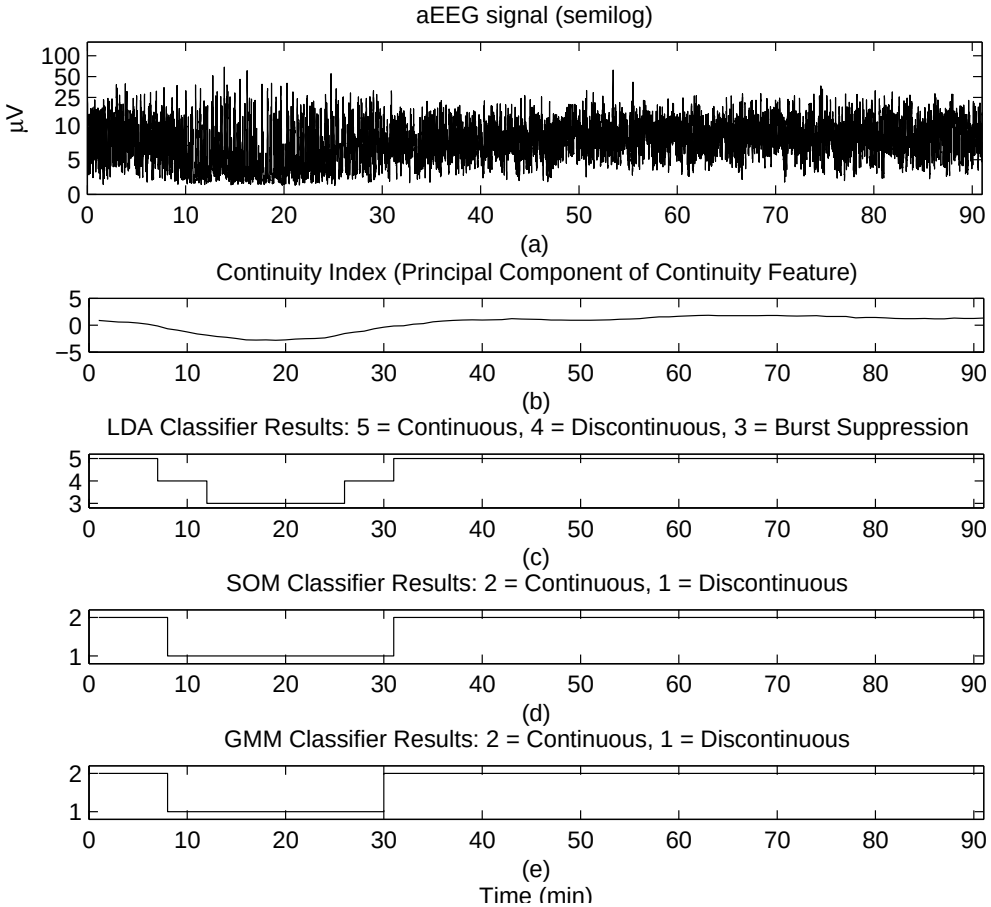

Figure 8.10: Comparison of the three classification methods and the continuity index

current continuity guidelines makes it difficult for human observers to agree on labels at times. The supervised method, namely the LDA based classifier, performs well in the task of determining the continuity of the signal based solely on the value of the continuity feature. No consideration was given to the distribution of the feature, or to any intrinsic clustering that may have been present in the signal. However, it does have the advantage of being consistent across all recordings, and, since the training data was collected across infants of different ages, the classifier can be used in a range of infants with different ages to assign objective labels according to the values of the continuity feature. This method is ideal for situations such as in a bedside monitoring system display, as a guide to aid clinicians. For example, one can determine the sleep state of the infant by looking at the background continuity of the EEG.

The two unsupervised classifiers, SOM and GMM based classification, behave similarily. One of the most important differences between the two classifiers is the

CA (weeks)	LDA results								
	Burst Sup.			Discont.			Cont.		
	max.	mean	min.	max.	mean	min.	max.	mean	min.
30.3	-1.70	-1.99	-2.22	0.10	-1.01	-1.80	0.53	0.34	0.21
31.2	-1.44	-2.10	-2.92	-0.10	-0.80	-1.44	0.98	0.47	-0.05
31.2	-1.71	-2.02	-2.42	-0.17	-1.11	-1.63	1.58	0.93	0.06
31.2	-1.77	-2.64	-3.25	0.03	-0.48	-1.81	1.82	0.89	-0.01
32.2	-1.61	-1.96	-2.20	-0.24	-0.95	-1.53	1.49	0.84	-0.16
32.2	N/A	N/A	N/A	-0.21	-0.65	-1.26	1.55	0.73	-0.19
32.2	N/A	N/A	N/A	-0.15	-0.99	-1.40	1.91	0.93	-0.21
33.8	-1.32	-1.65	-2.05	-0.52	-1.03	-1.32	1.57	0.85	-0.34
34.2	-1.62	-2.79	-3.57	-0.21	-0.79	-1.23	1.44	0.89	-0.20
34.2	-1.38	-2.72	-3.47	-0.34	-0.67	-1.23	1.17	0.55	-0.22

CA (weeks)	SOM results					
	Discont.			Cont.		
	max.	mean	min.	max.	mean	min.
30.3	-0.67	-1.51	-2.22	0.53	-0.36	-1.02
31.2	-1.00	-1.94	-2.92	0.98	0.09	-0.94
31.2	-0.65	-1.57	-2.42	1.58	0.89	-0.17
31.2	-0.88	-2.33	-3.25	1.82	0.44	-0.72
32.2	-0.53	-1.56	-2.20	1.49	0.81	-0.24
32.2	0.26	-0.36	-1.26	1.55	0.86	0.24
32.2	0.01	-0.86	-1.40	1.91	1.00	0.13
33.8	-0.08	-1.17	-2.05	1.57	1.00	-0.02
34.2	-1.00	-2.50	-3.57	1.44	0.77	-0.80
34.2	-1.23	-2.65	-3.47	1.17	0.36	-1.05

CA (weeks)	GMM results					
	Discont.			Cont.		
	max.	mean	min.	max.	mean	min.
30.3	-0.72	-1.46	-2.22	0.53	-0.24	-0.93
31.2	-1.23	-2.02	-2.92	0.98	-0.21	-1.74
31.2	-0.65	-1.57	-2.42	1.58	0.89	-0.17
31.2	-1.36	-2.45	-3.25	1.82	0.40	-1.04
32.2	-0.24	-1.51	-2.20	1.49	0.84	-0.16
32.2	0.20	-0.45	-1.26	1.55	0.80	-0.04
32.2	-0.15	-0.96	-1.40	1.91	0.95	-0.08
33.8	0.02	-1.18	-2.05	1.57	0.98	-0.17
34.2	-1.00	-2.50	-3.57	1.44	0.77	-0.80
34.2	-0.96	-2.57	-3.47	1.17	0.37	-1.05

Table 8.1: Statistics of continuity index in continuity states determined by different classification methods.

clustering method. SOM uses a mapping of the feature, and it is the mapping prototypes that are clustered, rather than clustering the data point directly. SOM can be very helpful in understanding multi-dimensional datasets, since it arranges the multi-dimensional data into a 2-dimensional map, which can be displayed and analysed further to find correlations between different dimensions. However, since the continuity feature has only two dimensions, the advantage of SOM over conventional clustering methods is not very significant. The advantage of an unsupervised classifier is the fact that the clustering method takes into account the distribution of the feature as well as the values of the feature. Any clusters that may be present in the dataset are taken into account by the clustering method. This makes unsupervised methods ideal for further experimental work in understanding the underlying states of the EEG signal, as well as analysing the states of the recording to determine how the EEG states change during the maturation process. Between the two methods examined here, the GMM is recommended over the SOM method, as the results are very similar, but GMM performs the clustering using the continuity feature directly, compared to SOM where the prototypes, representing a number of data points, are used for classification. Therefore, GMM is less computationally expensive, while giving similar performace. GMM also provides soft labelling, which is not available in the SOM classifier. This makes GMM more versatile than the SOM method.

8.7 Summary

The continuity feature described in Chapter 6 can be used as the input to classifiers which assign continuity labels to the signals. Both supervised and unsupervised classifiers were tested in this study. An LDA based supervised classifier was trained with labelled data, and provided better labelling than the previous system based on rEEG. An SOM based classifier and a GMM based classifier were developed to test the two unsupervised methods. Both yielded good results and correctly identified the change of states. The GMM classifier has an advantage over the SOM classifier since it also provides an option of soft labelling that is unavailable in the SOM classifier. It is recommended that both supervised and unsupervised classifiers with the continuity feature described in Chapter 6 be used to determine the states of the EEG signal. Supervised classifiers are ideal for applications where consistent classification is desired, as the decision boundary is identical for every recording. Unsupervised classifiers are preferred for experimental work and to aid further data analysis, since the classifier

attempts to analyse the intrinsic clusters occurring in the recording. This can develop understanding of the state transition by defining the continuity states of the signal in terms of the distribution of continuity feature, rather than on the value of the feature alone. Of the two methods studied in this work, the GMM based classifier is recommended over the SOM based classifier for the option of soft labels, being less computationally expensive, while providing similar performance.

Chapter 9

Discussion

Quantifying the EEG is advantageous, providing a measurement for further correlation with both maturation and clinical outcomes. The quantified continuity measurement can also be used to help classification systems to provide continuity labels familiar to clinicians. This chapter discusses the advantages of the system, both the quantified continuity measurement and qualitative labels as assigned by the classifiers. It also outlines areas where improvement is needed. Other possible applications that the continuity features can be used for are also discussed in this chapter.

9.1 Quantitative Continuity Feature Versus Existing Continuity Measurements

As stated in the literature review, quantitative continuity related features have been developed in medical research to correlate EEG with maturation and clinical outcomes. However, these features often only provide information on a single aspect of continuity, and only serve as a global measurement. For very young preterm infants, a global measurement would not cause problems, since the sleep-wake cycle and state

changes have not yet developed in their EEG signals. However, starting at around 33 weeks CA, differences in sleep state start emerging in the EEG signals. Using a global measurement will either average out the continuity in these states (as in the case of intensity values) or only reflect the continuity of one single state (as in the case of maximum inter-burst intervals). By using a quantified feature that is related to continuity itself rather than an aspect of continuity, the degree of continuity can be quantified. The feature is also localised, and reflects the continuity of a relatively short time window (10 minutes) instead of the continuity of the whole recording (as for the global continuity measurement currently used in medical research). This means that the changes between states can also be tracked within a recording, eliminating the issue of global measurement for recordings that span several states with various degrees of continuity.

One of the biggest issues with a quantitative continuity measurement is that there is currently no continuity index that can be compared with this measurment. Other continuity related features and labels from manual EEG or aEEG inspection could conceivably be compared with this measurement, but would require adjustments to allow for the difference between global (conventional) and localised (proposed) features. The existing state detection algorithm that automatically detects background states was designed for term infants, and is not as accurate on preterm infants, whose EEG continuity behaviour is different from that of term infants. Visual comparison with aEEG appears to be a good way to verify the continuity feature, as there are standard guidelines for interpreting aEEG, which are more appropriate for long recordings where interpreting the raw EEG may be too time consuming. Although guidelines for aEEG interpretion are somewhat quantitative, it is still largely subjective, as the "lower band" and "upper band" of the aEEG (the lower edge and the upper edge of the aEEG signal as displayed on the aEEG plot) have to be estimated from the plot, and are usually visually estimated as a smoothed envelope of the aEEG signal. The aEEG gives a rough idea of the continuity, which is related to the "bandwidth" (difference between the lower and upper bands) and the position of the lower and upper bands. From the results demonstrated in Chapter 6, the aEEG bandwidth is analogous to the estimated standard deviation of the log amplitude vector ($\hat{\sigma}$). This is expected as the bandwidth of the aEEG is related to the degree of variation present in the signal envelope. The lower and upper band are analogous to the values $\hat{\mu} \pm \hat{\sigma}$. Since the bands represent the lowest and the highest points in the signal envelope, they are related to $\hat{\mu} \pm \hat{\sigma}$, which are the values one standard deviation away from the

mean of the log amplitude vector. As the aEEG is plotted on a "semilog" scale (linear from 0 - 10μV, log scale for values higher than 10μV), the use of the log amplitude vector is justified. Using the standard deviation rather than the extreme values also minimises the effect of outliers. However, for muscle artifact, the feature is affected as it would be on aEEG, particularly for burst suppression and discontinuous signals, where the regions with low amplitude activities have a high frequency EMG signal superimposed on them. Since the continuity feature is derived from the logarithm of the amplitude vector, a small addition to the envelope will affect areas with lower amplitude more than areas with higher amplitude. This will cause the estimated standard deviation to be lower than would have been derived from signals without muscle artifact. Therefore, the feature extracted from signals with muscle artifact may not accurately reflect the continuity state of the underlying EEG signal, and will appear more continuous than the underlying EEG signal without the EMG artifact, since the distributions of the log amplitude vector derived from continuous signals have lower standard deviations and higher means. Since noise rejection or filtering is beyond the scope of this project, presently signals used for verification are manually inspected with the impedence values examined to ensure that signals are free from muscle artifact. For a practical implementation of a monitoring system that utilises the continuity feature, consideration is needed to ensure that the continuity feature is not affected by muscle artifact.

9.2 Quantified Continuity Measure

In order for maturation and clinical outcomes to be correlated with continuity, continuity needs to be clearly defined into some form of quantified value. As discussed in the literature review in Chapter 3, clinicians have been using simple measurements such as interburst interval and burst amplitude to correlate continuity information with either maturation or adverse clinical outcomes. This approach only offers a global view of continuity in the recording, rather than using localised information, and does not measure changes in continuity within the signal. Using a quantified version of continuity measurement, the amount by which the continuity changes within the recording can also be compared, as well as the continuity of the signal as a global value. Using a quantified continuity measure also aids in examining the continuity changes as a time series and analysing the changes of continuity in both a short and long term scale. This also removes the subjectivity involved in determining the

continuity state for infants via either the raw EEG signal or the aEEG representation.

Although there are advantages to using a quantified continuity index, it cannot completely replace the qualitative labels. While the quantitative index can be used for correlation with maturation and clinical outcomes, having a qualitative description of the continuity can be useful and essential to applications such as sleep cycle detection. A set of threshold rules can be used to translate the quantitative feature into qualitative labels, however a threshold approach does not take into account any intrinsic clusters present in the feature dataset. Since the states may "stabilise" at different places in the feature space, using a threshold that will apply to all recordings may not reflect the clustering of the data, and may label the data differently.

9.3 Monitoring Display

One of the aims of having an automatic state detection system is to assist clinical staff who are unfamiliar with neonatal EEG signals to gather information from the EEG states. Current aEEG is displayed alongside the raw EEG traces in bedside monitoring systems, as it provides an overview of the EEG signal, as well as offering a way to use simple guidelines to determine various EEG background states. As shown in Chapter 6 and discussed in Section 9.1, plotting the features $\hat{\mu}$ and $\hat{\mu} \pm \hat{\sigma}$ provides a plot analogous to the aEEG plot, by providing a smoother line than the aEEG lower and upper bands. Since the standard deviation of the log amplitude vector is used instead of the maximum and minimum values (as is the case with lower and upper bands of aEEG), the effect outlier values present in the raw EEG have on the display is reduced. Because the line plot is analogous to the aEEG display, existing guidelines can be easily adapted to suit the new plot, without requiring the clinical staff to learn a set of new guidelines for plot interpretation. This eliminates the need to estimate the values of the lower and upper bands required to interpret the aEEG plot.

Although the new feature plot behaves in a similar manner to the well established aEEG and gives smoother edges, the guidelines will need to be adjusted to match the aEEG plot. The current threshold used will not translate directly to the new feature plot, since the thresholds in the aEEG plot refer to the actual voltage rather than the log voltage. Further study will be required to determine appropriate thresholds to be used for the proposed plot. However, since the thresholding method does not

take into account the intrinsic states that appear in the EEG signal, it is possibly best to use the feature plots as a visual overview for the EEG recording, while using a clustering method to determine the background states of the EEG signal.

9.4 Qualitative Labels

Clinicians currently classify continuity states into categories such as continuous and *tracé discontinu*, and an expanse of medical literature uses these conventional labels (e.g. [10]). Therefore, the qualitative labels have their place and should be considered as well as the quantified values. A lot of the morphological characteristics of the neonatal EEG maturation process are associated with different continuity states of the EEG signal. A classifier that gives labels corresponding to the existing categories can help automate the recognition of these landmark patterns in an EEG signal.

As discussed in the Section 9.3, a classifier for the qualitative labels should take into account the natural clustering that may occur in the recording. Using the quantitative feature ($\hat{\mu}$ and $\hat{\sigma}$) helps to make the label assignment less subjective, and the resulting labels can better reflect the continuity states that are present in the recording. As shown in Chapter 6, the continuity feature has a tendency to cluster around the extreme ends of the feature distributions. This reflects the way the continuity of the EEG signal remains relatively stable for the duration of the signal where the continuity label remains unchanged. In a normal preterm infant, the signal either remains in one state for the duration of the recording, or the continuity of the signal oscillates between continuous and discontinuous (a sign of the sleep-wake cycle). As shown in Chapter 7, as the infant grows towards term, the states stabilise at different points in the continuity feature space. Using an unsupervised clustering method can ensure that the resulting labels reflect the intrinsic clusters and gives a better idea where the states stabilise in the feature space for a particular recording.

9.5 Continuity as the Context for EEG Analysis

Having a way to measure continuity helps to put the EEG recording into context. Because the behaviour of EEG signals varies a lot within a recording, a way to provide the context of the signal is very important when comparing multiple recordings. Features such as those relating to the time-frequency distribution of the signals can be

extracted from EEG, but without any context the variation of these features within a signal far exceeds the differences of the features between regions of the recordings in similar states. With the continuity as the context, features can be compared between recordings.

As continuity related features such as interburst interval have already been found to correlate with the maturation and clinical outcomes of infants, the continuity index and the qualitative labels can be used to help automatically track these features. Since interburst period is related to certain states of EEG signal, namely burst suppression, the continuity index or the qualitative label can be used to highlight areas of interest, as well as being the basis of an automatic burst and suppression detection algorithm.

9.6 Maturation Index

Continuity is one aspect of EEG that clinicians monitor to determine the neural maturation progress of preterm infants. Currently, this is done by EEG specialists who have been trained to recognise the continuity changes in EEGs of infants as they mature. The maturation index proposed in this research serves as a starting point to quantify these changes. Current medical literature describes the changes in continuity during the maturation process in preterm EEG as a shortening of interburst intervals; the age at which continuous signals start appearing (approximately 30 weeks); and the appearence of a sleep wake-cycle starting at approximately 34 weeks [11]. However, the more subtle changes are harder to describe using the qualitative labels currently assigned in medical literature. As shown in Chapter 7, the continuity feature appears to shift in a certain direction in EEG recordings as the infant grows towards term. This direction is the same as the minor component of a PCA transformation, and therefore a maturation index can be calculated as a result, solving the problem of using qualitative labels to describe changes in EEG.

Like the continuity index, the maturation index is also affected by muscle artifact, since it is derived from the same continuity feature. The maturation index is also only an estimate, and the relative changes in an infant's recording with respect to previous recordings are more representative of the maturation than the actual value of the maturation index. However, maturation indices of normal EEGs appear to fall within a certain range, suggesting that there may be a rate of change that is common to all healthy preterm infants. More investigation is needed to confirm this growth

pattern.

The continuity maturation index may be a good guide for maturation, but continuity is not the only measurement of maturation in neonatal EEG. Since the maturation index presented here is based only on the continuity of the infant, it can only be used as a guide to determine the maturation progress of the infant, and will need other aspects, such as synchrony and sleep-wake cycle, to comprehensively describe the maturation of the EEG.

9.7 Summary

The quantitative feature presented in this thesis has been determined to be a good candidate for use in continuity analysis. The two dimensions of the feature, the estimated mean and standard deviation of the log amplitude vector, are analogous to the current aspects of aEEG that are used to determine the continuity of the EEG signal, namely the lower and upper bands (the minima and maxima of the envelope of the EEG signal) and the bandwidth (the difference between the lower and upper bands). The feature provides a way to quantify continuity during an EEG recording, and can be used as the input in a classifier to produce qualitative labels for the continuity states. The quantitative continuity index helps clinicians to quantify continuity, and shows the local variations in continuity during an EEG recording, rather than giving a global indication similar to the values currently used as continuity related features. This continuity feature is also analogous to the aspects of aEEG used to determine continuity. Because the standard deviation of the amplitude distribution is used instead of the minima and maxima, as in the case of aEEG, the results are smoother with less local variation. Therefore, the feature can be used to produce a plot similar to the aEEG plot with smoother edges, making it easier to read. Using a clustering algorithm with the continuity feature, qualitative labelling helps translate the quantitative labels into labels familiar to clinicans, while labelling the signal to reflect any intrinsic clusters that may be present. A maturation index can also be derived from the continuity feature that increases as the infant grows towards term, since the continuity feature shifts in one direction in the feature space. This maturation index gives a rough idea of the maturation progress of the infant, and further study with other relevant features can help refine it further.

Chapter 10

Conclusions and Future Works

10.1 Conclusions

This thesis aimed to investigate the use of signal-processing techniques to aid and advance research in preterm neonatal EEG signal analysis. The feature proposed in this project is a quantitative continuity feature that can be used to produce a quantitative continuity index, and aid in qualitative continuity classification.

A literature review of the medical research in this field has revealed that medical researchers have been correlating continuity related features, such as inter-burst periods and burst amplitude, to maturation progress and clinical outcomes. There are also guidelines for visually inspecting EEG signals for continuity, symmetry, synchrony, and other landmark patterns to determine the maturation progress of infants. Infants whose EEG does not reflect the expected maturation progression are said to have dysmature EEG patterns and are considered abnormal. Continuity plays a critical role in EEG analysis not only because it is a good index of maturation, but also because it provides a context to other criteria for maturation evaluation, such as synchrony and symmetry. Currently, continuity is a qualitative measurement, with some related measurements used in clinical research. With a quantitative measurement, the process of evaluating the maturation of an EEG signal can be made more objective, as well as providing a context for further numerical EEG analysis.

Continuity refers to the variations in amplitude throughout the duration of a recording. For very young preterm infants, EEG recordings usually consist of high

amplitude bursts, separated by periods of very low voltage inactivity. This pattern is referred to as *tracé discontinu* or burst suppression in medical literature and is considered normal in very young preterm infants. The bursts gradually become lower in amplitude, while the inactivity between bursts becomes shorter and increases in amplitude. Around the age of 31 weeks, continuous signals, where the amplitude remains constant for a long period of time, start emerging. Continuity is normally determined by visual inspection of the raw EEG signal. Another common standard is the use of aEEG, where an estimate of the envelope of the signal is plotted on a compressed time scale, and simple threshold guidelines are established to classify the different patterns. In terms of normal preterm neonatal EEG, the three common patterns are continuous, discontinuous (*tracé alternant*) and burst suppression (*tracé discontinu*).

Initial investigation using joint time-frequency analysis has shown that different background continuities behave differently in the time-frequency domain. However, normal EEG does not form any recognisable pattern in the time-frequency domain. Since EEG signals are nonstationary, the behaviour of the signal changes over time. Therefore, without the context given by continuity, the features extracted from an EEG recording may vary depending on the state of the signal. This state dependent variation is likely to be larger than the variation between two recordings during similar background states. Therefore, the continuity state of the signal is crucial to further feature extraction and analysis.

In order to quantify continuity, the signal is first segmented, using the generalised likelihood ratio (GLR) method, into psuedo-stationary segments. The GLR segmentation method was one of three investigated in this project, and was found to offer the best compromise between segment boundary detection and false detection rates. The GLR segmentation method is based on the predictive error in an autoregressive model. Windows of test signal are joined to the reference windows to produce a pooled window, and the predictive error of the pooled window using an autoregressive model is compared with the predictive error for the testing window and the reference window. A threshold is defined for testing signals that belong to the reference window. If the predictive error exceeds the threshold, a segmentation boundary is defined at the position where the difference of the predictive errors is maximum. The result is a set of segment boundaries that divide the EEG signal into a series of psuedo-stationary segments.

The psuedo-stationary segments are the basis for the amplitude vector: a vector with the same length as the signal which estimates the envelope of the signal. Each segment is then rectified and averaged, and the mean absolute value is used as the amplitude for the duration of the segment. Since the continuity of the signal is essentially a description of the amplitude distribution, the statistical parameters of the amplitude vector were used as the continuity feature. The log-normal distribution was chosen as the distribution of the amplitude, and the estimated mean and standard deviation of the log value of the amplitude vector were used as the continuity feature. The features of segments with manually assigned background states were plotted against their assigned states, and a correlation between the continuity feature and the continuity of the signal was found. Furthermore, the continuity of the signal was found to coincide with the major axis of the feature. A weaker correlation was found between the minor axis and the age of the infant. The feature is analogous to the aspects of aEEG that clinicians use to determine continuity. Instead of using the minima and maxima of the signal envelope, the standard deviation of the log amplitude vector is used, and hence the feature is less prone to local variation and outliers. Plotting the feature ($\hat{\mu}$ and $\hat{\mu} \pm \hat{\sigma}$) is analogous to the aEEG plot, and provides a smoother plot than aEEG.

To quantify continuity and produce a continuity index, the continuity feature is transformed using principal component analysis (PCA). Since the continuity of the signal is correlated with the major axis of the feature dataset, by performing PCA on the dataset, the principal component will indicate the continuity of the signal. This is plotted against existing background detection methods for term infants and the aEEG signal to verify it as an indicator of the continuity of the signal.

The minor component of the recording is averaged to come up with one maturation index of the recording. The minor component does not change dramatically over the course of the recording, as it corresponds to the minor axis of the dataset. The mean of the minor component gives one value per recording to correlate with the conceptional age (CA) of the infant. This maturation index is plotted against the CA of preterm infants with healthy clinical outcomes. The maturation index increases as the CA increases, though this number seems to be infant specific, and the relative changes in the numbers during the maturation of the infant seem to be a more reliable indication for EEG maturation than the specific values that the maturation indices display.

Qualitative classification was also investigated, since clinicians are already familiar with the labels, and it provides a way to label the EEG recording into various states, which can be used to model the signal. Linear Discriminant Anlysis (LDA), Self Organising Maps (SOM) and Gaussian Mixture Model (GMM) have been investigated as possible classifiers. LDA is the only supervised classifier of the three methods investigated. The advantage of a supervised system is the fact that the output labels are controlled. Therefore, using a set of segments manually labelled with conventional background state labels as training inputs, the result would be a classifier that produces labels for the segments that are familiar to clinicians and researchers. However, this approach does not detect any intrinsic clusters that may occur in the recording. This means that, while it is a good system for producing labels for further study, it does not give any extra information about the distribution of the feature. SOM is an unsupervised classifier that produces a mapping of the training data. The mapping can be analysed to produce clusters that reflect any intrinsic clusters present in the training dataset. The intrinsic clusters that appear in the data, however, may not represent the labels currently used for continuity, especially if the recording only contains one continuity state (e.g. for very young preterm infants) or when movement artifact is present in the recording. Also, using SOM for clustering purposes makes soft labelling difficult, as the data being clustered are prototypes of the training data, and the classification is based on the best matched unit in the SOM rather than the data itself. GMM is another unsupervised classification method investigated. Each state is modelled as a Gaussian distribution in the feature space. Each data point is considered a mixture of the models and its likelihood of belonging to each of the states is calculated. This produces probabilities of the EEG signal belonging to each background state for each feature vector. These probabilities can be used as soft labels for the signal, or can be further processed as input for a "winner-takes-all" network that gives the labels of the most probable continuity state. Like the SOM classifier, the GMM is unsupervised and therefore suffers the same disadvantage of mislabelling when the recording is noisy or consists of only one state. However, the centres of the Gaussian distribution can be analysed to ensure that clusters found in the dataset reflect the continuity states of the signal.

It is therefore concluded that the proposed feature, using the parameters of the amplitude statistics of EEG, can be used as a quantitative description of neonatal EEG continuity. This feature has been shown to correlate with aspects of aEEG related to continuity. The continuity and maturation index derived from this feature

have been shown to correlate with the continuity and maturation states of the EEG respectively. The feature has also been shown to be a good candidate as the input to classification systems to automate continuity labelling.

10.2 Future Works

The continuity feature described in the previous chapters serves as a starting point for further investigation into preterm EEG maturation and abnormality detection. This section outlines some suggestions for further investigation that take advantage of the context provided by the continuity feature.

10.2.1 Automatic Sleep Pattern Detection

Preterm infants have very obvious sleep patterns, which are related to the continuity of their EEG signals. Normal sleep-wake cycles appear in infants starting at around 30 weeks CA. The sleep-wake cycle usually appears in EEG as alternating discontinuous and continuous periods, since EEGs during sleep periods usually appear as discontinous, andwakeful periods continuous. Knowing this, sleep-wake cycles can be detected by detecting fluctuations in continuity. The state of the infant can be inferred from the continuity of the EEG. This can be implemented on bedside monitoring systems to help identify the state of the infants and avoid disruption of sleep patterns for treatment and routine caring procedures.

10.2.2 Further Feature Analysis

One of the problems with EEG feature analysis is the non-stationary nature of the signal, and the marked difference in behaviour during different sleep and wakeful states. With the context provided by the continuity measurement, the feature extracted from the signal can be compared across different recordings. This ensures that the differences between the recordings are not caused by different states. For example, comparing the dominant frequency of randomly selected segments in different recordings might result in comparing segments from different states. Manually picking the segments can overcome the problem to some extent, but the selected segments will then be chosen subjectively. Using the continuity as a context, the distribution

of the states can be analysed, and the segments can be selected using the continuity measurements to ensure the selection is objective.

10.2.3 Burst Suppression Analysis

Burst suppression in preterm infants is considered a normal pattern. However, the interburst interval and other features such as burst amplitude have been shown to correlate with clinical outcomes. Burst suppression can be detected using the continuity feature as shown in Chapter 8. Features such as interburst interval and burst amplitude are related to continuity. These continuity features can be used to generate a quantitative measurement of continuity (see Chapters 6 and 7), which can then be correlated with the clinical outcomes of the infants, and the normal maturation process. The GLR segmentation method outlined in Chapter 5 can be used as a new way to define bursts and interburst periods of the burst suppression pattern, rather than using simple threshold methods. Since the GLR method is modelled using autoregressive models, any significant spectral changes in the signal will trigger a new segment boundary. This means the segments generated from this method are psuedo-stationary, and using the segments and their amplitude to identify interburst periods and bursts can give a more thorough definition than a simple threshold method.

10.2.4 Maturation Index

In the course of this investigation, the idea of using the minor component of the continuity feature was explored. The minor component is the second component of the continuity feature after performing PCA on the data. The mean of the minor component in a recording appears to increase as the CA of the infant increases. This serves as a starting point for further investigation into how continuity is affected by the maturation process, and more research can be done to correlate the change or lack of change in this minor component, to develop a maturation index that is based on the continuity of the recording. Since the feature proposed by this thesis in effect quantifies the traditionally qualitative aspect of continuity, the maturation index is analogous to the way clinicians estimate the maturation progress of an infant by visually inspecting the EEG signal for continuity information. Further investigation can also examine maturation during different states of infant EEG, since changes in state

(e.g. "continuous" and "discontinuous") may develop at different rates throughout the maturation process.

List of References

[1] B.J. Fisch. *Fisch and Spehlmann's EEG primer*. Elsevier Science Publishers B.V., Amsterdam, 3 edition, 1999.

[2] S. Marret, D. Parain, J.-F. Ménard, T. Blanc, A.-M. Devaux, P. Ensel, C. Fessard, and D. Samson-Dollfus. Prognostic value of neonatal electroencephalography in premature newborns less than 33 weeks of gestational age. *Electroencephalography and Clinical Neurophysiology*, 102(3):178–185, 1997.

[3] K. Watanabe, F. Hayakawa, and A. Okumura. Neonatal EEG: a powerful tool in the assessment of brain damage in preterm infants. *Brain and Development*, 21(6):361–372, 1999.

[4] D. Selton, Andre M., and Hascoët J.M. Normal EEG in very premature infants: reference criteria. *Clinical Neurophysiology*, 111(12):2116–2124, 2000.

[5] P.B. Colditz, L.J. Buck, K. Foster, and B.E. Lingwood. Can signal processing help prevent brain damage in the newborn? In *Proceedings of the Fifth International Symposium on Signal Processing and Its Applications, ISSPA '99*, volume 1, pages 345 – 349, 1999.

[6] BrainZ BRM2 brain monitor: a bedside continuous EEG and cerebral function monitor, 2005. [Online] Available: http://www.brainzinstruments.com/index.asp [30 June 2005].

[7] J. Fawke. Neurological outcomes following preterm birth. *Seminars in Fetal and Neonatal Medicine*, 12(5):374–382, 2007.

[8] J.J. Neil and T.E. Inder. Imaging perinatal brain injury in premature infants. *Semin Perinatol.*, (6):433 – 443, 2004.

[9] L. Hellström-Westas, L.S. de Vries, and Rosén I. *An Atlas of Amplitude-Integrated EEGs in the Newborn*. The Parthenon Publishing Group, New York, 2003.

[10] C.T. Lambroso. Neonatal polygraphy in full-term and premature infants: A review of normal and abnormal findings. *Journal of Clinical Neurophysiology*, 2(2):105–155, 1985.

[11] E.M Mizrahi, R.A. Hrachovy, and Kellaway P. *Atlas of Neonatal Electroencephalography*. Lippincott Williams & Wilkins, Philadelphia, 3 edition, 2003.

[12] S. Victor, R.E. Appleton, M. Beirne, A.G. Marson, and A.M. Weildling. Spectral analysis of electroencephalography in premature newborn infants: Normal ranges. *Pediatric Research*, 57(3):336–341, 2005.

[13] E. Biagioni, L. Bartalena, A. Boldrini, G. Cioni, S. Giancola, and A. E. Ipata. Background EEG activity in preterm infants: correlation of outcome with selected maturational features. *Electroencephalography and Clinical Neurophysiology*, 91(3):154–162, 1994.

[14] C.R. West, J.E. Harding, C.E. Williams, M.I. Gunning, and R.B. Malcolm. Quantitative electroencephalographic patterns in normal preterm infants over the first week after birth. *Early Human Development*, 82:43 – 51, 2006.

[15] V.F. Burdjalov, S. Baumgart, and A.R. Spitzer. Cerebral function monitoring: A new scoring system for the evaluation of brain maturation in neonates. In *Pediatrics*, volume 112, pages 855–861, 2003.

[16] L. Hellström-Westas and Rosén I. Specific monitoring of neonatal brain function with optimized frequency bands. *Engineering in Medicine and Biology Magazine, IEEE*, 20(5):40–46, 2001.

[17] E. Biagioni, L. Bartalena, A. Boldrini, R. Pieri, and G. Cioni. Constantly discontinuous EEG patterns in full-term neonates with hypoxic-ischaemic encephalopathy. *Clinical Neurophysiology*, 110(9):1510–1515, 1999.

[18] B. Boashash and M. Mesbah. A time-frequency approach for newborn seizure detection. *Engineering in Medicine and Biology Magazine, IEEE*, 20(5):54–64, 2001.

[19] Antonia Papandreou-Suppappola, editor. *Applications in time-frequency signal processing*. CRC Press, Boca Raton, 2003.

[20] P. Zarjam, G. Azemi, M. Mesbah, and B. Boashash. Detection of newborns' EEG seizure using time-frequency divergence measures. In *Acoustics, Speech, and Signal Processing, 2004. Proceedings. (ICASSP '04). IEEE International Conference on*, volume 5, pages V–429 – V–432, 2004.

[21] M.E. Saab and J. Gotman. A system to detect the onset of epileptic seizures in scalp EEG. *Clinical Neurophysiology*, Article in Press, Corrected Proof, 2004.

[22] Ö. Özdamar, I. Yaylali, P. Jayaker, and C.N. Lopez. Multilevel neural network system for EEG spike detection. In *Proceedings of the Fourth Annual IEEE Symposium Computer-Based Medical Systems, 1991.*, pages 272 – 279, 1991.

[23] Ö. Özdamar, Guanglong Zhu, L. Yaylali, and P. Jayakar. Real-time detection of EEG spikes using neural networks. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 1992. Vol.14.*, volume 3, pages 1022 – 1023, 1992.

[24] Ö. Özdamar, C.N. Lopez, and I. Yaylah. Detection of transient EEG patterns with adaptive unsupervised neural networks. In *Proceedings of the 1992 International Biomedical Engineering Days, 1992.*, pages 192 – 197, 1992.

[25] T. Kalayci, Ö. Özdamar, and N. Erdöl. The use of wavelet transform as a preprocessor for the neural network detection of EEG spikes. In *Proceedings of the 1994 IEEE Southeastcon '94. 'Creative Technology Transfer - A Global Affair'.*, pages 1 – 3, 1994.

[26] T. Kalayci and Ö. Özdamar. Wavelet preprocessing for automated neural network detection of EEG spikes. *Engineering in Medicine and Biology Magazine, IEEE*, 14(2):160 – 166, 1995.

[27] P. Zarjam, M. Mesbah, and B. Boashash. An optimal feature set for seizure detection systems for newborn EEG signals. In *Circuits and Systems, 2003. ISCAS '03. Proceedings of the 2003 International Symposium on*, volume 5, pages V–33 – V–36, 2003.

[28] P. Zarjam, M. Mesbah, and B. Boashash. Detection of newborn EEG seizure using optimal features based on discrete wavelet transform. In *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03). 2003 IEEE International Conference on*, volume 2, pages II–265 – II–268, 2003.

[29] D. Hoyer, R. Bauer, K. Conrad, M. Galicki, A. Doring, H. Hoyer, B. Walter, H. Witie, and U. Zwiener. Specific monitoring of neonatal brain function with optimized frequency bands. *Engineering in Medicine and Biology Magazine, IEEE*, 20(5):40–46, 2001.

[30] V. Goel, A.M. Brambrink, A. Baykal, D.F. Koehler, R.C.and Hanley, and N.V. Thakor. Dominant frequency analysis of eeg reveals brain's response during injury and recovery. *IEEE Transactions on Biomedical Engineering*, 43(11):1083 – 1092, 1996.

[31] N. Löfgren, K. Lindecrantz, A. Flisberg, R. Bågenholm, I. I Kjellmer, and M. Thordstein. Spectral distance for arma models applied to electroencephalogram for early detection of hypoxia. *Journal of Neural Engineering*, 3:227–234, 2006.

[32] F. Itakura. Minimum prediction residual principle applied to speech recognition. *IEEE Trans. Acoust. Speech Signal Process*, 23:67 – 72, 1975.

[33] R. Agarwal and J. Gotman. Adaptive segmentation of electroencephalographic data using a nonlinear energy operator. In *Proceedings of the 1999 IEEE International Symposium on Circuits and Systems, 1999. ISCAS '99.*, volume 4, pages 199 – 202, Orlando, FL USA, 1999. IEEE.

[34] Rajeev Agarwal, Jean Gotman, Danny Flanagan, and Bernard Rosenblatt. Automatic EEG analysis during long-term monitoring in the icu. *Electroencephalography and Clinical Neurophysiology*, 107(1):44–58, 1998.

[35] R. Agarwal and J. Gotman. Long-term EEG compression for intensive-care settings. *Engineering in Medicine and Biology Magazine, IEEE*, 20(5):23 – 29, 2001.

[36] J.P. Turnbull, K.A. Loparo, M.W. Johnson, and M.S. Scher. Automated detection of tracé alternant during sleep in healthy full-term neonantes using discrete wavelet transform. *Clinical Neurophysiology*, 112:1893 – 1900, 2001.

[37] M.A. Navakatikyan. New compressed representation of EEG for the monitoring of brain function. *Acta Pædiatrica*, 96:144, 2007.

[38] M.A. Navakatikyan, J.D.E. Barks, G. Greisen, and M.C Mathur, A.and. Toet. Automatic background classification of the amplitude-integrated EEG in infants. *Acta Pædiatrica*, 96:144, 2007.

[39] B. Boashash, editor. *Time frequency signal analysis-methods and applications*. Longman Cheshire Pty Limited, Melbourne, 1992.

[40] B. Boashash, editor. *Time frequency signal analysis and processing : a comprehensive reference*. Elsevier, Boston, 2003.

[41] L. Cohen. *Time-frequency analysis*. Prentice Hall PTR, Englewood Cliffs, N.J, 1995.

[42] H.I. Choi and W.J. William. Improved time-frequency representation of multicomponent signals using exponential kernels. *IEEE Trans on Acoust., Speech, Signal Processing*, 37:862–871, 1989.

[43] B. Barkat and B. Boashash. A high-resolution quadratic time-frequency distribution for multicomponent signals analysis. *IEEE Transactions on Signal Processing*, 49(10):2232 – 2239, 2001.

[44] Z.M. Hussain and B. Boashash. Multi-component IF estimation. In *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing, 2000.*, pages 559 – 563, 2000.

[45] Y. Zhao, L.E. Atlas, and R.J. Marks. The use of cone-shaped kernels for generalized time-frequency representations of nonstationary signals. *IEEE Trans. Acoust. Speech, Signal Processing*, 38:1084 – 1091, 1990.

[46] Steven R. Parker and Ashwin Natarajan. Optimal time-frequency representations for the analysis of EEG waveforms. Technical report, Arizona State University Ira A. Fulton School of Engineering, 2002.

[47] B. Boashash, A.M. Zoubir, and M. Roessgen. On-line detection of seizure in newborn EEG using signal processing tools. In *Digital Signal Processing Proceedings, 1997. DSP 97., 1997 13th International Conference on*, volume 1, pages 79 – 82, Santorini Greece, 1997.

[48] B. Boashash, P. Barklem, and M. Keir. Detection of seizure signals in newborns. In *Acoustics, Speech, and Signal Processing, 1999. ICASSP '99. Proceedings., 1999 IEEE International Conference on*, volume 4, pages 2351 – 2354, 1999.

[49] B. Boashash, H. Carson, and M. Mesbah. Detection of seizures in newborns using time-frequency analysis of EEG signals. In *Statistical Signal and Array Processing, 2000. Proceedings of the Tenth IEEE Workshop on*, pages 564 – 568, Pocono Manor, PA USA, 2000.

[50] B. Boashash, M. Mesbah, and P. Colditz. Newborn EEG seizure pattern characterisation using time-frequency analysis. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP '01). 2001 IEEE International Conference on*, volume 2, pages 1041 – 1044, 2001.

[51] G. Bodenstein and H.M. Praetorius. Feature extraction from the electroencephalogram by adaptive segmentation. *Proceedings of the IEEE*, 55(5):642 – 652, 1977.

[52] R.M. Rangayyan. *Biomedical signal analysis : a case-study approach*. IEEE Press, New York, 2002.

[53] U. Appel and A.V. Brandt. Adaptive sequential segmentation of piecewise stationary time series. *Information Sciences*, 29:27 – 56, 1983.

[54] L. Wong and W. Abdulla. Time-frequency evaluation of segmentation methods for neonatal EEG signals. In *Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1303 – 1306, 2006.

[55] P.F. Prior and D.E. Maynard. *Monitoring Cerebral Function: Long-Term Monitoring of EEG and Evoked Potentials*. Elsevier Science Publishers B.V., Amsterdam, 1986.

[56] E. Limpert and M. Stahel, W.A.and Abbt. Log-normal distributions across the sciences: Keys and clues. *BioScience*, 51(5):341–352, 2001.

[57] L. Wong and W. Abdulla. Automatic detection of preterm neonatal EEG background states. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2008)*, pages 421 – 424, 2008.

[58] A.M. Martínez and A.C. Kak. Pca versus lda. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:228–233, 2001.

[59] PRTools: The matlab toolbox for pattern recognition., 2008. [Online] http://www.prtools.org/ [30 June 2008].

[60] SOM toolbox, 2005. [Online] http://www.cis.hut.fi/projects/somtoolbox/ [30 June 2008].

Appendix A

Database of Preterm Infant EEG

The following table contains a list of the infants enrolled in the VIBeS study whose EEG recordings were made available for this research. Only infants with known GA were listed here. Note that the neurological follow-up information is not available for every infant, and some infants have fewer than four recordings on file. The attributes of the infants recorded are:

- **VIBeS ID**: ID of the infant at the VIBeS study
- **GA at Birth**: Number of weeks since conception at birth
- **Birth Weight**: Weight of infant at birth measured in grams (g)
- **WMI Index**: White Matter Injury (WMI) score as determined from MRI scan at term (1 = no WMI, 4 = Severe WMI)
- **IVH Grade**: Intraventricular Hemorrhage (IVH) grade as determined from MRI scan at term (0 = no IVH, 4 = severe IVH)
- **Dis. Score**: Disability score as determined by neurological (MDI and PDI) tests performed at 2 years (1 = normal in both test, 2 = abnormal in one of the tests, 3 = abnormal in both tests)
- **MDI**: Mental Development Index (MDI), a measurement for language and logic skills in children, as determined at the 2-year neurological clinical followup
- **PDI**: Psychomotor Development Index (PDI), a measurement for motor skill in children, as determined at the 2-year neurological clinical followup

- **CA1, CA2, CA3, CA4**: Conceptional age of the infant at the first, second, third and fourth EEG recording, respectively.

VIBeS ID	GA at birth	Birth Weight	WMI Index	IVH Grade	Dis. Score	MDI	PDI	CA1	CA2	CA3	CA4
40	29.2	1075	1	0	3	94	88	30.20	31.20	32.20	34.20
41	29.2	1130	1	0	3	104	100	30.34	31.20	32.20	33.77
42	29.2	1126	NA	0	3	104	96	30.20	31.20	32.20	34.20
45	28.0	1140	2	0	3	102	84	31.00	31.86	32.86	NA
47	29.1	940	1	0	2	74	65	30.10	31.10	32.10	33.10
48	24.1	684	2	0	3	86	84	31.10	32.10	32.96	34.96
49	24.1	520	2	0	2	76	84	31.10	32.81	35.81	NA
52	27.0	950	3	0	3	94	100	29.14	31.00	33.00	NA
55	26.3	739	3	0	1	52	81	26.59	27.59	28.73	30.44
56	28.0	1067	3	0	1	45	57	28.14	29.14	30.00	32.00
59	25.1	819	2	0	3	122	107	25.39	26.39	27.53	29.39
60	28.4	1091	1	0	3	100	100	28.97	29.97	30.97	32.97
61	28.4	1206	1	0	3	98	92	28.97	29.97	30.97	32.97
62	28.4	1160	NA	0	3	102	96	28.97	29.97	30.97	32.97
74	25.0	956	2	0	2	82	88	25.57	26.43	27.57	29.57
75	27.0	844	1	0	NA	NA	NA	27.57	28.14	29.00	31.43
76	27.0	1134	1	0	2	72	73	27.57	28.14	29.00	31.43
77	25.0	414	3	0	2	74	84	25.71	26.71	27.71	28.71
79	31.0	950	3	1	2	72	96	31.14	31.86	32.86	34.71
80	28.0	890	NA	0	2	80	80	28.43	28.86	29.86	31.71
81	22.0	620	3	0	3	94	100	23.00	24.57	25.57	26.57
82	28.0	1086	2	0	2	76	84	28.29	29.00	30.00	32.00
85	24.0	635	1	2	3	106	117	24.00	25.00	26.00	28.29
98	25.2	780	NA	0	3	94	103	26.34	27.34	28.20	30.06
99	25.0	665	NA	0	3	106	113	25.71	26.57	27.71	29.71
100	25.2	750	2	0	3	92	107	26.20	27.34	28.91	31.77
103	27.3	1220	3	0	3	106	92	30.87	31.59	32.59	34.59
107	26.6	1156	1	0	3	100	110	30.46	32.60	34.89	NA
108	26.6	1120	1	0	3	112	110	29.89	30.89	32.03	33.89
109	31.0	965	1	0	3	98	107	31.14	31.86	32.86	NA
114	27.0	1040	1	0	3	98	88	27.57	27.86	28.86	30.86
115	30.0	1120	4	2	2	NA	NA	30.71	31.00	32.00	34.00
116	26.0	874	3	0	3	102	92	26.29	26.86	27.86	29.86
118	28.6	576	4	0	NA	NA	NA	29.74	30.46	31.17	38.60
121	26.6	806	2	0	3	96	84	27.03	27.46	28.46	30.46
122	28.5	970	1	0	3	93	94	29.64	30.21	31.07	32.50
123	27.0	1040	2	4	2	68	84	27.29	27.86	28.86	34.43

VIBeS ID	GA at birth	Birth Weight	WMI Index	IVH Grade	Dis. Score	MDI	PDI	CA1	CA2	CA3	CA4
124	31.0	1240	2	0	2	76	88	31.43	32.00	33.00	35.00
125	27.0	1084	2	0	3	90	103	27.43	28.00	29.00	31.00
126	27.0	1025	2	0	3	88	88	27.43	28.00	29.00	31.00
130	29.0	764	1	0	1	NA	NA	29.71	30.00	31.00	33.00
131	29.0	1135	2	0	3	92	84	29.43	30.00	31.00	33.00
136	25.0	660	2	0	2	68	69	25.86	26.29	27.00	29.14
137	28.0	1044	2	0	3	100	113	28.29	28.86	29.86	31.86
138	28.0	1025	2	0	3	89	96	28.14	28.86	29.86	31.86
140	29.6	1395	2	0	3	110	84	30.31	31.03	31.60	33.46
141	29.6	1130	1	0	3	100	113	30.31	31.03	31.60	33.46
150	26.0	815	2	0	2	78	103	26.14	27.14	27.86	28.86
152	28.0	1114	2	0	2	70	84	28.29	28.86	29.86	31.86
155	32.0	1200	1	0	2	82	92	32.14	34.14	NA	NA
159	24.3	600	2	0	2	80	84	24.59	25.16	26.16	28.30
160	28.0	1085	3	0	1	40	40	28.29	28.86	29.86	33.43
163	27.4	745	2	0	2	80	88	27.83	28.40	29.40	31.40
164	27.4	1096	3	4	2	76	77	27.83	28.40	29.40	31.40
173	31.0	900	2	0	2	80	84	31.71	32.29	33.00	35.00
174	31.0	1265	3	0	2	78	80	35.00	NA	NA	NA
175	25.1	756	2	2	2	68	96	25.39	26.10	27.10	29.10
186	24.3	NA	NA	NA	NA	NA	NA	25.16	26.16	NA	NA
187	29.0	NA	NA	NA	NA	NA	NA	29.14	29.86	30.86	32.86
200	27.2	790	3	0	1	45	84	27.20	28.20	29.20	31.20
203	27.6	NA	NA	NA	NA	NA	NA	29.89	30.31	31.60	NA
214	26.4	NA	NA	NA	NA	NA	NA	26.83	27.40	28.40	30.40
221	31.6	NA	NA	NA	NA	NA	NA	32.60	33.46	34.46	35.74
228	29.0	1180	2	0	3	96	100	29.43	30.14	31.14	33.00
234	30.4	NA	NA	NA	NA	NA	NA	30.97	31.40	32.69	34.54
250	27.2	NA	NA	NA	NA	NA	NA	27.77	28.20	29.20	NA
268	31.4	NA	NA	NA	NA	NA	NA	31.97	34.11	35.26	37.26
280	25.0	NA	NA	NA	NA	NA	NA	26.57	27.43	28.57	30.43
284	28.6	NA	NA	NA	NA	NA	NA	28.60	29.89	30.46	32.46
285	29.0	NA	NA	NA	NA	NA	NA	29.43	30.00	NA	NA
286	24.0	NA	NA	NA	NA	NA	NA	24.14	25.00	26.00	26.14
289	28.0	NA	NA	NA	NA	NA	NA	28.71	29.00	30.00	32.00
290	28.0	NA	NA	NA	NA	NA	NA	28.86	29.14	30.00	32.00
293	30.2	NA	NA	NA	NA	NA	NA	30.77	31.49	32.06	34.06
296	28.3	NA	NA	NA	NA	NA	NA	28.59	29.30	30.30	32.30
297	25.4	NA	NA	NA	NA	NA	NA	25.83	NA	NA	NA

VIBeS ID	GA at birth	Birth Weight	WMI Index	IVH Grade	Dis. Score	MDI	PDI	CA1	CA2	CA3	CA4
300	25.5	NA	NA	NA	NA	NA	NA	26.21	27.07	NA	NA
304	25.6	NA	NA	NA	NA	NA	NA	27.74	28.89	29.74	30.74
307	27.4	NA	NA	NA	NA	NA	NA	27.97	30.69	31.54	32.11
314	25.6	NA	NA	NA	NA	NA	NA	27.89	NA	NA	NA
326	26.5	865	3	0	2	70	84	26.64	27.21	27.21	28.21
337	26.2	635	NA	0	1	40	45	27.20	27.63	28.20	30.06
341	27.0	NA	NA	NA	NA	NA	NA	27.86	28.71	29.71	30.71
348	27.0	NA	NA	NA	NA	NA	NA	27.71	28.57	NA	NA
351	28.0	NA	NA	NA	NA	NA	NA	28.43	28.71	29.86	31.86
358	29.1	895	2	0	3	90	96	29.53	30.10	31.10	33.10
361	28.0	1180	NA	0	3	88	92	28.43	29.00	30.00	NA
363	30.0	830	2	0	2	84	100	30.43	31.00	NA	NA
364	26.0	690	2	0	2	66	85	26.00	26.86	27.86	29.86
367	24.0	680	2	3	2	78	96	24.29	25.00	25.86	28.00
370	26.5	1060	2	0	3	86	96	26.93	27.21	28.50	30.36
371	26.5	930	2	0	3	86	84	26.93	27.21	28.50	30.36
386	29.0	NA	NA	NA	NA	NA	NA	29.29	29.86	30.86	32.86
389	27.0	675	2	0	1	45	103	27.29	27.86	28.86	30.86
390	27.0	830	2	2	2	62	92	27.29	27.86	28.86	30.86
418	24.4	741	3	3	1	45	88	24.69	25.40	26.40	28.40
419	26.0	575	2	0	1	40	40	26.57	27.00	28.00	32.00
420	24.4	675	1	0	2	74	96	24.40	25.40	26.40	28.40
422	25.3	821	2	0	2	72	96	25.30	26.16	27.16	29.16
565	31.3	NA	NA	NA	NA	NA	NA	31.87	32.30	33.30	35.30
566	29.4	950	2	0	2	66	45	29.40	30.26	31.26	33.26
571	28.6	1191	2	0	3	108	97	28.89	29.60	30.46	NA
772	23.5	592	2	0	3	96	98	23.79	24.21	25.50	27.36
780	28.6	NA	NA	NA	NA	NA	NA	28.89	29.46	30.46	31.74
803	28.3	1157	4	0	1	40	40	29.01	29.59	NA	NA
804	28.3	1036	2	0	3	100	100	29.01	29.59	NA	NA
811	26.2	886	1	0	2	76	96	26.34	27.20	28.06	30.06
820	29.4	NA	NA	NA	NA	NA	NA	29.83	30.69	31.26	NA
821	25.0	660	1	0	3	112	121	25.86	26.43	27.29	29.14
822	29.4	916	2	0	2	80	96	29.97	30.54	31.11	32.97
823	27.4	1100	2	1	3	86	88	27.54	28.26	29.26	31.26
835	26.2	920	2	0	2	72	96	26.49	27.20	28.20	30.06
836	26.2	875	1	0	2	60	92	26.49	27.20	28.20	30.20
853	27.6	1295	1	9	3	106	96	28.03	28.74	NA	NA
854	27.6	1165	1	9	3	100	88	28.17	28.74	29.46	31.46

VIBeS ID	GA at birth	Birth Weight	WMI Index	IVH Grade	Dis. Score	MDI	PDI	CA1	CA2	CA3	CA4
864	25.2	820	2	0	2	70	45	25.77	26.20	27.06	29.06
865	25.2	820	1	0	2	72	88	25.91	26.34	27.20	29.20
866	28.1	915	1	0	2	76	96	28.24	29.10	29.96	31.96
867	27.0	870	2	0	2	84	108	27.71	28.14	29.00	31.00
891	26.5	625	1	0	3	88	84	26.93	27.50	28.50	NA
892	26.5	890	1	0	3	90	88	26.93	27.50	28.50	30.50
900	25.3	800	2	0	3	110	103	25.87	26.87	27.87	29.16
903	24.4	730	1	0	3	112	103	24.83	25.40	26.40	28.40

Appendix B

Cerebral Function Monitoring and Amplitude-Integrated EEG

The CFM, which stands for cerebral function monitor, is a hardware solution for long-term EEG monitoring that was proposed in the late 1960s [55]. The machine was designed to condense information on an EEG into a shorter recording, before the time of electronic displays. In essence, it condenses the EEG data into what the medical community refers to as aEEG - Amplitude-Integrated EEG. This is basically the envelope of the EEG signal. The raw EEG signal is amplified, and filtered with a bandpass filter to filter out the frequencies below 2Hz and above 15Hz. A high frequency boost is present in the bandpass filter, as higher frequency contents have relatively lower amplitude and this bias is added to compensate for this. The filtered signal is then processed with a semi-log amplitude compression, which compresses the signal logarithmically above 10μV. This is done to accommodate the wide range of amplitudes present in the EEG signal. The compressed signal is then rectified and smoothed to give the aEEG signal. The algorithm used is shown as a block diagram in Figure B.1 [55].

Although the algorithm was proposed in the early 60s, it still remains one of the main ways of displaying the EEG signal. This is due to the knowledge accumulated over the years on how to interpret aEEG, and it remains one of the most efficient ways to display an overview of long EEG recordings [9]. For this reason, most electronic monitoring systems have digitally implemented the process, and provide the aEEG signal as one of the display options.

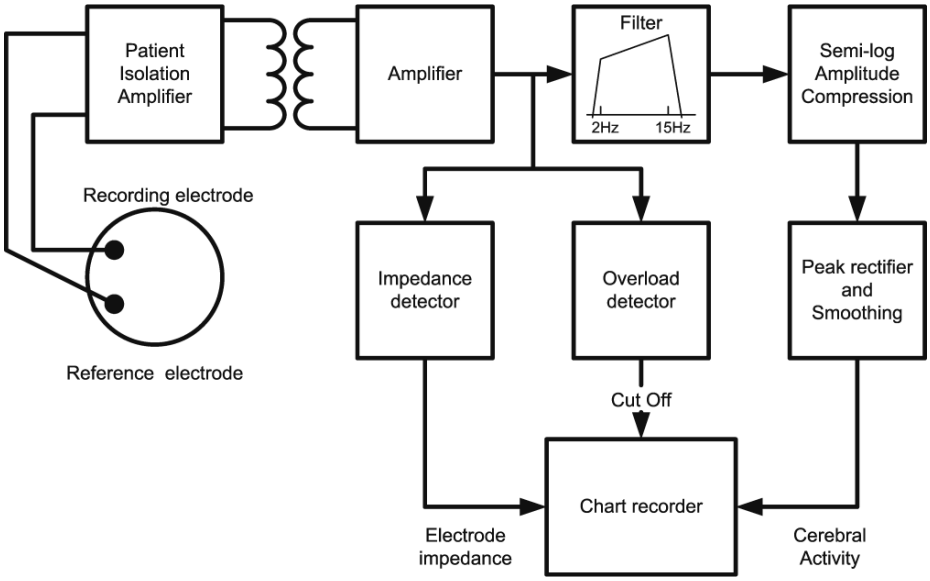

Figure B.1: Block diagram showing the algorithm for aEEG.

B.1 Matlab Code for TFD Calculation: with Time-Lag Kernels

Included here is the Matlab function used for time-frequency distribution calculation. The implementation is based on the book chapter "Computation of Discrete Quadratic TFDs" in [40].

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%function [tfd, ambi_kern, ambi_func, tf_kern] = tlkern(s, N, tr, tf,
%                            lf, af, ap, tf_length, lf_length, bypass)
%
%Function to calculate different TFD with kernel defined as inputs
%
%INPUTS:
%s : signal vector
%N : assumed period.
%tr: time resolution
%tf: time factor (i.e.: g1[n], options are 'delta' (delta function),
     '1' (unity function), and options from the MATLAB function 'window')
%lf: lag factor (i.e.: g2[m], options are 'delta' (delta function),
     '1' (unity function), and options from the MATLAB function 'window')
%af: auxiliary factor (i.e.: g3[n,m], options are 'mb' (Modified-B), '
     cw' (Choi-William), 'rihaczek' (Rihaczek), 'zam' (ZAM), '1' (unity),
     'delta' (delta function at (0, 0)) and 'delta_n' (delta function in
     time direction, equivalent to Wigner-Ville))
%ap: auxiliary parameter (e.g.: in MBD case: beta)
%tf_length: length of time factor window
%lf_length: length of lag factor window
%bypass: variable to bypass the conversion into analytic signal
%        (1 = bypass conversion) defualt = 0
%
%OUTPUT:
%tfd(1:Mpad+1, 1:Nsel): the TFD
%ambi_kern: the ambiguity kernel
%ambi_func: the ambiguity function of s
```

```
%tf_kern: time-frequency kernel
%
%Algorithm as included in Time Frequency Signal Analysis and
%Processing: A Comprehensive Reference (B. Boashash ed.), 2003
%Elsevier, Ch. 6.5
%
%Adapted by Lisa Wong (Jan 2005)
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
function [tfd, ambi_kern, ambi_func, tf_kern] = tlkern(s, N, tr,  ...
tf, lf, af, ap, tf_length, lf_length, bypass)

if N ==0
    N = 2*length(s); %to avoid wrap-around effect for non-periodic
                     %signals
else
    N = length(s); %Assume period = length of signal (for analytic
      %signals)
end

M = N; %default support length of kernel.

%begin algo init
Mpad = 2^ceil(log(M)/log(2));
Ncut = min(N, length(s));       % duration of TF plot
Nsel = ceil(Ncut/tr);           % no. traces in TF plot

Moff = fix(M/2);
%limit frequency resolution (for speed)
if Mpad > 1024
    Mpad = 1024;
    Moff = fix(Mpad/2);
end

if nargin <10
    bypass = 0;
```

```
end

Noff = fix(N/2);

if ~bypass
    %make analytic signal for calculation
    z = fft(real(s), N);            % s truncated or padded
    z(2:N-Noff, :) = 2*z(2:N-Noff, :);    % positive frequencies
    z(Noff+2:N, :) = 0;                % negative frequencies
    z = ifft(z);                    % analytic function
else
    z = s;
end



%time dependent kernel
g1(1:N) = 0;
if strcmpi(tf,'delta')     %delta function
    g1(1) = 1;
elseif strcmpi(tf,'1')
    g1(1:N) = 1;
else
    if nargin > 7
        glength = odd(tf_length);
    else
        glength = odd(N/10);
    end
    gtemp = window(glength,tf);
    Lg = floor(length(gtemp)/2);
    gindex = -Lg:Lg;
    g1(1+rem(N+gindex, N)) = gtemp;
    clear gtemp;
end

%lag dependent kernel
```

```
g2(1:Mpad) = 0;
if strcmpi(lf, 'delta')
    g2(1) = 1;
elseif strcmpi(lf, '1')
    g2(1:Mpad) = 1;
else
    if nargin > 8
        hlength= odd(lf_length);
    else
        hlength= odd(N/4);
    end
    htemp = window(hlength, lf);
    Lh = floor(length(htemp)/2);
    hindex = -Lh : Lh;
    g2( 1 + rem(Mpad+hindex, Mpad)) = htemp;
    clear htemp;
end

%auxiliary factor in kernel
g3(1:N,1:Mpad) = 0;
if strcmpi(af, 'mb')  %modified b-distribution
    temp(1:N) = 0;
    for n_vec = -Noff:Noff
        temp(1+rem(N+n_vec,N)) = (cosh(n_vec))^(-2*ap);
    end
    temp = temp/sum(temp);      %normalise
    for m_vec = -Moff:Moff
        g3(:,1+rem(Mpad+m_vec, Mpad)) = temp.';
    end

elseif strcmpi(af, 'cw') %choi-william
    pi_sigma = pi * ap;
    ntemp = -Noff:Noff;
    n_vec = [ntemp(Noff+1:end), ntemp(2:Noff)];
    for m_vec = -Moff:Moff
```

```
        if m_vec == 0
            g3(1,1+rem(Mpad+m_vec, Mpad)) = 1;
        else
            m_term = 4 * m_vec ^ 2;
            g3(:,1+rem(Mpad+m_vec, Mpad)) = sqrt(pi_sigma/ ...
            (m_term + pi_sigma)) .* exp((-pi * pi_sigma * n_vec.^2) ...
            /(m_term + pi_sigma))';
        end
    end


elseif strcmpi(af, 'rihaczek')
    g3 = eye(size(g3));

elseif strcmpi(af, 'bd')
    ntemp = -Noff:Noff;
    n_vec = [ntemp(Noff+1:end), ntemp(2:Noff)];

    for m_vec = -Moff:Moff
        g3(:, 1+rem(Mpad+m_vec, Mpad)) = ((abs(m_vec)./ ...
                                ((cosh(n_vec)).^2)).^ap)';
    end

elseif strcmpi(af, 'zam')
    ntemp = -Noff:Noff;
     n_vec = [ntemp(Noff+1:end), ntemp(2:Noff)]';

    for m_vec =  -Moff:Moff
        win = abs(ap*n_vec) <= abs(2*m_vec);
        g3(:, 1+rem(Mpad+m_vec, Mpad)) = win .* ...
                        g2(1+rem(Mpad+m_vec, Mpad));
    end
    g2(1:Mpad) = 1;
```

```
elseif strcmpi(af, '1')
    g3(1:N,1:Mpad) = 1;
elseif strcmpi(af,'delta')
    g3(1,1) = 1;
elseif strcmpi(af,'delta_n')
    g3(1,:) = 1;
end


%Calculating IAF K(1:N, 1:Mpad)
for n_vec = 1:N
    taumax = min([N-n_vec, n_vec-1, round(Mpad/2)-1]);
    m_vec = -taumax:taumax;
    if size(z, 2) == 1
        K(n_vec,1+rem(Mpad+m_vec, Mpad)) = (z(n_vec-m_vec).* ...
                                           conj(z(n_vec+m_vec)))';
    else
        K(n_vec,1+rem(Mpad+m_vec, Mpad)) = (z(n_vec-m_vec, 1).* ...
         conj(z(n_vec+m_vec, 2)))';
    end
    tau = round(Mpad/2);
    if (n_vec<= N-tau) & (n_vec >=tau+1)
        if size(z, 2) == 1
            K(n_vec, tau+1) = 0.5 * (z(n_vec + tau)*conj(z(n_vec-tau))...
                                    + z(n_vec - tau)*conj(z(n_vec+tau)));
        else
            K(n_vec, tau+1) = 0.5 * (z(n_vec + tau, 1)* ...
            conj(z(n_vec-tau, 2)) + z(n_vec - tau, 1)*conj(z(n_vec+tau, 2)));
        end

    end
end
```

```
if nargout >2 %find ambiguity function
    ambi_func = ifft(K);
    ambi_func = [ambi_func(Noff+2:N, :) ; ambi_func(1:N-Noff, :)];
    ambi_func = [ambi_func(:, Moff+2:Mpad) ambi_func(:, 1:Mpad-Moff)];
    ambi_func = fliplr(ambi_func);
end

%Convolution in time: R[n,m] = K[n,m]*_n G[n,m]
%use FFT instead to make things quicker (duality)
for m_vec = -Moff:Moff
    mcorr = 1+rem(Mpad+m_vec, Mpad);
    K(:,mcorr) = ifft(fft(K(:,mcorr)).*fft(g3(:,mcorr)).*...
     fft(g1.'*g2(mcorr)));
    temp_kern(:,mcorr) = ifft(fft(g3(:,mcorr)).*fft(g1.'*g2(mcorr)));
end


%find kernels
if nargout >1
    temp_akern = ifft(temp_kern);

    ambi_kern = [temp_akern(Noff+2:N, :) ; temp_akern(1:N-Noff, :)];
    ambi_kern = [ambi_kern(:, Moff+2:Mpad) ambi_kern(:, 1:Mpad-Moff)]';

    if nargout >3
        temp_tfkern = (fft(temp_kern'))';
        tf_kern = [temp_tfkern(Noff+2:N, :) ; temp_tfkern(1:N-Noff, :)];
        tf_kern = [tf_kern(:, Moff+2:Mpad) tf_kern(:, 1:Mpad-Moff)]';
    end
end

%apply time resolution
for nsel = 1:Nsel
    n_vec = 1+tr*(nsel-1);
```

```
    r(:,nsel) = K(n_vec,:).';
end

r = fft(r);

%final normalisation

tfd = [real(r);real(r(1,:))].*(Ncut/Nsel/Mpad);
```