AN EXPERT SYSTEM FOR EEG MONITORING IN THE PEDIATRIC ICU

ANITHA K. PASUPATHY

Department of Electrical Engineering
McGill University
Montreal, Canada.

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of
Master of Engineering

August 1994

ISBN 0-612-05469-1

ABSTRACT

A knowledge-based expert system was developed to assess the level of abnormality in the brain electrical activity of pediatric patients monitored in the intensive care unit. Six hours of an 8-channel EEG record serves as the input to the monitoring device based on which the brain activity is classified as being normal, mildly abnormal, moderately abnormal or severely abnormal.

Spectral band activity is computed for each channel for every 30-second epoch. Artifact rejection is accomplished by a median filter with a hard-limiter thresholder. Quantitative variables reflecting possible abnormality : a measure of amplitude depression, a measure of assymmetry, a measure of anterio-posterior differentiation and a measure of EEG variability over time are extracted from each EEG record. Statistical distributions of these measures are established for a control "normal" population of about ten patients so classified by a neurologist on visual interpretation. New EEGs to be analysed are statistically compared with the control population and a probability measures of normality for the various measures are determined. The expert system learns from prior examples of classification done by the neurologist by a technique of inductive machine learning. The monitor is trained and tested using sixty examples using the rotation method of error estimation.

The monitor had a tendency to classify the EEGs with a higher level of abnormality than the expert. Possible reasons and potential solutions are discussed.

RESUME

On a développé un système expert pour déterminer le niveau d'anormalité de l'activité électrique du cerveau chez des enfants placés sous surveillance à l'unité de soins intensifs. Une section de 6 heures d'un EEG à 8 voies est analysée et classée par le système dans les catégories suivantes: normal, anormalité mineure, anormalité moyenne et anormalité sévère.

On commence par calculer les bandes d'activité spectrale pour chaque section de 30 secondes. La réjection d'artéfacts se fait ensuite à l'aide d'un filtre médian et d'un seuil fixe. On calcule à partir de l'enregistrement complet les variables suivantes, reflétant divers aspects de l'anormalité électroencéphalographique: une mesure de dépression d'amplitude, une mesure d'asymétrie, une mesure de différentiation antéro-postérieure, et une mesure de fluctuations temporelles. On obtient ensuite une distribution statistique de ces variables pour une population de sujets contrôles dont l'EEG a été considéré totalement normal par interprétation visuelle. Les EEGs à analyser sont alors comparés à cette population, permettant d'obtenir une mesure de probabilité de normalité pour chacune des variables. Le système expert apprend par la méthode de logique inductive la classification de chaque EEG faite par le neurologue. L'apprentissage et l'évaluation des erreurs sont faits par la méthode de rotation avec 60 exemples.

Le moniteur a tendance à classer les EEGs dans un niveau d'anormalité plus élevé que le neurologue. On discute les raisons possibles et les solutions à envisager.

ACKNOWLEDGEMENTS

I would like to extend my sincere thanks to my supervisor, Dr. Gotman, Montreal Neurological Institute for his motivation and continuous guidance. His invaluable comments and suggestions inspired me to perform well as a researcher and helped me improve the quality of my thesis.

I would like to thank Dr. Rosenblatt, Montreal Children's Hospital, for his help while I tried to understand children's EEG. His efforts to improve the performance of the monitor by painstaking re-reading the EEGs is very much appreciated.

I appreciate the help extended by Dr. Lowther, McGill University and Dr. Gottesman, Montreal Children's Hospital, for the many interesting and thought-provoking discussions. I wish to thank Dr. Amsel, McGill University, for the suggestions she made regarding statistical analysis.

The help offered by Elie Saade, computer programmer, Montreal Children's Hospital, with programming and other computer-related problems is deeply appreciated. I also thank Elyse Chevrier, technician, Montreal Children's Hospital, for teaching me the fundamentals of children's EEG and also carefully documenting all the data collected for this project.

Finally I wish to acknowledge the support provided by the Medical Research Council of Canada and Stellate Systems through the university - industry program of the Medical Research Council. (grant MRC -UI - 11514).

CONTENTS

ABSTRACT.. ii

RESUME... iii

ACKNOWLEDGEMENTS.. iv

1. INTRODUCTION 1

1.1 Neurological Monitoring.. 1

1.2 The Electroencephalogram.. 3

1.2.1 Generation and Recording of the EEG.. 3

1.2.2 Why the EEG for monitoring... 6

1.2.3 Quantitative EEG 9

1.3 Cardiac Surgery and Brain Injury 10

1.4 Project Definition 12

1.5 Literature Review... 14

2. THE ABNORMAL EEG.. 18

3. DATA PRE-PROCESSING 23

3.1 DATA Acquisition 23

3.2 Artifact Management 29

3.2.1 Artifacts.. 29

3.2.2. Artifact Rejection... 38

4. FEATURE SELECTION 49

4.1 Measure of Amplitude.. 50

4.2 Measure of Left / Right Symmetry 52

4.3 MEASURE OF FRONT / BACK DIFFERENTIATION 54

4.4 MEASURE FOR VARIABILITY 58

5. STATISTICAL ANALYSIS 62

5.1 POPULATION DISTRIBUTIONS 64

5.1.1 Amplitude Normality 64

5.1.2 Symmetry 68

5.1.3 Front / Back Differentiation 70

5.1.4 Variability 70

5.2 STATISTICAL TESTS 72

5.2.1 T - Statistic 74

5.2.2 Assessment of Variability 80

5.3 DATA REDUCTION 82

6. KNOWLEDGE-BASED EXPERT SYSTEM 85

6.1. MACHINE LEARNING FROM EXAMPLES 86

6.2. CONCEPT LEARNING 87

6.2.1. Concept Learning by Generalization 88

6.3. NEURAL NETWORKS 91

6.4 KNOWLEDGE-BASED EXPERT SYSTEM FOR EEG CLASSIFCATION 91

6.4.1 Neural Network or Induction by Generalization? 91

6.4.2 Defining the boundary 93

6.4.3. Dealing with bad instances 96

7. RESULTS AND CONCLUSION 98

7.1 RESULTS OF STATISTICAL ANALYSIS 98

7.2 TRAINING AND TESTING DATA 106

7.3 RESULTS OF THE CLASSIFIER 108

7.4 DISCUSSION 111

8. REFERENCES 117

1. INTRODUCTION

1.1 Neurological Monitoring

The purpose of any monitoring facility in the Intensive Care Unit (ICU) is to equip the physician with an additional pair of eyes; to enhance his/her powers of observation to detect abnormalities at a reversible stage so that timely and medically sound decisions can be made. In the present day ICUs, innumerable catheters, transducers, digital read-outs and alarms keep vigil on the heart, lungs and kidneys. The brain however, is monitored only by clinical observations and by measurement of intra-cranial pressure (ICP).

Invasive ICP monitoring devices provide the most accurate and reliable information for clinical determination of ICP. Non-invasive methods of determining ICP are promising but none are clinically useful at present. Invasiveness is an unfavorable quality for any monitoring device. Further, irrespective of the method of ICP monitoring used, the user requires a level of expertise for optimal interpretation of the data it produces. Spurious ICP measurements that go unrecognized as erroneous may lead to management decisions with potentially catastrophic consequences (Miller et. al., 1986).

Monitoring the Central Nervous System level of functioning is a difficult task as it is not feasible to have a neurologist at all times, doing serial neurologic examinations. As a result, neurologic examinations are often delegated to intensive care nurses. Even when performed conscientiously, serial neurologic examinations are discontinuous and subjective. These wait for a clinical manifestation of a functional deterioration. As a result, often times they are unable to anticipate capricious clinical deterioration and identify them only after they occur. This defeats the purpose of a monitor itself. Also, clinical observations are very limited in a comatose or unresponsive patient, or an artificially paralyzed patient. A single or even repeated electrophysiologic study may suffer from the same problems as does the neurologic examination - it is a discrete, brief sample of data that may not reflect the patients overall condition. Particularly, it would miss the very gradual changes that may be taking place

A successful ICU monitoring system should meet the following criteria:

- be more sensitive and specific than clinical observations.
- be non-invasive.
- be easily operated and interpreted by non-experts.
- be usable at the patients bedside
- not interfere with medical or nursing care of the patients.

Examples of successful non-neurologic monitoring units are the bedside electrocardiographic (EKG) and transcutaneous pulse oximetry.

Many neurointensivists believe that the electroencephalogram (EEG) can become an integral part of monitoring in the ICU (Emmerson and Chiappa, 1988). The EEG is very sensitive to change in physiological state and its value as a prognostic indicator of functional recovery has been demonstrated (Arroyo et. al., 1993). The technique itself is non-invasive and technological advances have made possible the collection, storage and analysis of continuous EEG. However, raw EEG recorded continuously over several hours generates cumbersome amounts of data and its complexity discourages interpretation by non-experts.

To overcome this shortcoming and to make it easily operable by non-experts, several attempts have been made at compressing and simplifying the data with definite benefits. Recent studies have examined the continuous EEG as a tool for making on-line clinical management decisions (Jordon 1990, 1992).

1.2 The Electroencephalogram

1.2.1 Generation and Recording of the EEG

Electroencephalography involves the recording and analysis of the electrical signals generated by the brain. The electrical activity of the brain consists of ionic currents generated by biochemical sources at the cellular levels. These ionic currents cause electric and magnetic fields that can be measured in the brain and surrounding tissues.

The EEG is recordable from the scalp surface after being picked up by metal electrodes and conductive jelly. The arrangement of these electrodes on the surface of the scalp is done based on the international 10-20 System (Jasper 1958). EEG is recorded as a potential difference between pairs of electrodes and each such pair is referred to as a "channel". Figure 1.1 illustrates a 16-channel EEG. The amplitude of the EEG typically ranges from 10 to 100 μV and is amplified about ten thousand times for recording and display. The EEG is corrupted by artifacts caused by various sources such as patient movement, power supply interference and poor electrode contact. Amplification of these signals by ten thousand times causes immense artifactual distortion in the EEG.

Figure 1.1 - 20 seconds of a 16-channel EEG.

The combination of electrodes examined at a particular point in time is referred to as a "montage". For instance, potential difference between successive pairs of electrodes in each hemisphere may be recorded as shown in Figure 1.2. This is referred to as a "bipolar anterior posterior montage".

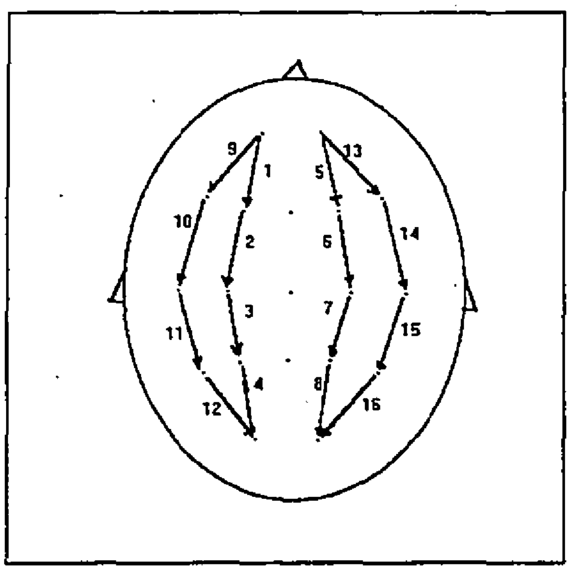

Figure 1.2

Continuous analog EEG can be displayed using paper write-out or oscilloscopic display. Computer-based digital EEG systems are now in use. In addition to efficient means of data storage and transmission, they also lend themselves to subsequent data processing.

The EEG is used in the evaluation of infectious diseases of the nervous system, head trauma, cerebral vascular accidents, epilepsy and brain tumors. It can also contribute as an indicator of brain function in metabolic disorders and in the evaluation of organic causes of psychiatric problems and behavioral and adjustment problems of children. More recently, EEG has come to play a major role in the evaluation of the cerebral death of donors where organ transplants are considered. EEG can also be recorded from the cortical surface or from depth probes. The former is called electrocorticogram while the latter is called depth electrogram.

1.2.2 Why the EEG for monitoring

A few of the more important reasons for using the EEG as a monitor of brain activity are discussed below.

1. *EEG reflects cerebral neuronal function.* The EEG represents the temporal and spatial summation of potentials, at the junctions of nerve cells called neurons. These potentials reflect the underlying state of cerebral metabolism which in turn depends on multiple factors including synthesis of enzymes and energy (Siegel et. al., 1989). Thus, the EEG is a composite reflection of complicated intracellular activity and inter-neuronal communication. A disturbance in one or more of these components will produce a disturbance in the EEG. This multi-layered system makes the EEG a highly sensitive, although non-specific, indicator of cerebral dysfunction.

2. *The EEG is sensitive to ischemia (reduced cerebral blood flow) and hypoxia (reduced level of oxygen to brain) -- the most common causes of brain injury.* Predictable EEG changes occur with cerebral ischemia. The EEG is mainly generated by the pyramidal neurons of the cortex, called pyramidal because of their shape. Following hypoxic-ischemic injury the pyramidal neurons typically demonstrate severe neuronal dropout, suffering a large drop in neuronal population. The relatively selective vulnerability of these cells makes the EEG particularly sensitive to these common insults. The sensitivity

of the EEG to ischemic injury has been demonstrated by Jordon and Stringer (1991). In their studies they observed a correlation between amplitude of the background EEG activity and the volume and severity of the ischemic damage.

3. *EEG correlates with cerebral topography.* The international 10-20 System used for electrode placement establishes a consistent relationship between the electrode placement on the surface of the scalp and the underlying cerebral topography (Homan et. al., 1987). Therefore, inferences maybe drawn about disease localization based on the EEG abnormalities detected at the scalp. This was a more important attribute before the introduction of the tools capable of more specific disease localization such as computed tomography (CT) or magnetic resonance imaging (MRI). However for patients in the ICU transport for imaging studies could be hazardous. The EEG has the advantage of being able to serve as a bedside monitor and aid in decision making about disease localization.

4. *EEG detects neuronal dysfunction at a reversible stage.* The EEG deteriorates before cell membrane failure. Astrup et al. (1981) have demonstrated that EEG abnormalities due to diminished blood flow set-in much before cell-death or energy failure occurs. This implies the existence of a "therapeutic window" following EEG abnormalities. Gross et al. (1981) and Wood et al. (1984) have reported that EEG abnormalities reversed as cerebral blood flow increased in patients with cerebral ischemia thus demonstrating this theory.

5. *EEG is the best available method for detecting seizure activity*. A seizure is the result of occasional , excessive and disorderly electrical discharges of the gray matter which may be detected by EEG monitoring.. Clinically, it could manifest itself as sudden, brief attacks of loss of consciousness, motor activity, sensory phenomena or inappropriate behavior. A significant incidence of acute seizures among ICU patients, subsequent to head injury, spontaneous intracranial hemorrhages and ischemic strokes has been reported by Engel (1989). A systematic study of seizures among patients in the ICU has documented a high incidence of non-convulsive seizures (Jordan, 1992). The clinical features of these seizures, unlike their convulsive counterparts, are subtle, ambiguous or absent and hence diagnosis of these is quite difficult without EEG monitoring. EEG monitoring has aided seizure management in the ICU in two ways. Non-convulsive seizures with very subtle or no clinical accompaniments are detected earlier thereby allowing the initiation of timely treatment (Maynard and Jenkinson, 1984). ICU patients exhibit a wide variety of involuntary and semi-purposeful movements which could be mistaken to be clinical accompaniments of a seizure. EEG monitoring helps distinguish semi-purposeful jerks, spasms, head deviations, etc. from true clinical manifestations of seizures.

In addition to the neurobiologic reasons discussed above, the technique of EEG itself is non-invasive and an EEG monitoring system can operate at the patient's bedside without any interference to examination or patient care. However, interpretation of raw EEG data requires skilled personnel who may not be available at all times in the ICU.

Quantitative EEG discussed in the following section may be able to overcome this shortcoming.

1.2.3 Quantitative EEG

A neurophysiological EEG monitor should record brain activity of the patient in the ICU over several hours continuously in order to provide authentic diagnostic information. Thorough review and interpretation of the EEG requires the presence of a neurologist on-site throughout the period of recording. This is quite unreasonable. In practice, there is a considerable time lag between recording of the EEG and actual interpretation. Also, visual interpretation of a 24 hour EEG recording is quite tedious and time-consuming.

Quantitative EEG (QEEG) relies on the transformation of digitized EEG signals into mathematically derived parameters thereby performing a frequency analysis by computing the Fast Fourier Transform, a period analysis for measuring half or full waves or an amplitude analysis for determining the average amplitude, its variance skewness and kurtosis (Pronk, 1987). These parameters characterize the EEG thus analyzed. This information can be further interpreted by statistical analyses based on a select population considered to be "normative". It can also be displayed as topographic maps or graphs in a form intelligible to the non-expert. Pattern recognition.techniques may also be applied to the parameters derived to classify them as normal or abnormal recordings.

Several factors in addition to the ones mentioned above, favor QEEG analysis. Bickford (1973) has demonstrated that experienced electroencephalographers make considerable errors in estimating amplitude values and consistently read higher than the computer estimate. In the presence of artifacts the human estimates are further degraded. QEEG using computer-assisted frequency analysis has been reported to be more sensitive than visual interpretation in several situations. QEEG monitoring systems can compress data and thereby identify modulating backgrounds, intervals of physiologic sleep and gradual shifts in dominant frequency. Long-term EEG trends of prognostic significance such as those mentioned above may be missed by visual analysis. Also, the earliest manifestation of cerebral ischemia subsequent to carotid surgery is a 5 to 15% drop in amplitude (Chiappa et. al., 1979). This goes unnoticed during visual interpretation but is reported by QEEG analysis.

It would be quite useful to have a monitoring unit that would not only record EEG but also interpret it on-line. Such a system would then extract features from the EEG and based on those characteristics, raise an alarm at the appearance of abnormal EEG activity. At this time a neurologist may be called in to confirm and act upon the finding of the monitoring system. One place where such a system would make a definite impact on clinical decision-making is the pediatric ICU.

1.3 Cardiac Surgery and Brain Injury

Innumerable children are born each year with congenital heart disease serious enough to require surgery early in life. The use of deep hypothermic Cardiopulmonary Bypass (CPB) with or without deep hypothermic Circulatory Arrest (CA) has improved operating conditions for pediatric cardiac surgery. This has resulted in improved survival and reduced cardiac morbidity.

Hypothermic CA has been a widely used technique since its introduction in the early 1960's. CA creates an operative field free of perfusion cannulae and blood and this proves to be a definitive advantage. The use of this technique is based on the premise that there is a "safe" duration of total circulatory arrest, which bears an inverse relation to body temperature; the organ with the shortest "safe" circulatory arrest time is the brain.

The perioperative period is an important time for the occurrence of brain injury serious enough to be followed by neurologic sequelae. Results of a study at Duke University Medical Center indicated that intracellular brain oxygenation decreases significantly during circulatory arrest and remains impaired after rewarming and CPB, despite normalization of oxygen availability (Greeley et. al., 1989). Deep hypothermic CA seems to be a factor in the delayed recovery of cerebral blood flow and metabolism in patients. Some patients exhibit a normal response to low cerebral blood flow following surgery, i.e. they demonstrate an increased oxygen extraction from the blood. However, the cerebral oxygen metabolism in some patients is stunned and unable to exert a protective response of increased oxygen extraction. It is therefore likely that low cardiac output and pressure-passive cerebral blood flow potentiate brain ischemia after CPB and surgery in some patients. Ischemia - a state of diminished blood supply to the brain, is the initiating event for the development of several brain lesions.

As the overall outcome of corrective surgery for congenital heart disease has improved, the focus is now on neurologic dysfunction due to surgery. The incidence of significant neurologic deficits following cardiac surgery in infants varies approximately between 5% and 25% depending upon the sophistication of the follow-up measures.

Extensive research is going on, both to improve the existing methods of CA and bypass and also to try and reverse ischemic damage when it occurs.

What the pediatric ICU could use is a neurophysiologic monitor which could interpret the EEG and raise an alarm at the appearance of abnormal activity. As discussed in section 1.2.2, the EEG is very sensitive to ischemia and EEG abnormalities set in prior to neuronal dysfunction. Such a monitor would then detect abnormalities at a reversible stage and provide a scope for therapeutic assistance. This in essence, is the aim of this project and is further discussed below.

1.4 Project Definition

The aim of the project is to build an automated Neurophysiological Monitor for the pediatric ICU. The primary utility of the system is to serve as a bedside diagnostic tool for pediatric cardiac patients subsequent to surgery. This may also be used to monitor other ICU patients as the system basically detects abnormal EEG patterns. The EEG of cardiac patients in the ICU is recorded for about 20 hours starting an hour or so after surgery. The monitor being built is an off-line device which accepts several hours of raw EEG data as input, performs quantitative analysis and classifies it as normal or abnormal based on its characteristics. Such a monitor aims at mimicking the neurologist, the expert. The features used by the monitor for interpretation of an EEG will then be the quantitative equivalents of the qualitative features used by the expert himself. Knowledge is acquired in collaboration with the expert about the features that are crucial in differentiating a normal from an abnormal EEG. This is discussed in Chapter 2.

A continuous EEG recorded over long hours is subject to various kinds of artifacts. Rejection of artifacts of physiologic and environmental origin is quite essential for accurate interpretation of the EEG. Once artifacts are rejected, the EEG is ready for quantification. Chapter 3 discusses the EEG data acquisition system at the Montreal Children's Hospital, the source and nature of artifacts that are encountered during a long-term recording in the ICU and the design of filters for artifact rejection.

Knowledge Engineering involves the design and development of quantitative parameters equivalent to the qualitative features crucial in classifying an EEG. Chapter 4 discusses the various mathematically derived parameters extracted from the EEG. EEG is a qualitative tool and no clear-cut quantitative boundaries exist between normal and abnormal recordings. Statistical analyzes of the quantitative features extracted, discussed in Chapter 5, provides information about the normality of the EEG by comparing it with a population considered normative.

A Knowledge-Based Expert System accepts the probability measures associated with each of the parameters extracted and classifies the EEG into one of four categories, normal, mildly abnormal, moderately abnormal, or severely abnormal. Machine Induction principles are used to train the system with prior examples. Chapter 6 discusses this in detail. The performance of the automated monitor is discussed and improvements are suggested in Chapter 7.

1.5 Literature Review

The literature reveals several instances of EEG monitoring in the ICU. EEG changes correlate with regional ischemia during carotid artery endarterectomies (Trojaborg and Boysen, 1973). As a result EEG monitoring is currently in extensive use during carotid surgery (Chiappa et. al., 1979, Cho et. al., 1986) and this has made a definite impact on clinical decision making. EEG monitoring is currently in practice during cardiac surgery as well (Salerno et. al., 1978). Jordon and Stringer (1991) have reported a decisive impact of EEG monitoring on clinical management decisions in 85% of the patients monitored subsequent to cardiac surgery. EEG monitoring of patients in coma has provided clues to the cause of coma and prognostic information (Cant and Shaw, 1984). EEG monitoring has been used for monitoring barbiturate therapy for increased intra cranial pressure (Ropper and Rockoff, 1983).

Much of the results presented in the literature referenced above relies on visual analysis rather than on computer analysis. This is somewhat limited in scope and enormously time-consuming. Computer-assisted analysis has been used by several researchers, however the most sensitive variables to be monitored have not been determined. The tool most commonly used is the Compressed Spectral Array (CSA) devised by Bickford et. al. (1971). Here the frequency spectra of EEG activity are computed and plotted against a vertical time scale for successive epochs. Gross changes can therefore be identified visually. Other display formats such as trend plotting of various EEG and physiologic parameters have also been used.

Pronk (1987) and Prior (1987) have reviewed some EEG features that may be useful in prognostication and clinical decision-making. These include spectral features, period and amplitude features, measures of mobility and complexity, and autoregressive filtering coefficients. Thomas et al. (1985) have reported that, subsequent to cardiopulmonary bypass, the most significant changes concerned amplitude rather than frequency parameters.

Devices designed to provide an automatic EEG monitoring system have been few. Bickford (1950) estimated the depth of barbiturate and ether anesthesia with such a device and automatically regulated the rate of drug administration by changes produced in the EEG. Similar servomonitor mechanisms have been developed for other anesthetic agents as well (Verzeano, 1951). The Cerebral Function Monitor (CFM) described by Prior (1973), is the first device capable of artifact rejection in addition to automated interpretation. The system also has an in-built electrode impedance monitoring system which maybe quite important for long-term EEG monitoring systems.

Bickford et. al. (1971), in their monitoring system for diagnosis of irreversible coma, have concerned themselves with ECG artifact rejection using template subtraction techniques. MacGillivray and Kennedy (1968) describe a monitoring system for patients with grossly disturbed metabolic states. This system, involving analog devices, produces a regular plot at short intervals of the relative proportions of different frequencies in the EEG. They report that clinicians understand these frequency plots and interpret them in order to assist management of comatose patients.

Maulsby (1973) has described a multi-channel EEG analysis tool that displays information extracted from the EEG in an anatomical format. This can be easily comprehended by physicians having little formal training in EEG. The display consists of four figures of the head upon each of which the activity corresponding to one of four frequency bands is plotted as bar graphs.

Maynard and Jenkinson (1984) developed a system called the Cerebral Function Analyzing Monitor (CFAM) an improvement of the CFM described above. This device bandpass filters (2-20 hz) the EEG, performs amplitude rectification and smoothing and then computes, every 2 sec, five amplitude measures and the percentage of activity in nine frequency bands These are displayed as a function of time.

The Vital Signs Monitoring System (VSMS) (Chiappa and Hoch, 1993) is another computer-assisted diagnostic tool. This system plots the trend over several hours, of various EEG and physiological parameters such as peak, median power and spectral edge frequencies, frequency bin activity totals, frequency bin ratios, blood pressure, heart rate and intra-cranial pressure. In addition to data reduction, the plots generated by the CFAM and the VSMS make visual interpretation of EEG trends easier and discernible by the non-expert.

Artifact rejection is quite important for accurate interpretation of the EEG. Most of the systems in the literature that rely on computer analysis do not reject artifact completely. The literature does not show evidence for the presence of systems capable of completely automated EEG interpretation. During feature extraction from the EEG, it would probably be wise to extract quantitative equivalents of the qualitative features used by the neurologist for interpretation. The literature reveals some work in this direction.

This project attempts at complete automation of EEG interpretation. Rejection of all kinds of artifacts is attempted prior to extraction of features that correspond to the qualitative features used by the neurologist.

2. THE ABNORMAL EEG

EEG abnormalities in children can be of three kinds.

* Background abnormalities - Background activity refers to the basic EEG rhythm which are present at all times. Sometimes, major EEG patterns and changes in activity are superimposed on these basic rhythms. EEG background abnormalities correspond to aberrations in the amplitude and frequency composition of these basic rhythms.

* Ictal abnormalities - Ictal abnormalities are those associated with seizures. These appear most often as sharp waves and spikes, hypersynchronous rhythmic activity or as paroxysmal slow wave activity.

* Abnormalities of organization in states and maturation - Composition of EEG activity in children varies considerably with age. Children who exhibit sleep rhythms that are uncharacteristic of their age are said to have abnormalities of organization in sleep states and maturation.

Background abnormalities appear to be most suitable for diagnostic applications especially in long-term prognosis (Lombroso 1985). Studies indicate that background EEG is a good indicator of prognosis subsequent to hypoxic-ischemic injury, if recorded for sufficient duration in all states (Watanabe et.al., 1980). Serial recordings and follow-up studies have disclosed a graded series of background abnormalities from maximally depressed EEG to normal background EEG. Each grade of abnormality correlates closely with the subsequent neurological outcome. The prognostic significance of the background abnormality depends on the time of the recording as well. Studies indicate that recordings within the first 48 hours of the illness serve as the best prognostic indicators (Kayser-Gatchalian and Neundorfer, 1980).

Ictal abnormalities in children could be caused by the involvement of the central nervous system or by metabolic derangement and are either transient or sustained. These provide prognostic information as well. However, background abnormalities are, in general, more reliable prognostic indicators (Lombroso, 1985). For instance, the prognosis of a child with ictal abnormalities accompanied by a sustained background abnormality is worse than that of another, whose ictal abnormalities are accompanied by a transient background abnormality. The third category, abnormalities associated with maturation and organization of states, consists of more subtle deviations of certain bioelectric parameters and their usefulness is not well-established.

An automated EEG neurophysiological monitor should be capable of detecting both background abnormalities and ictal abnormalities in order to obtain maximum prognostic information. An analysis of the background activity involves the study of long-term trends and therefore superimposed bursts of activity which may represent ictal abnormalities are smoothed out. Similarly, a system designed to detect ictal abnormalities pays no attention to long-term trends. Therefore the detection of these two should be handled independently. The intent of this project is to develop a monitoring system capable of identifying background EEG abnormalities.

Background EEG abnormalities manifest themselves in several forms.

1. *Depression* - A depressed EEG is characterized by a 10-50 μV activity with mixed frequencies, persistent through all states, sleep and wakefulness. This background pattern requires caution in interpretation since transient depression could be caused by various reasons and is not an indicator of unfavorable prognosis. Detection of depressed activity is quite crucial since the first manifestation of cerebral ischemia is a 5 to 15% drop in amplitude (Chiappa, 1979).

2. *Inactive Pattern* - This is characterized by cerebral activity below 10 μV almost continuously, unreactive to stimulus. It occurs in disparate clinical conditions and carries quite an unfavorable prognosis.

3. *Burst Suppression* - Periods of inactive background (lower than 10 μV) interrupted by synchronous or asynchronous bursts of activity characterize this abnormal pattern. The intermittent bursts themselves last between 0.5 to 6 sec and contain one or more irregular slow waves with or without sharp transients. Studies show that a burst suppression pattern heralds an unfavorable outcome and has a very high statistical predictability associated with it (Lombroso, 1985).

4. *Interhemispheric Amplitude Asymmetry* - A persistent amplitude asymmetry in background rhythms between corresponding channels of the two hemispheres is considered abnormal. Here again transient or mild asymmetries are of no pathological significance. A persistent voltage asymmetry recognizable in various states could denote a depression in one hemisphere or large amplitude activity in the other, both of which

point to an underlying abnormality. Bricolo et al.(1978), assessed the significance of asymmetry and found that in most cases, the lower amplitude corresponded to the side of the lesion.

5. *Monotonous Pattern* - This consists of an almost invariant diffuse pattern present at all times, poorly reactive to stimuli. Studies show that the second level response to a diminished cerebral blood supply, after generalized depression, is a diffuse monotonous delta pattern from all cerebral regions (Chiappa, 1979). Although the presence of sleep patterns in comatose patients may have prognostic significance, the presence of spontaneous alteration of the EEG is more important. To study spontaneous variations adequately, the EEG is to be recorded over an extended period of time. Bricolo and co-workers (1973) and Bergamasco and associates (1968) have found that a invariant EEG carries worse prognosis than a cycling (alternating) EEG. In another study, approximately 95% of the patients with a slow and monotonous CSA had unfavorable outcomes against only 30% of those with a changing CSA (Bricolo et.al., 1978). Although the cycling EEG patterns only weakly correlate with the clinical state of the patient, they still have significant independent prognostic value (Rumpl et.al., 1979).

6. *Absence of Amplitude Gradient* - The EEG of a normal child shows a sharp decline in voltage from the posterior to the anterior head regions. This is usually accompanied by a marked decrease in low frequencies in the same posterior anterior direction. Substantial clinical evidence correlates an absence of such a gradient with severe neurologic injury (Slater and Torres, 1979).

The monitoring unit built aims at identifying these above-mentioned background abnormalities by extracting quantitative features from the raw EEG. However, prior to this the raw EEG data recorded is to be pre-processed in order to make interpretation easier. This is discussed in chapter 3.

3. DATA PRE-PROCESSING

3.1 DATA Acquisition

At the Montreal Children's Hospital, EEG is recorded from 8-channels by a montage called "Little H" illustrated in Figure 3.1 below. Twelve electrodes including a ground are glued onto the surface of the scalp and the electrode scalp junction is filled in with conductive jelly. EEG recorded from channels 1 and 4 correspond to activity from the frontal head regions of the left and right hemispheres respectively. Similarly, channels 2 and 5 represent right and left central parietal head regions, channels 3 and 6 represent the posterior head regions and 7 and 8 record the central temporal head regions of the left and right hemispheres. This is illustrated in Figure 3.1.

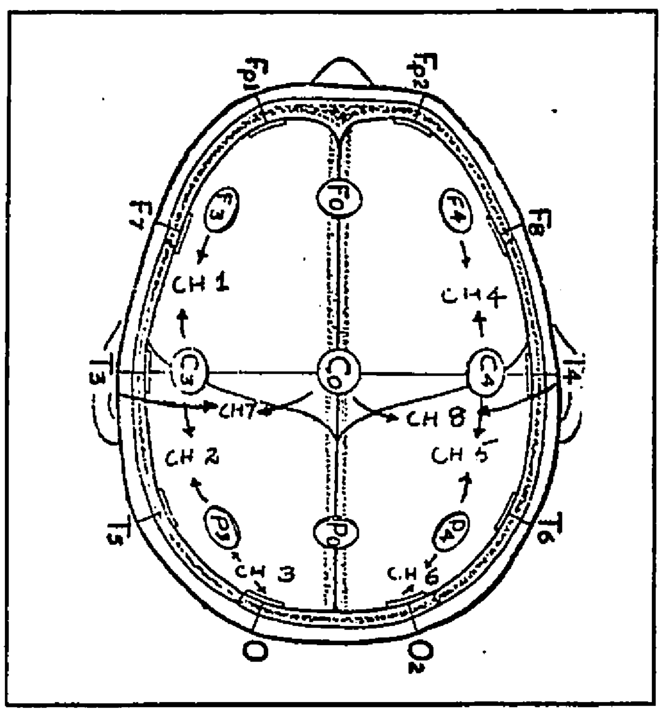

Figure 3.1 - "Little H" montage.

After being picked up by the electrodes, the analog signal is amplified by a factor of 10000 and then digitized at 200 Hz sampling frequency. Prior to digitization, the signal is filtered by a low pass filter (cut-off frequency 30Hz and attenuation 6db/octave) to prevent aliasing and a high pass filter to remove artifacts due to respiration, sweating, etc. EEG is recorded and displayed on a computer monitor using a software called MONITOR. Figure 3.2, on the next page, illustrates 20 seconds of such an 8-channel EEG record.

As explained earlier, analysis of the background activity involves a study of long-term trends. EEG recorded continuously over a period of six hours produces approximately 60 MB of data. The first step toward trend and pattern analysis is data reduction. Further, while detecting background EEG abnormalities, sustained rather than impulsive changes in the EEG are of significance. Therefore, loss of detail as a result of data reduction does not distort EEG interpretation.

Visual interpretation of the EEG involves for most parts, assessment of its frequency composition. The neurologist estimates visually the amounts of EEG activity in the various frequency bands, delta (1-3 Hz), theta (3-7 Hz), alpha (7-14 Hz) and beta (14-30 Hz), compares and correlates them and arrives at a decision about the normality of the EEG. It would therefore be quite useful to transform the entire EEG into the frequency domain. This would provide information about the EEG activity in the various frequencies ranges over a period of time. The EEG activity at various frequencies at different instants of time could then be readily compared and the state of the EEG record could be arrived at. In addition, it would also perform the task of data-reduction and facilitate trend analysis.

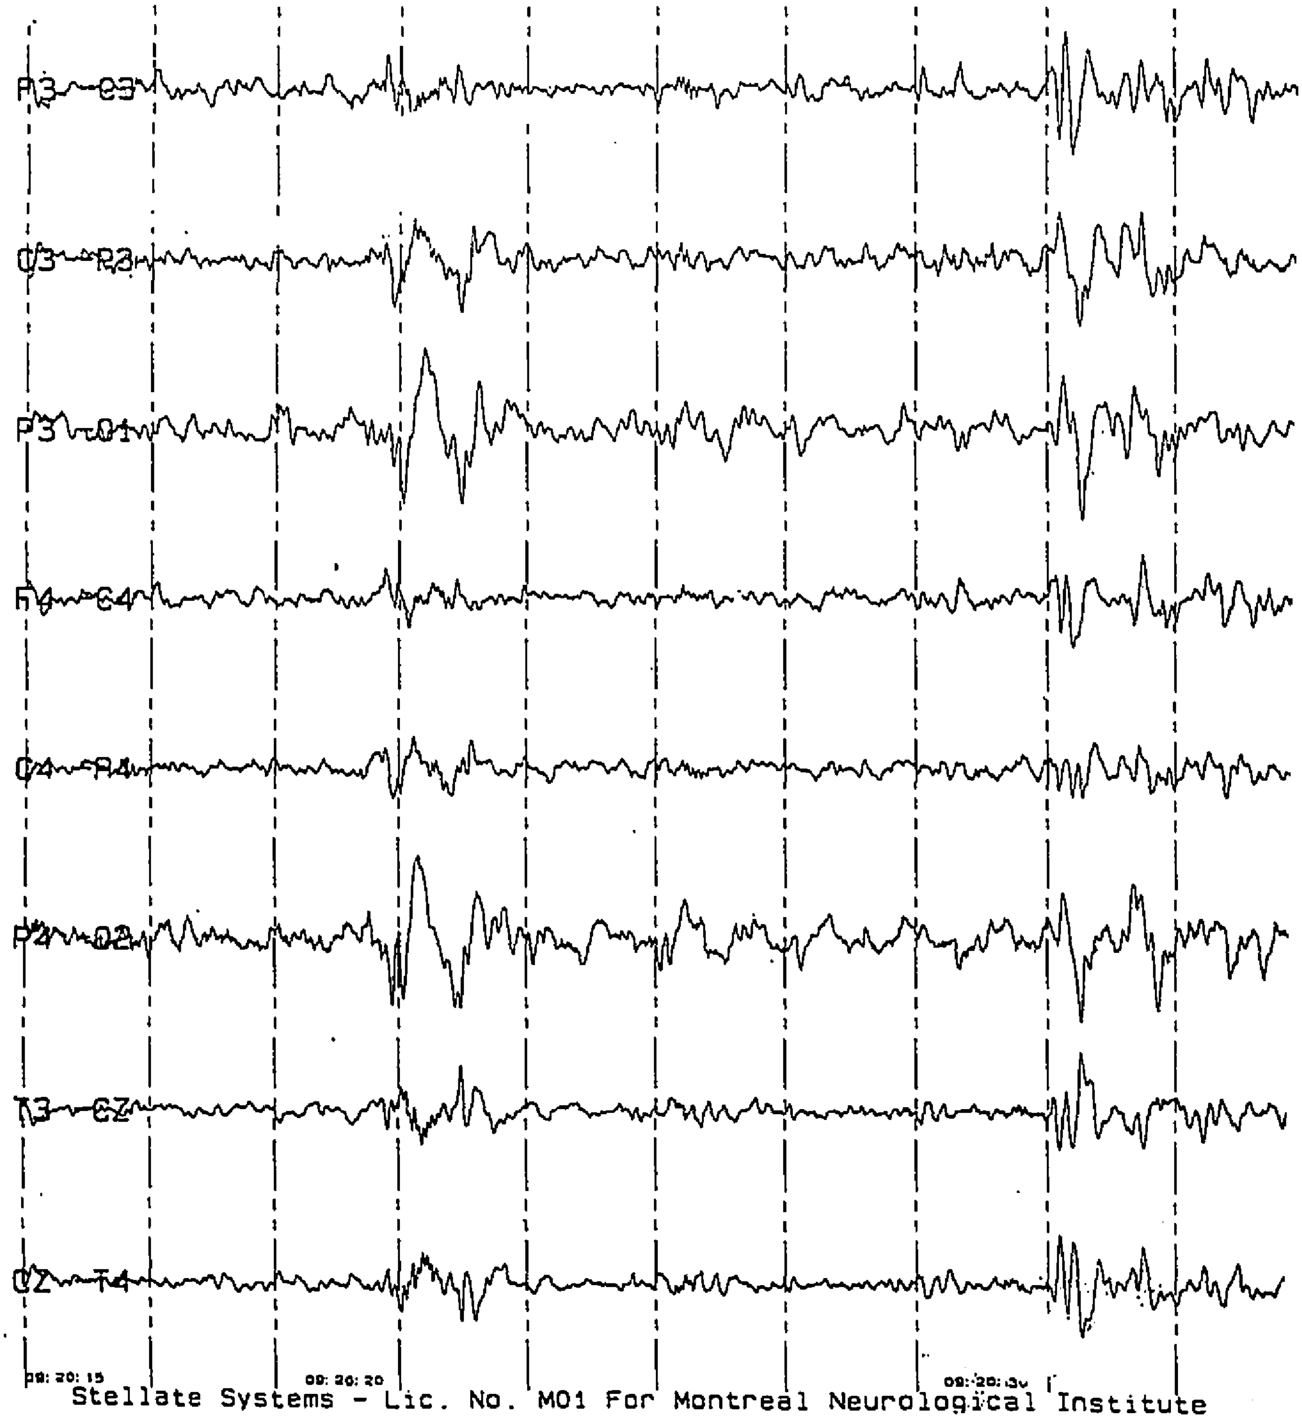

Figure 3.2 - 20 seconds of an 8-channel EEG recorded using the "little H" montage. The names written on the left corner indicate electrode - pair corresponding to that channel of EEG.

The frequency analysis of the EEG record is performed using a software called ECLIPSE. The program divides the entire EEG recording into epochs of 30 second duration. Each such epoch is referred to as a 'page'. The choice of a 30-sec epoch is based on the premise that any rhythm that lasts for a duration less than 30 seconds is impulsive rather than sustained activity. After filtering using a cosine window, Fast Fourier Transforms are computed for every 512 sample points (corresponding to 2.56 seconds) and this amounts to eleven frequency domain distributions per channel per page. The epoch duration is in fact 28.06 seconds and not exactly 30 seconds to facilitate this. The spectral resolution is 0.39Hz. The frequency distributions for each page of a channel is obtained by averaging the 11 distributions corresponding to it. The band activity is then the average of the amplitudes of the activity within the frequency range of the corresponding band. The frequency ranges of the various frequency bands are defined below.

Band Titles	Frequency Range(Hz)
Delta	1.17 to 3.13
Theta	3.52 to 7.03
Alpha	7.42 to 13.28
Total	1.17 to 14.06

Table 3.1 - Definition of EEG frequency bands.

The frequency of genuine cerebral rhythms could be as high as 50 Hz. However, the total band spans only up to 14 Hz. The rationale behind such a narrow band is discussed in the section on Artifact Rejection. Once the "activity" of the EEG record is computed, each channel is associated with four average amplitudes of activity corresponding to the four frequency bands for every page of the record. The activity values for a frequency band of a channel are then plotted against time thus forming a time series as shown in Figure 3.3. Such a plot is called a band array and it represents the trend in activity in a particular frequency range arising from a particular head region. For instance, Figure 3.3 represents the activity over a 12 hour period in the delta, theta, alpha and total bands recorded from the left posterior head region. Activity computation by ECLIPSE reduces 6 hours of raw EEG data from 60 MB to about 500 KB. The next hurdle to overcome prior to trend analysis is artifact management discussed in the following section.

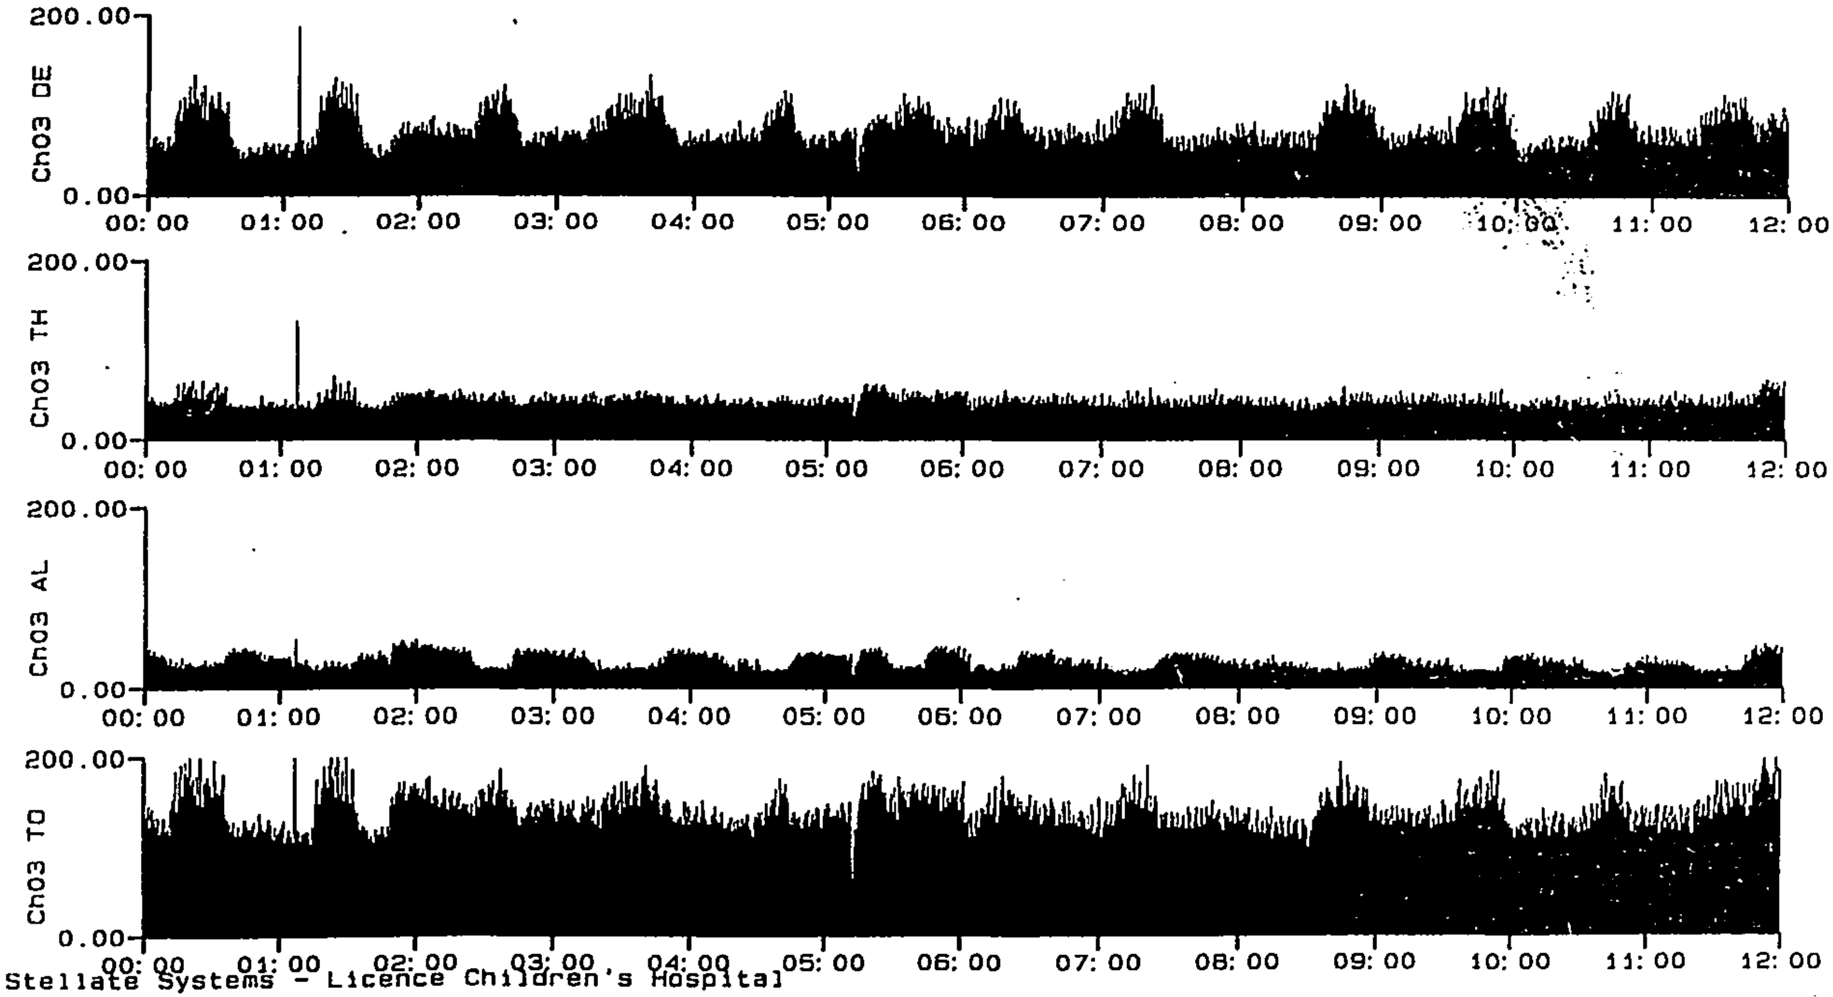

Figure 3.3 - Delta, theta, alpha and total bands of the left posterior channel of a 12 hour EEG record.

3.2 Artifact Management

3.2.1 Artifacts

Artifacts are frequent and often intractable during long-term EEG recordings in the ICU. Here, unlike in the EEG laboratory, the technician is deprived of a controlled environment. This is so since the prime concern in the ICU is the support of patient's life and, particularly in the early hours after surgery, many procedures both diagnostic and therapeutic are in progress. The effect of these artifacts on computer analysis is quite serious and their rejection is essential for authentic interpretation of the recording. In the ICU the sources and characteristics of artifacts are numerous and varied. The more common types of EEG artifact are listed below.

* 60 Hz artifact from the mains power.
* Patient movement.
* Poor electrode contact.

60 Hz artifact is due to interference from nearby equipment or the very common ground loop. Figure 3.4 illustrates 20 seconds of an EEG with 60 Hz artifact most prominently visible in the right frontal head regions (channel f4-c4). In the ICU where several electrical devices are operated simultaneously, power supply interference is quite common. This artifact also occurs when a patient is grounded more than once and there is a difference between the grounds. Yet another source of 60 Hz artifact is when a ground electrode is shorted to one of the active electrodes.

Figure 3.4 - 20 seconds of an EEG record corrupted by 60 Hz artifact. The artifactual distortion is present in channels c3 - p3, p3 - o1, f4 - c4, c4 - p4 and p4 -o2.

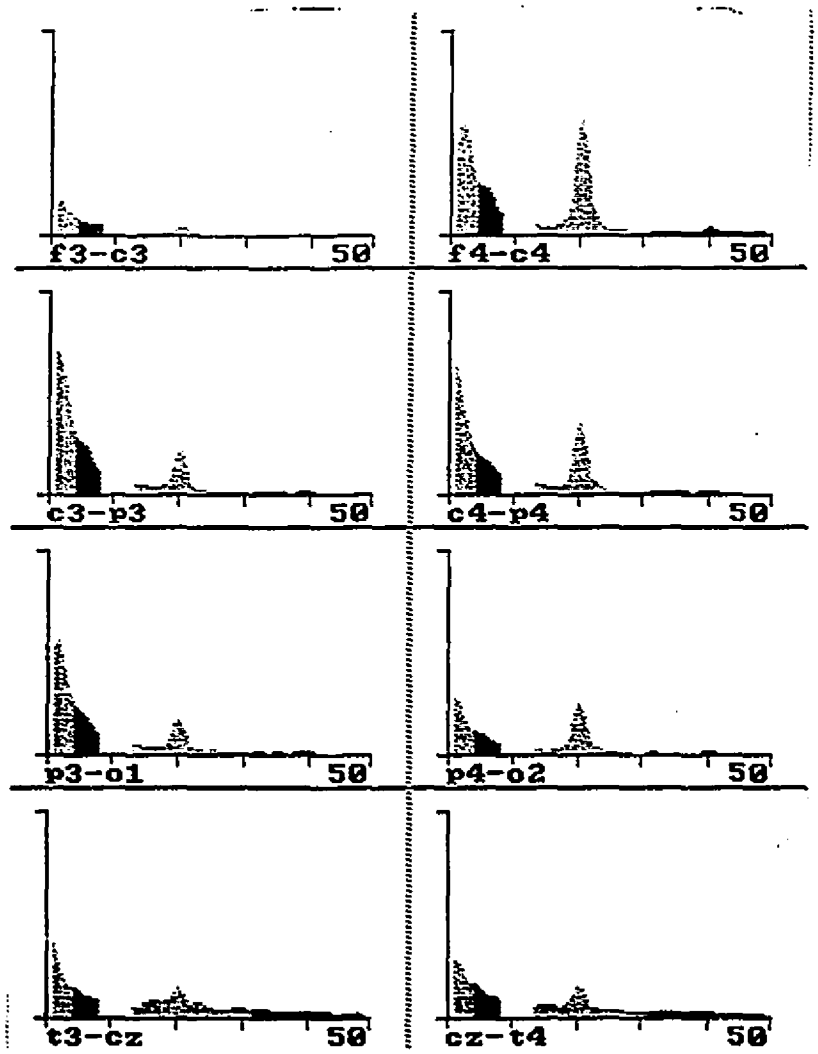

Figure 3.5 - Spectra (1.17 Hz to 50Hz) of EEG illustrated in Figure 3.4. Peaks in activity at 20 Hz are observed in channels corrupted by 60 Hz artifact due to harmonics.

In addition to 60 Hz, the above-mentioned sources also generate harmonics at 120 Hz, 180 Hz, etc. which contaminate the EEG. Since the signal is digitized at 200 Hz, the activity at 120 Hz aliases and appears as an activity peak at 80 Hz and 180 Hz activity aliases as 20 Hz. The 20 Hz activity contaminates the beta band. Figure 3.5 above illustrates the amplitude spectra (1.17 to 50 Hz) of the EEG shown in Figure 3.4. The peak at 20 Hz is due to aliasing as explained earlier.

Patient-generated artifacts include body movement, muscle contraction of the scalp, blinking, chewing, coughing, swallowing, hiccoughing and involuntary myoclonic jerks. Contractions of the scalp muscles and other muscles due to chewing, coughing and swallowing produces broad band activity referred to as the Electromyogram (EMG). The EMG is quite useful for behavioral studies, however, while monitoring cerebral function these rhythms are not of interest and they contaminate the EEG making interpretation difficult. Another source of patient-generated artifact is the EOG (Electro-occulogram) generated by eye-movement.

Patient movement causes artifactual signals both due to mechanical movements of electrical contacts and movement of the conductors that carry current from the scalp electrodes to the amplifiers. This conductor motion results in induction of an electric current due to the earth's magnetic field. The current through the conductor due to the potential recorded at the scalp is very small due to the magnitude of the potentials themselves and the very large impedances of the amplifiers. The induced currents are thus comparable to the currents of cerebral origin and therefore contaminate the EEG signals. Patient examination and physiotherapy also cause rhythmic artifact due to discharges of static current.

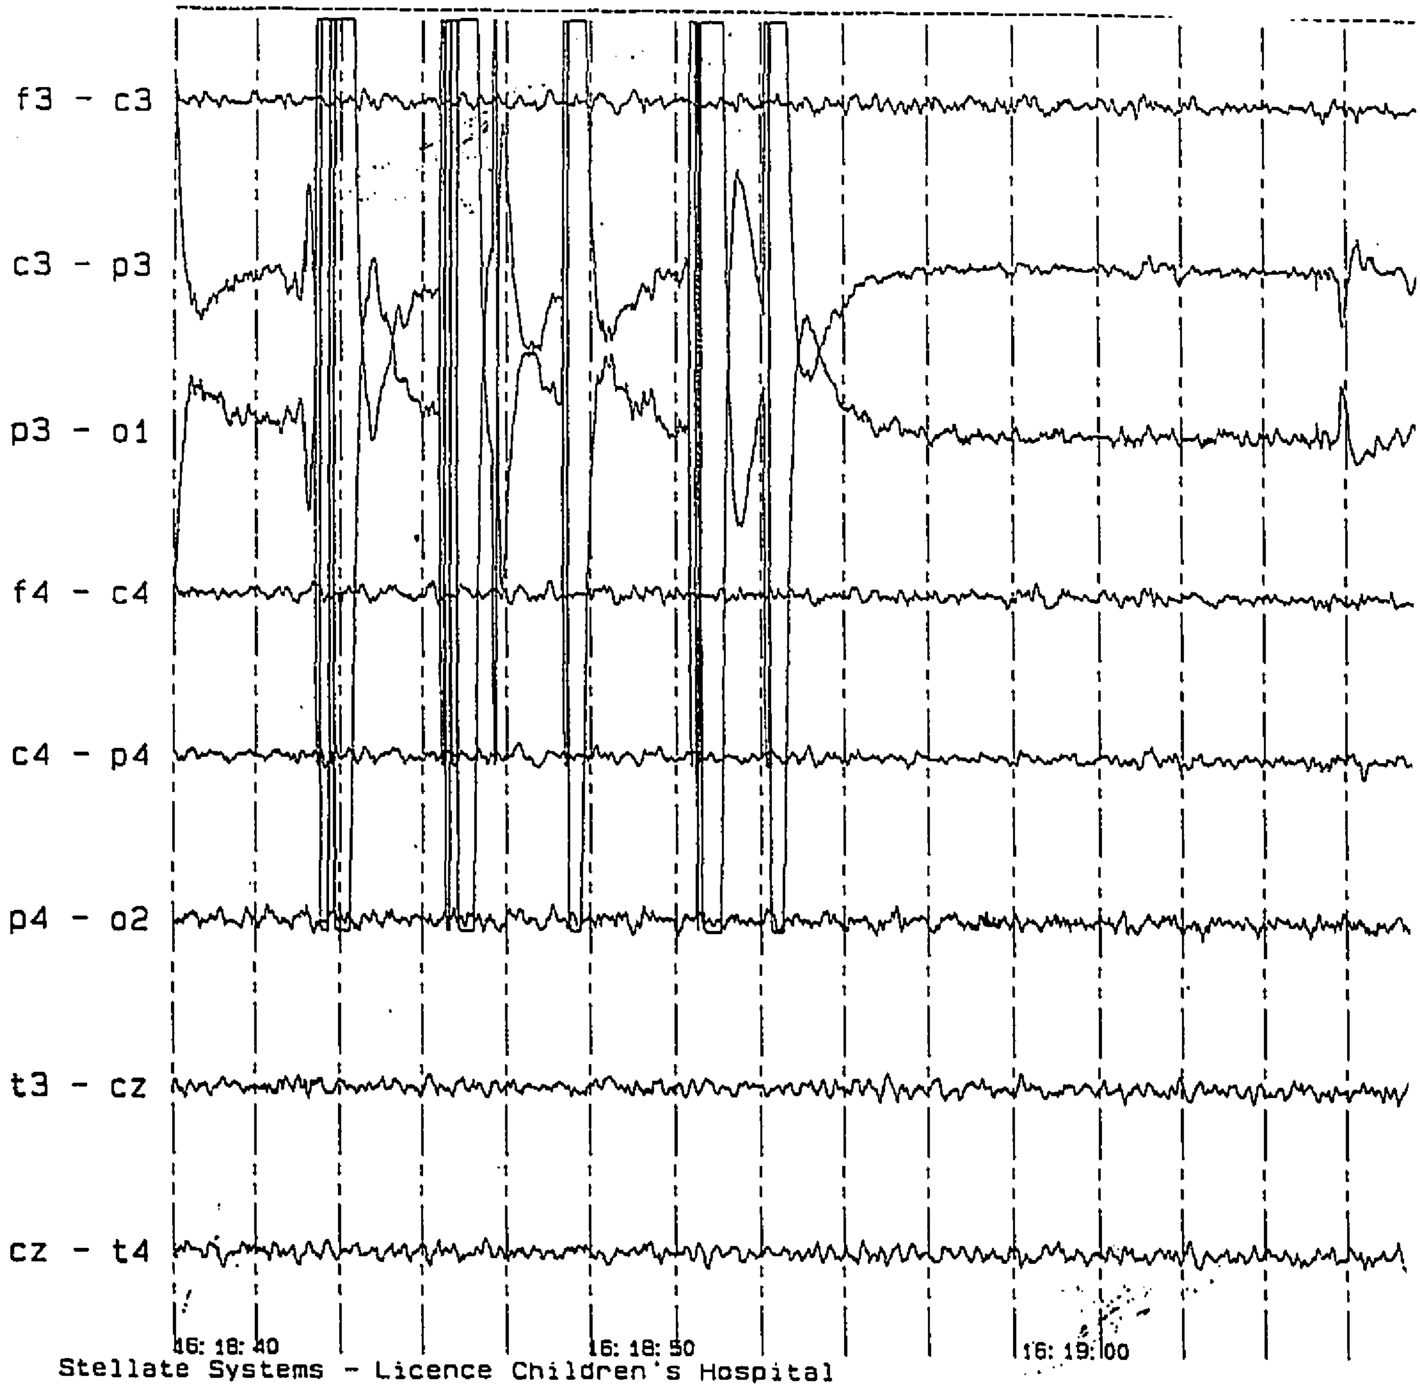

Figure 3.6 - Patient-generated artifact. Artifact is observable in channels c3 - p3 and p3 - o1 due to movement of conductor from the common electrode p3.

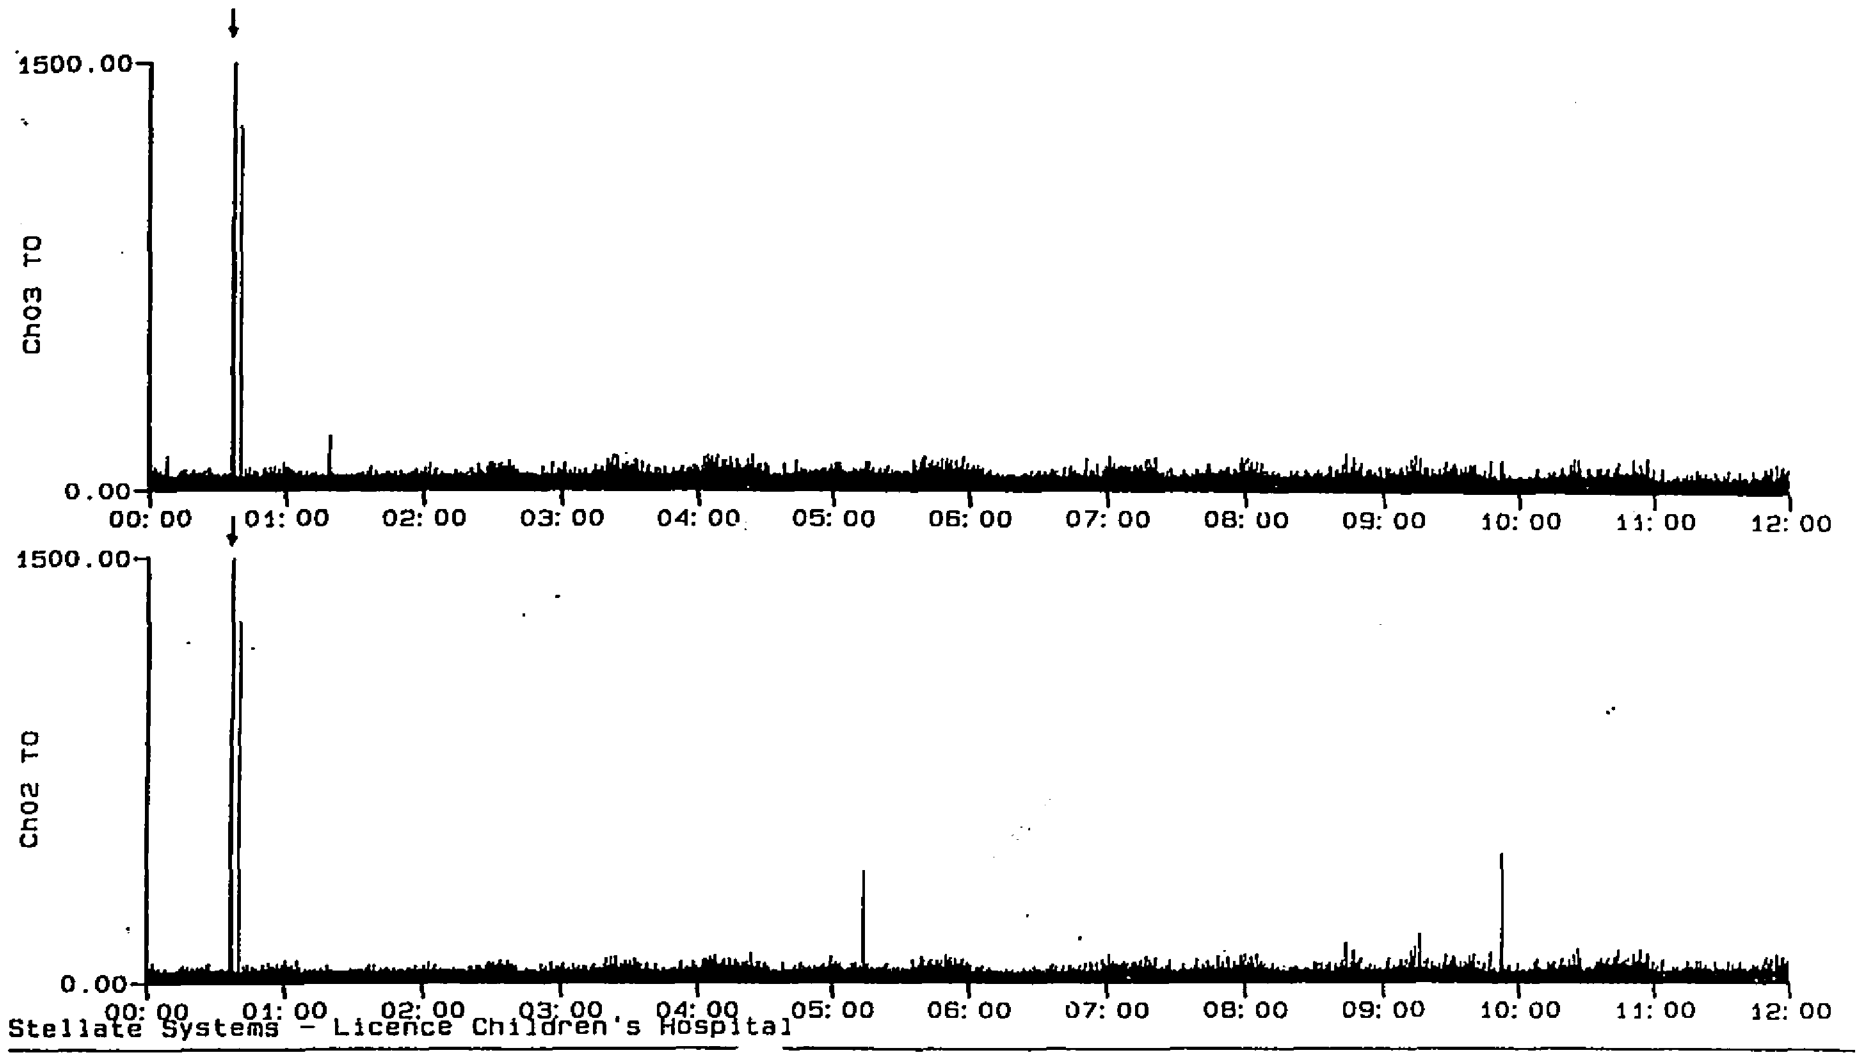

Figure 3.7 - Total band array of channels with artifact in Figure 3.6. Arrow points to the activity corresponding to the EEG in Figure 3.6. The magnitude of the patient-generated artifact as compared to the normal background is appreciable from this figure.

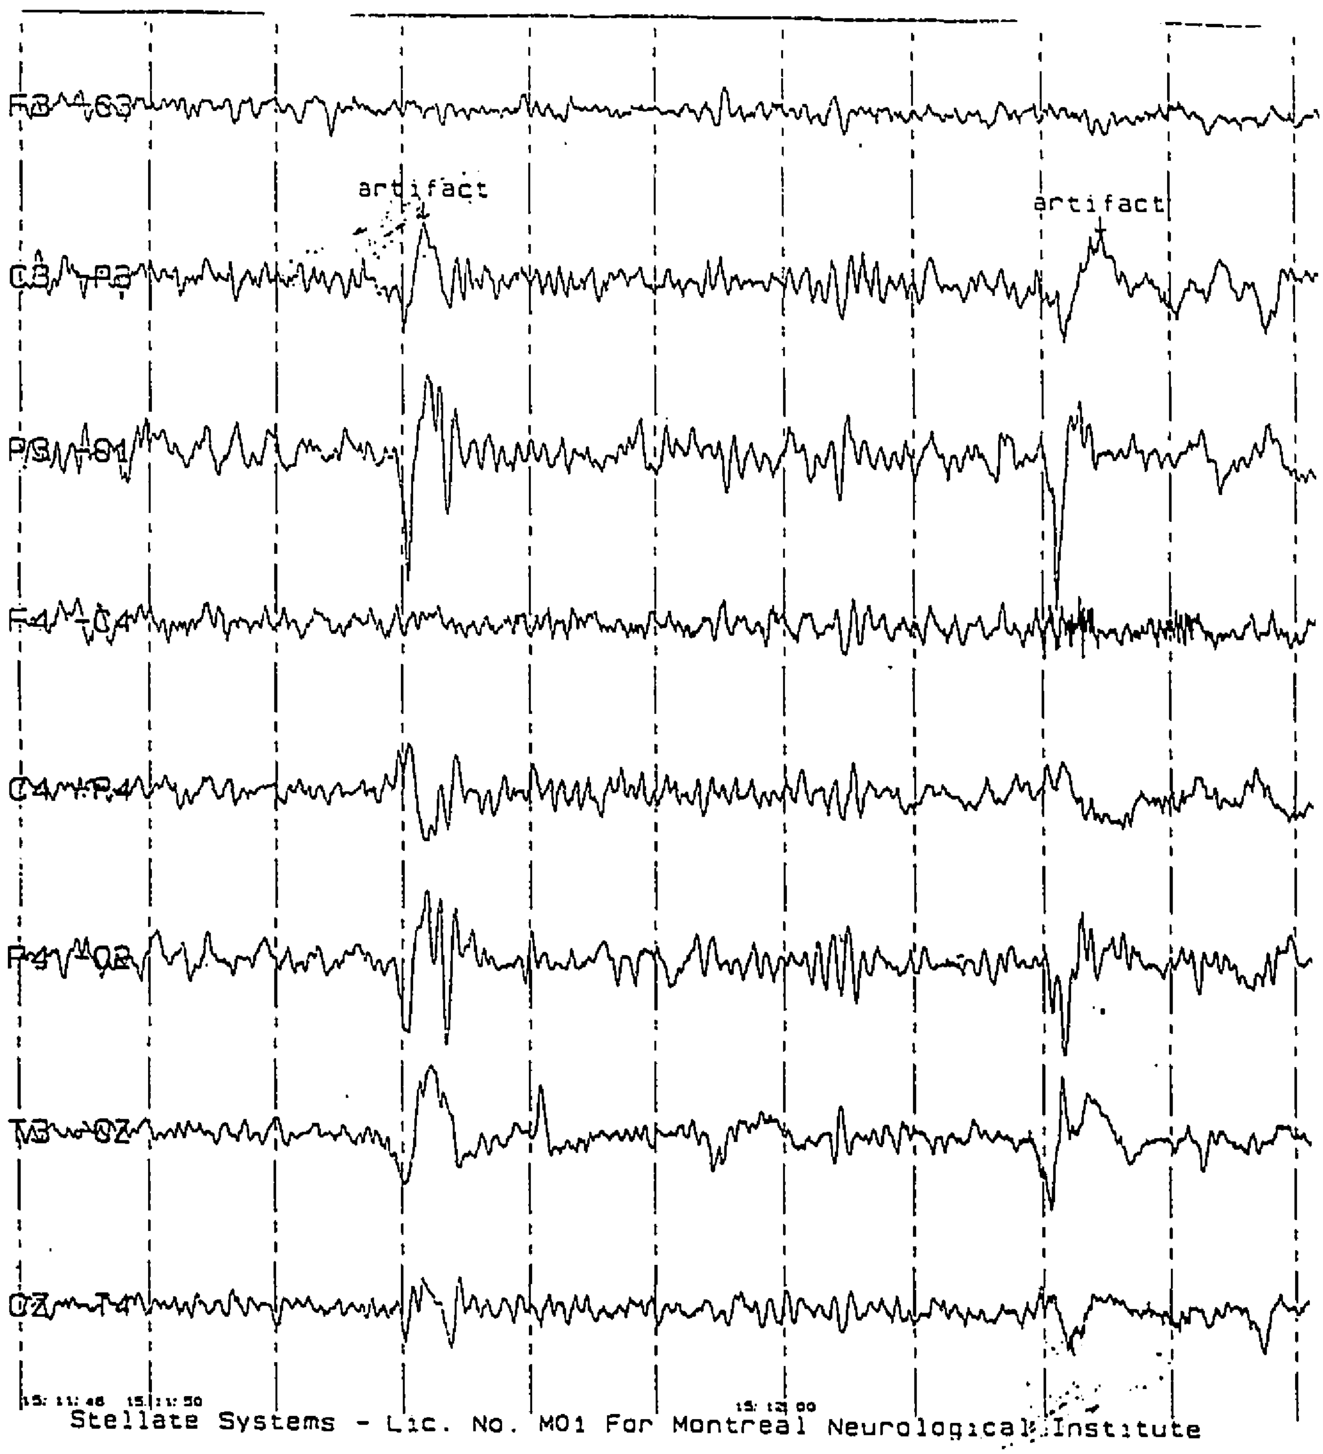

Figure 3.8 - Artifact due to poor electrode contact.

Patient-generated artifacts are usually spiky, containing sharp elements and very large amplitudes. They stand-out quite clearly from the background. These are usually impulsive not lasting for more than a second in duration. Figure 3.6 illustrates an EEG contaminated by patient-generated artifact. In the band arrays, these artifacts clearly stand out from the background and are usually solitary. Figure 3.7 illustrating the total band array of the EEG in Figure 3.6, helps appreciate the difference in amplitude between EEG of cerebral origin and patient-generated EEG artifact. Here the very large amplitude spiky components correspond to patient generated artifact.

Artifact due to poor electrode contact is quite inevitable during a long-term recording and this is illustrated in Figure 3.8. Electrodes are glued on to the surface of the head and a jelly fills the gap between the electrode and scalp to make contact. During long-term recordings, the jelly could dry up and this may impair the contact. Also, the electrode itself may not be glued properly. Artifact due to poor contact is characterized by low frequency and moderately high amplitude. Such artifacts however, are not impulsive and remain until the electrode is glued again or fresh jelly is injected into the junction. In the band arrays, the artifact is usually seen in the delta and total bands and occasionally in the theta band as they are characterized by very low frequencies. The band array corresponding to an EEG with poor electrode contact is illustrated in Figure 3.9a. The humps in activity here correspond to electrode contact artifact. Figure 3.9b illustrates the band array of another EEG where the humps correspond to genuine fluctuations in the frequency composition of the EEG.

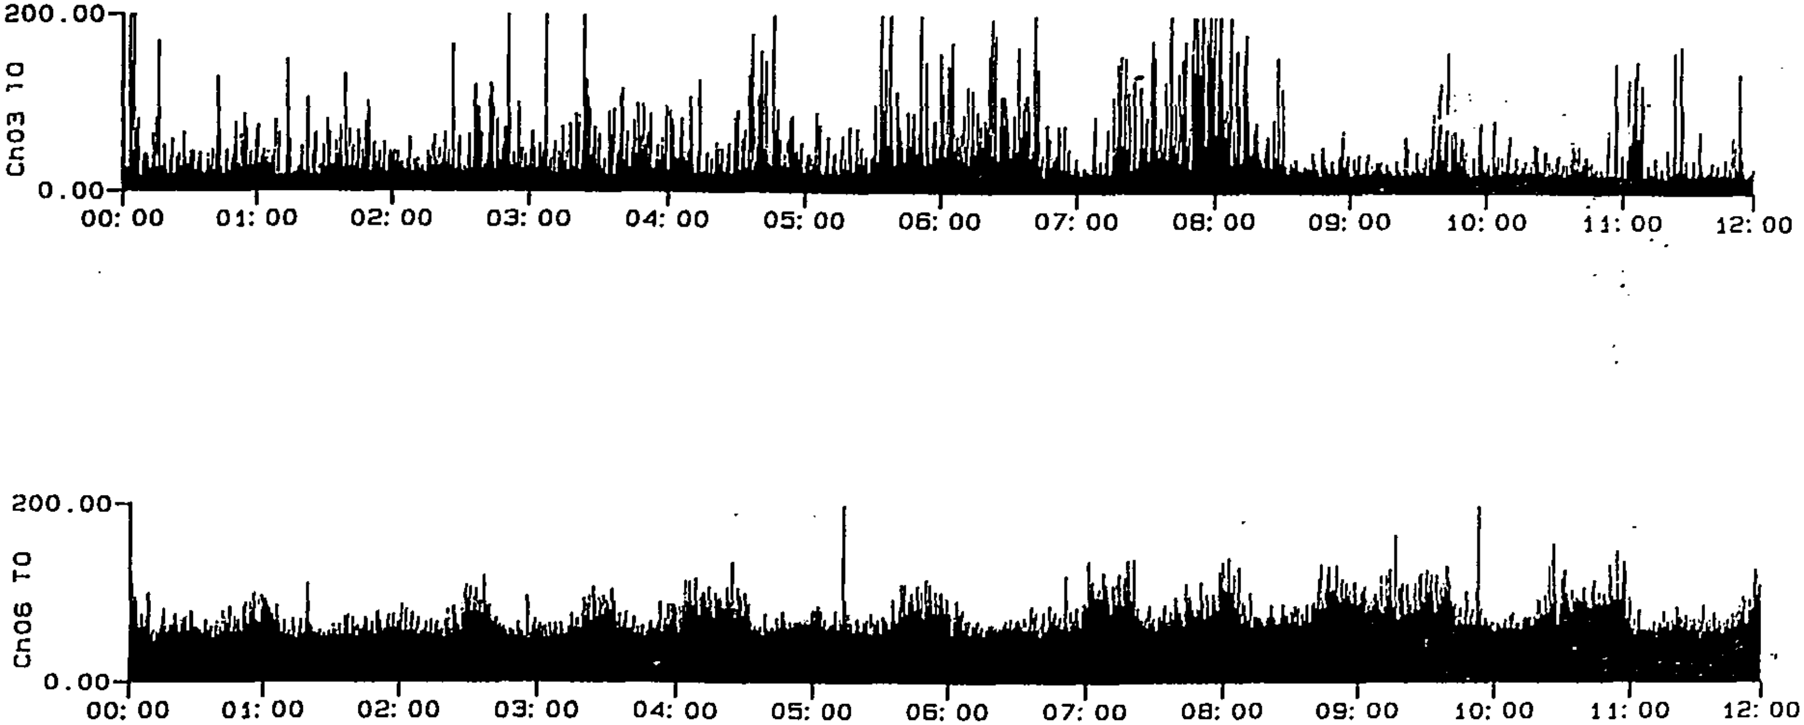

(a)

(b)

Figure 3.9a - Total band array of EEG with sustained artifact.The humps in activity at 6 hrs and 8 hrs are due to sustained artifact. 3.9b - EEG with no sustained artifact. The humps in activity at 7 hrs and 9 hrs are non-artifactual spontaneous variations of the EEG. Although generally of smaller amplitude, these are not very different from the ones in (3.9a).

The humps in both 3.9a and 3.9b look quite alike even though one is artifactual while the other is not. Removal of the artifactual humps while preserving the genuine humps could be a difficult problem to solve.

3.2.2. Artifact Rejection

As discussed earlier, the aim of this diagnostic tool is to study gross changes of the EEG background over several hours. The state of any EEG is determined mainly on the basis of its frequency composition over several hours, average amplitude and symmetry in activity between different regions of the head. It is therefore sufficient to reject artifacts large enough to manifest themselves in the frequency spectrum. It is quite unnecessary to identify and reject specific artifactual EEG waves in the time domain. Each point in the band array of the activity file represents average activity over a 30 second period in the corresponding frequency band for the particular channel. Artifact identification and rejection from these band arrays would be sufficient for our purposes. It is important to realize that artifact identification can be done visually only by examining the raw EEG record in the time domain. However, suppression of the artifact is done in the frequency domain. Therefore all examples of band arrays of artifactual or non-artifactual EEG presented in this section are chosen subsequent to visual examination of the raw EEG itself and not by examining the band activities.

The frequency range of background EEG of cerebral origin is 0-30 Hz. Some faster rhythms of cerebral origin up to 50 Hz may be present but usually correspond to ictal activity. Therefore any activity at 60 Hz could be treated as artifact and rejected. The artifact due to harmonics however, alias and appear at lower frequencies and contaminate the beta frequency band (14 to 30 Hz). This band itself is altered by medication and administration of anesthetics (Mahla) thereby making interpretation difficult. To remove the effect of these non-pathological confounders, the 60 Hz artifact and its harmonics, the total band is defined from 0 to 14 Hz. As seen in the previous section, most background EEG abnormalities manifest themselves in the delta (1 to 3 Hz) and theta (3 to 7 Hz) frequency ranges. Therefore, such a definition of the total band is not likely to distort or handicap the interpretation in a significant way.

To reject the artifacts due to patient movement and poor electrode contact discussed above, a filter capable of rejecting both impulsive very high-amplitude artifact and sustained moderately high amplitude artifact is required. Linear filters are used extensively, in several signal processing applications today. However, they fail to perform well with high amplitude impulsive noise components. Linear filters merely smear the effect of impulsive noise components. The performance of a 5-point moving-average filter on an artifactual EEG is illustrated in Figure 3.10. (3.10a) illustrates the unfiltered band array with impulsive EEG artifact and (3.10b) is the band array subsequent to filtering using the 5-point linear filter. (3.10d) illustrates the filtered output of the band array in Figure 3.10c with sustained artifact. The filter averages the band array values within a 5-point window. Impulsive high-amplitude artifact is merely smeared rather than removed as observed in the illustration (3.10b). The linear filter does not reject sustained artifact either as is evident from Figure 3.10d.

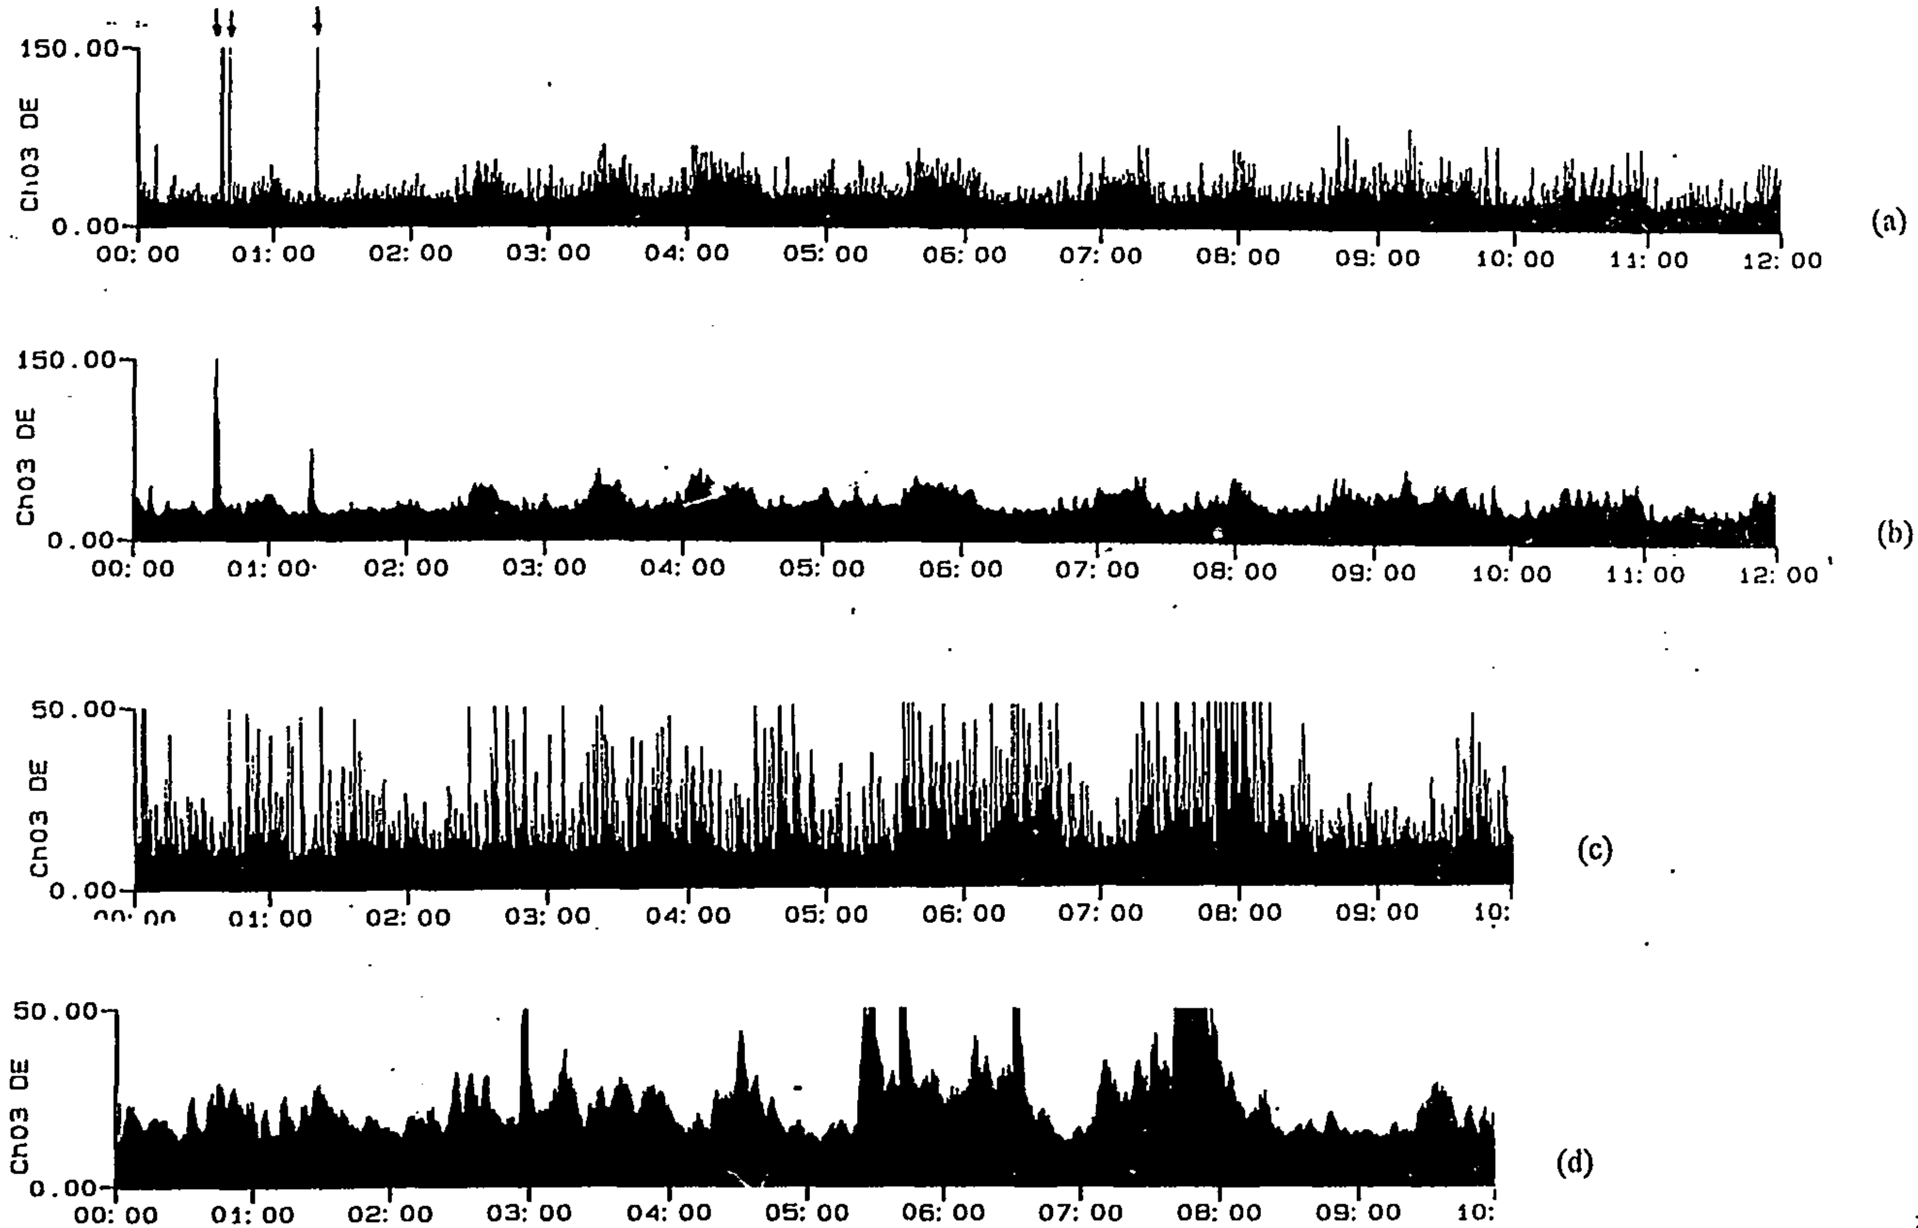

Figure 3.10a - Delta band array of EEG with impulsive artifact. Arrows point to the artifact. 3.10b - 5-point linear filter output of (3.10a). 3.10c - Delta band array of EEG with sustained artifact. 3.10d - 5-point linear filter output of (3.10c).

A non-linear filter called the median filter has been used successfully with impulsive noise components. The median filter replaces the value at each point by the median of the signal values in some finite neighborhood about that point. The non-linear characteristics of this filter help remove impulsive noise components while maintaining sharp edges at the same time. However, it fails to reject sustained noise components. The performance of a 5-point median filter on an EEG with impulsive artifact and on one with sustained artifact is illustrated in Figure 3.11. Figure 3.11b illustrates the effect of median filtering on the band array illustrated in Figure 3.11a. The median filters removes the artifact leaving no trace of it. (3.11c) represents the band array of an EEG with sustained EEG artifact and (3.11d) shows the effect of the median filter on it. In the case of sustained artifact the median filter performs rather poorly and most artifact passes right through and is unrejected as is evident from (3.11d).

From the discussion above, it is clear that the median filter is efficient at rejecting impulsive noise components. With sustained noise however, the performance of neither of the filters discussed above is remarkable. Sustained artifacts in EEG recordings can extend over several hours. In such cases, both linear and median filters dampen the artifactual activity initially. After a certain period of time (depending on the size of the window) when the entire neighborhood is artifactual and similar, these filters fail to recognize and reject them as artifact. Therefore to get rid of artifacts that extend over long periods of time, a hard-limiter threshold should be used.

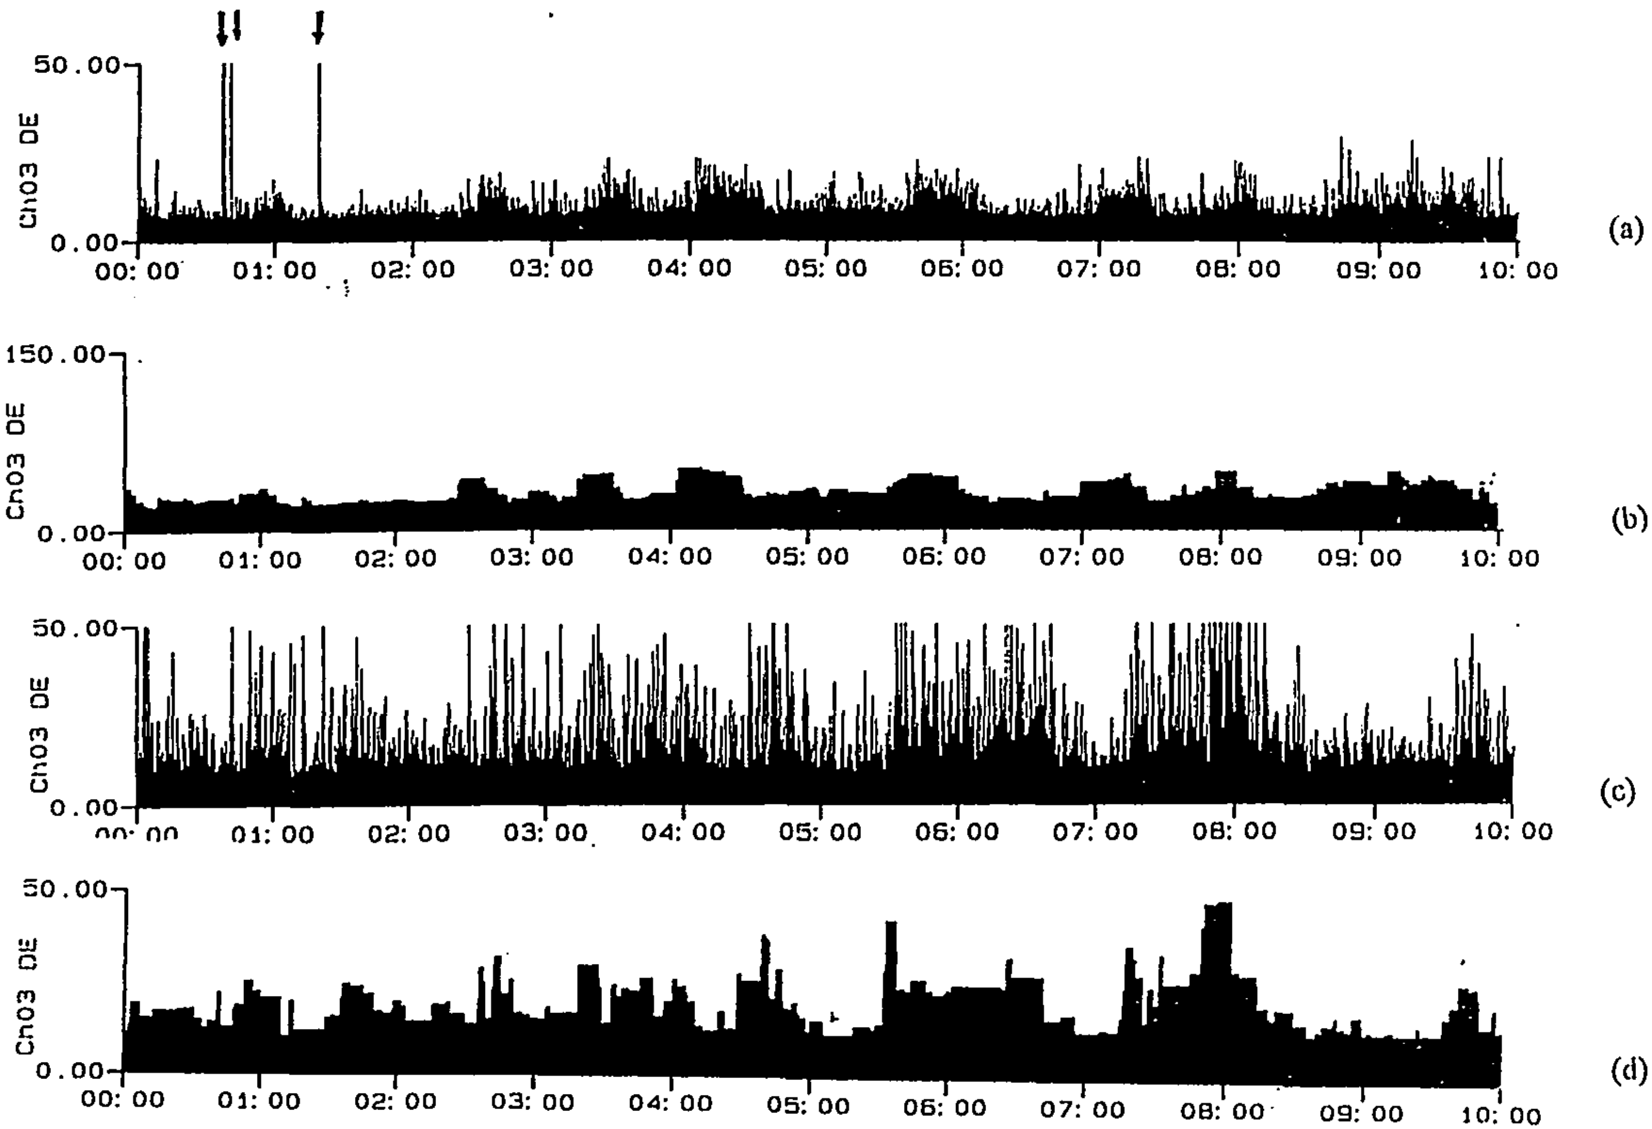

Figure 3.11a - Delta band array of EEG with impulsive artifact. Arrows point to artifact. 3.11b - 5-point median filter output of (3.11a). 3.11c - Delta band array of EEG with sustained artifact. 3.11d - 5-point median filter output of (3.11c).

A median filter with a hard-limiter threshold performs well with the rejection of sustained artifact. The data is first filtered with a 5-point median filter. If the median of any window of total band activity is greater than a certain threshold, then the activities in all frequency bands of that particular epoch for that channel are replaced by the respective average band activities averaged up to that point in time, thus rejecting the artifact. The choice of the threshold is quite critical for the performance of the filter. An absolute threshold may not be a good choice since the amplitude of the artifact itself is proportional to the average amplitude of the background EEG. A high threshold may fail to reject artifacts of depressed EEGs while a low threshold may reject genuine high amplitude rhythms of cerebral origin. The threshold should therefore be relative to the average total band activity.

The amplitude of most artifacts including that due to poor electrode contact is at least 1.5 times the average EEG background activity. Extensive EEG review indicates that genuine fluctuations of cerebral origin in total band activity, i.e. average EEG activity in the 0 to 14 Hz range over a 30 sec. period, is rarely greater than 25% of the average background activity. Therefore, it is safe to assume that total band activity greater than the overall average total band activity by at least 1.5 times is artifactual.

As explained earlier, the band arrays are first filtered with a 5-point median filter which rejects impulsive artifacts. If the median in any window of total band activity is greater by 1.5 times than the corresponding running average up to that point in time, then the activity values in all band arrays corresponding to that epoch and channel are replaced by their respective running averages in activity. This requires that the first window that is filtered is artifact free in order to build an artifact free running average. This should be ensured by the user. The monitoring system prompts the user to input the time to start analysis. The user then should ensure that at least 3 minutes of "clean" EEG forms the start of analysis.

The performance of a median filter with a hard-limiter threshold on the signals in Figure 3.12a and 3.12c are illustrated in the Figures 3.12b and 3.12d respectively. Clearly, the hard-limiter threshold rejects both sustained and impulsive artifacts quite effectively. All other filters discussed above fail to perform well when the artifact fills the entire window. All data points are then similar and the filter does not have a sample of true unartifactual band array values. Further these filters reject as artifact, what does not conform with the rest of the members of the same window. In the case of sustained artifacts therefore, nothing is rejected. The hard-limiter threshold filter however, has a sample of true band array values in the form of a running average. The filter has a reference uncontaminated by artifact and is therefore able to reject artifact by comparison even when the artifact is sustained and fills the entire window.

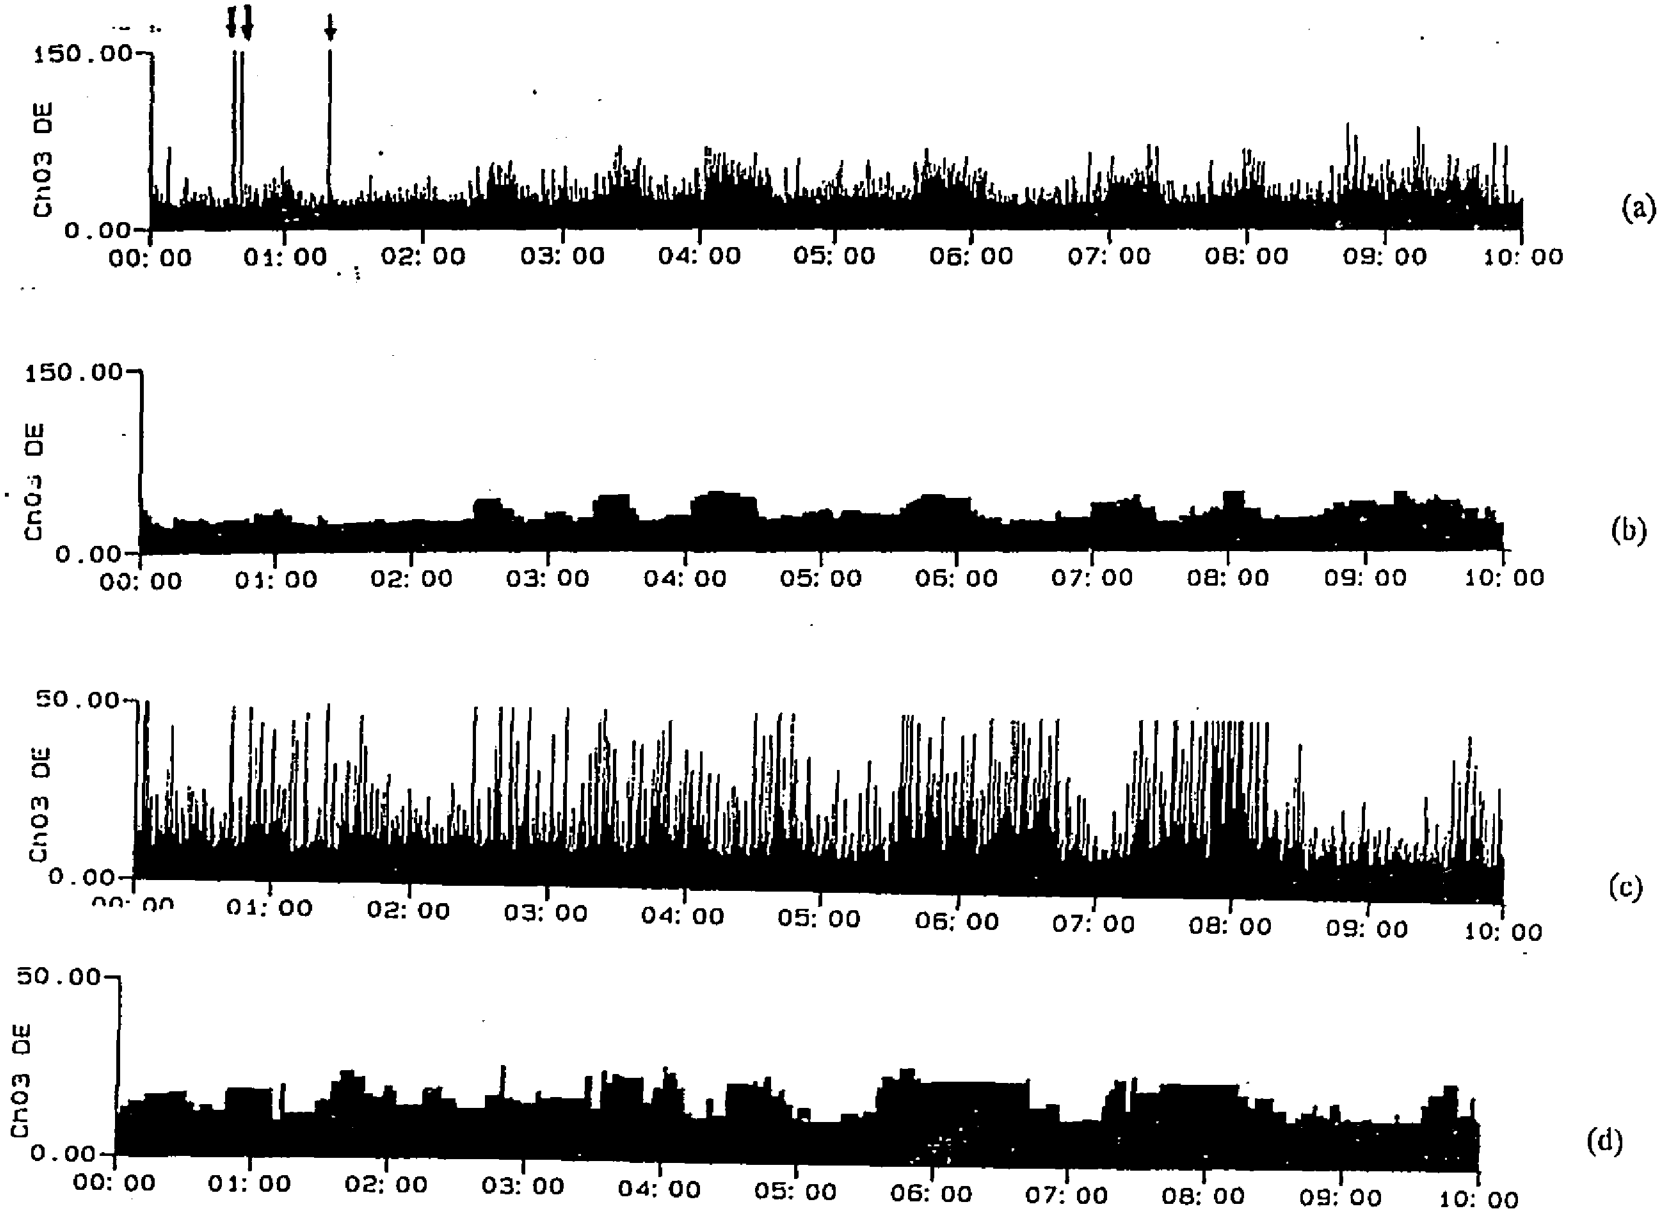

Figure 3.12a - Delta band array of EEG with impulsive artifact. Arrows point to artifact. 3.12b - Output of a 5-point median filter with hard-limiter threshold at 1.5 of the band array in (3.12a). 3.12c - Delta band array of EEG with sustained artifact. 3.12d - Output of a 5-point median filter with a har-limiter threshold at 1.5 of the band array in(3.12c)

Other threshold values have also been tried. A lower threshold tends to reject true EEG as artifact. This is illustrated in Figure 3.13, where a threshold of 1.25 rejects genuine EEG as artifact. Figure 3.13a represents the original band array and figure 3.13b the result subsequent to filtering. A high threshold on the other hand accepts artifactual EEG as real EEG. This is seen in Figure 3.14 which illustrates the performance of the filter on the band array shown in Figure 3.14a for different threshold values. Figure 3.14b shows output for a threshold value of 2.0 and figure 3.14c shows the output for a threshold of 1.5. Comparing the two outputs, the superior performance of a threshold at 1.5 is clearly demonstrated.

Once artifact rejection is complete, the band array is ready for feature extraction discussed in the following Chapter.

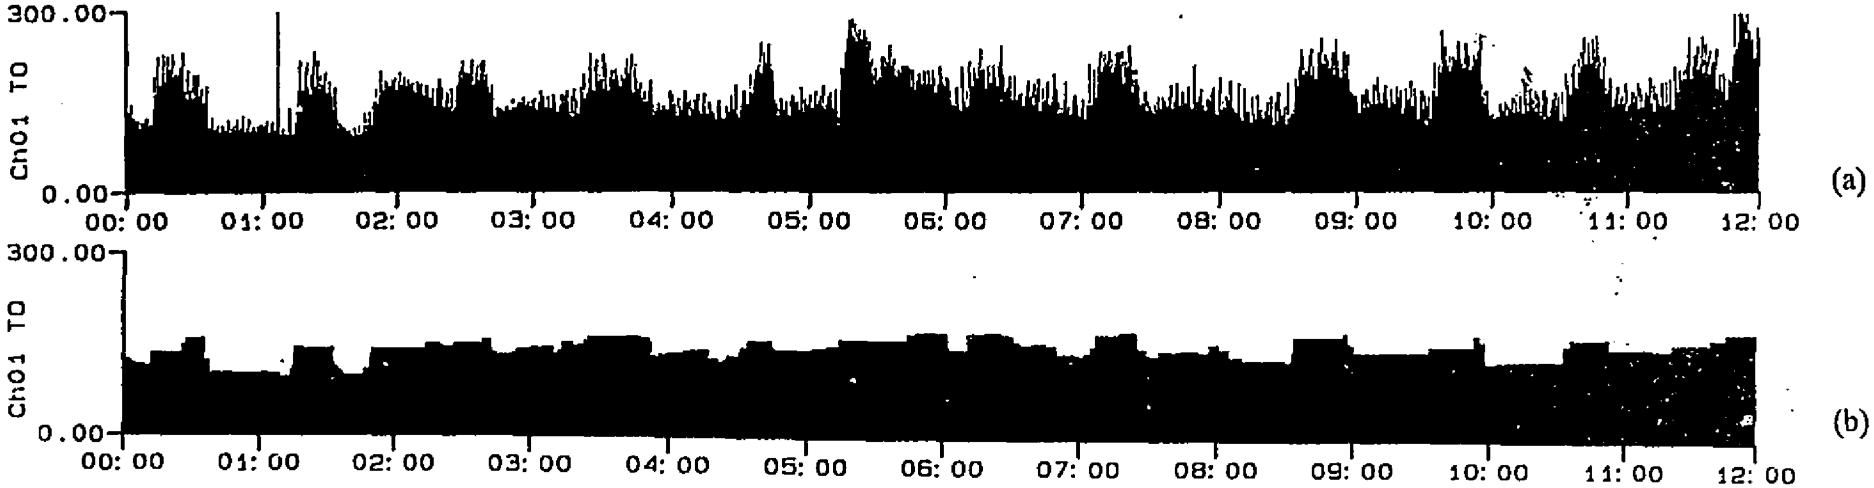

Figure 3.13a - Total band array of the frontal channel. 3.13b - Filtered version of (3.13a) with a 5-point median filter and a hard-limiter threshold at 1.25. The activity humps due to genuine EEG variability visible in (a) at 2 hrs, between 5 and 6 hrs and between 9 and 10 hrs are decapitated in (b). This is due to a very low threshold limit.

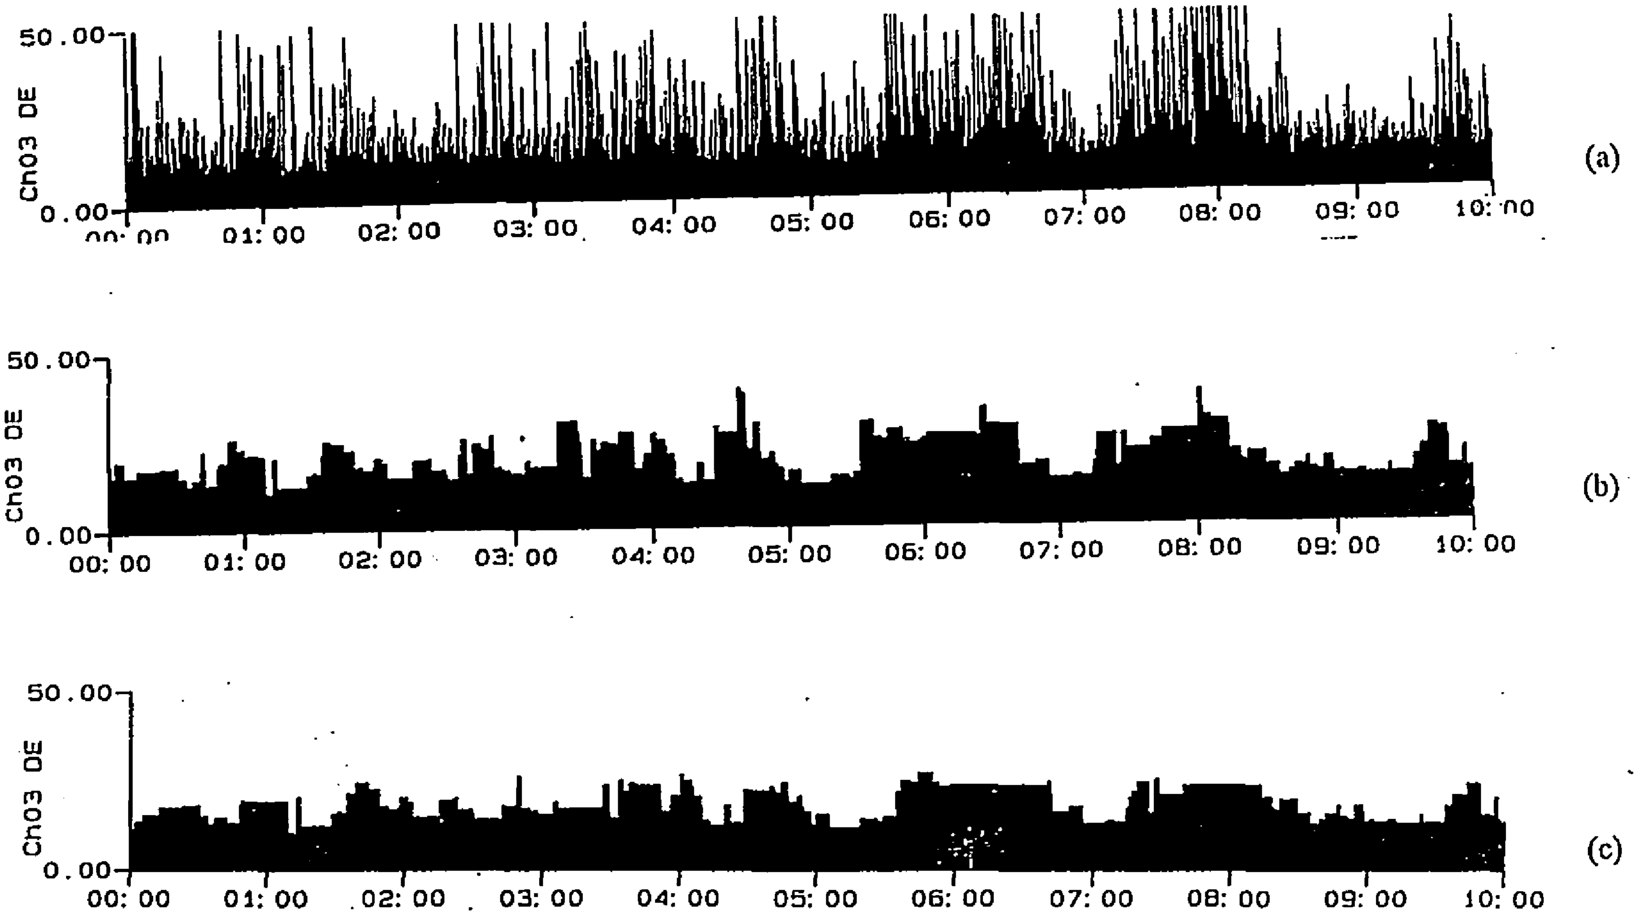

Figure 3.14a - Delta band activity of EEG with sustained artifact. Filtered version of (a) with a 5-point median filter and hard-limiter threshold at 2.0 (3.14b) and 1.5 (3.14c). Artifact rejection is superior with threshold at 1.5 as evident with the humps at 6 hrs and 8 hrs.

4. FEATURE SELECTION

From the discussion on abnormal EEG patterns it is apparent that the nature of background EEG activity is primarily assessed on the basis of four aspects, namely,

* Amplitude
* Left / Right symmetry
* Variability of the EEG
* Anterior / Posterior gradient

A depressed EEG record is characterized by very low amplitude activity. The inactive pattern discussed in Chapter 2 as an indicator of bad prognosis is also associated with low amplitude values, lower than that of a depressed record. An amplitude measure is therefore quite important to detect these patterns. The prognostic significance of a variable EEG pattern, a gradient in slow activity from the anterior to the posterior head regions and symmetry in activity between the two hemispheres is discussed in Chapter 2. These aspects are crucial in classifying an EEG as normal or otherwise.

Quantitative equivalents of the above-mentioned four features are to be derived in order to obtain information about the EEG record that would facilitate its classification as normal or otherwise. To study spontaneous alterations and assess long-term trends effectively, it is necessary to analyze an extended EEG record (> 6 hours). To detect or evaluate variability of the EEG, data recorded for several hours is required. However, it is important to obtain information about an abnormal EEG pattern at the earliest in order to be most useful. Taking both these factors into consideration, classification of the EEG is done based on 6 hours of recording. Features are extracted from 6 hour long EEG records and the classification is done based on these characteristics. The mathematical derivations of the quantitative features extracted from an EEG prior to classification are discussed below.

4.1 Measure of Amplitude

Figure 4.1 illustrates total band activities of two EEG records, figure 4.1a corresponding to a normal EEG and figure 4.1b corresponding to a depressed EEG, after artifacts have been rejected from them. A depressed EEG record is characterized by low amplitude values, and the severity of depression is inversely proportional to the amplitude of the activity. A simple measure of the total band activity quantifies amplitude normality.

Sustained rather than impulsive amplitude abnormalities are bad prognostic indicators. The amplitude measure is mathematically derived as the logarithm of the average over a 5 - minute period of the total band activity of any channel of the EEG record. In a band array, this represents the logarithm of the average of 10 points of total band activity since basic calculations are made every 30 seconds. The reason for using the logarithm of the average rather than a simple average is explained in the chapter on Statistical Analysis.

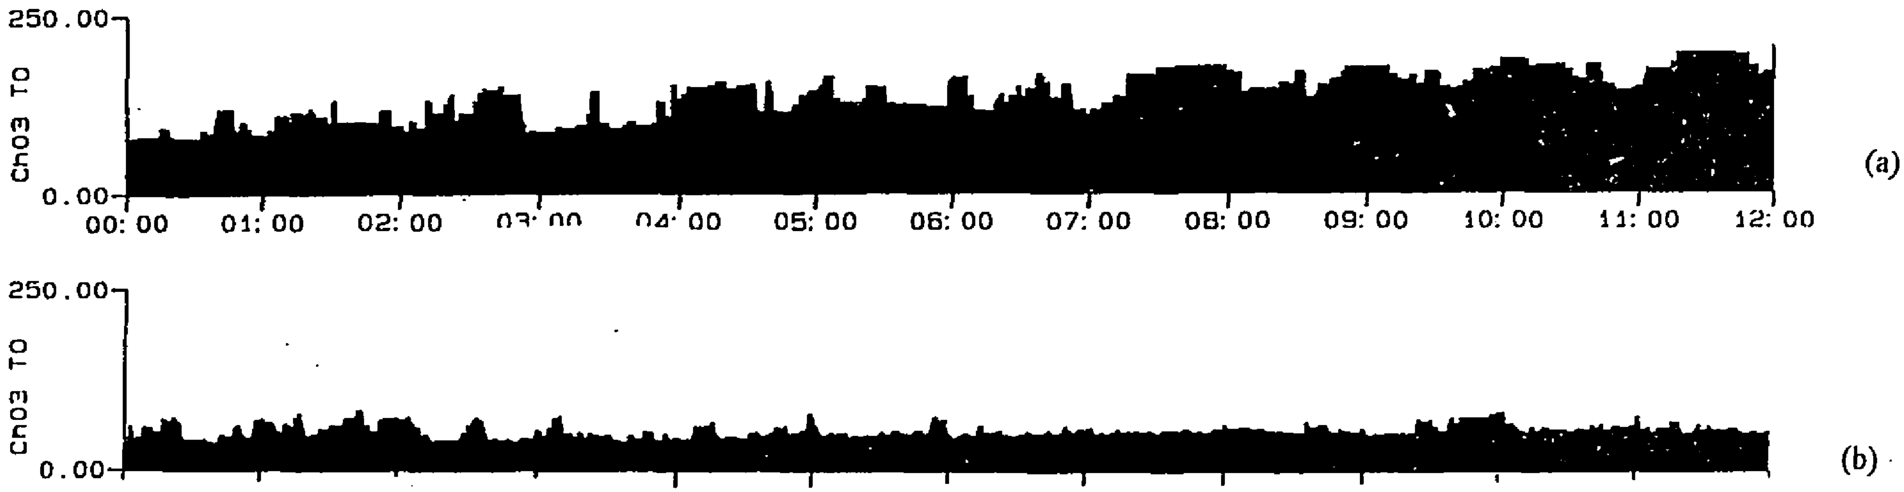

Figure 4.1 - Total band array of left posterior channel of a normal EEG (a) and depressed EEG (b).

Depressed EEG activity may arise from one or more head regions. The amplitude of the EEG activity recorded from the various channels should therefore be ascertained and evaluated for depression. Eight measures of amplitude, corresponding to the eight channels, are extracted from every EEG record that is to be evaluated. An amplitude measure extracted every 5 minutes for every channel therefore amounts to 12 x 6 x 8 or 576 amplitude values for a 6 hour EEG record.

4.2 Measure of Left / Right Symmetry

Figure 4.2 on the following page, depicts delta band activities from the posterior head regions of the left and right hemispheres of two patients, one corresponding to a symmetrical EEG illustrated in figure 4.2a and a second to an asymmetric pattern in figure 4.2b, according to the interpretation of the neurologist. A simple ratio of activity between corresponding channels of the left and right hemispheres could quantify the level of symmetry. A ratio value closer to "1" implies a symmetrical EEG pattern. A value greater than 1 implies a right hemispheric depression and a value less than 1, a left hemispheric depression. The extent of the depression itself is directly proportional to the absolute difference between 1 and the ratio value. With Left / Right symmetry, as in the case of depression, sustained rather than impulsive abnormalities are of concern. Therefore, the delta band activity is smoothed to remove impulsive transients prior to computing the ratio. The formula for the Left / Right Symmetry parameter is given below.

$$r = \text{logarithm}\left[\frac{\text{average activity of channel in the left hemisphere}}{\text{average activity of corresponding channel in right hemisphere}}\right]$$

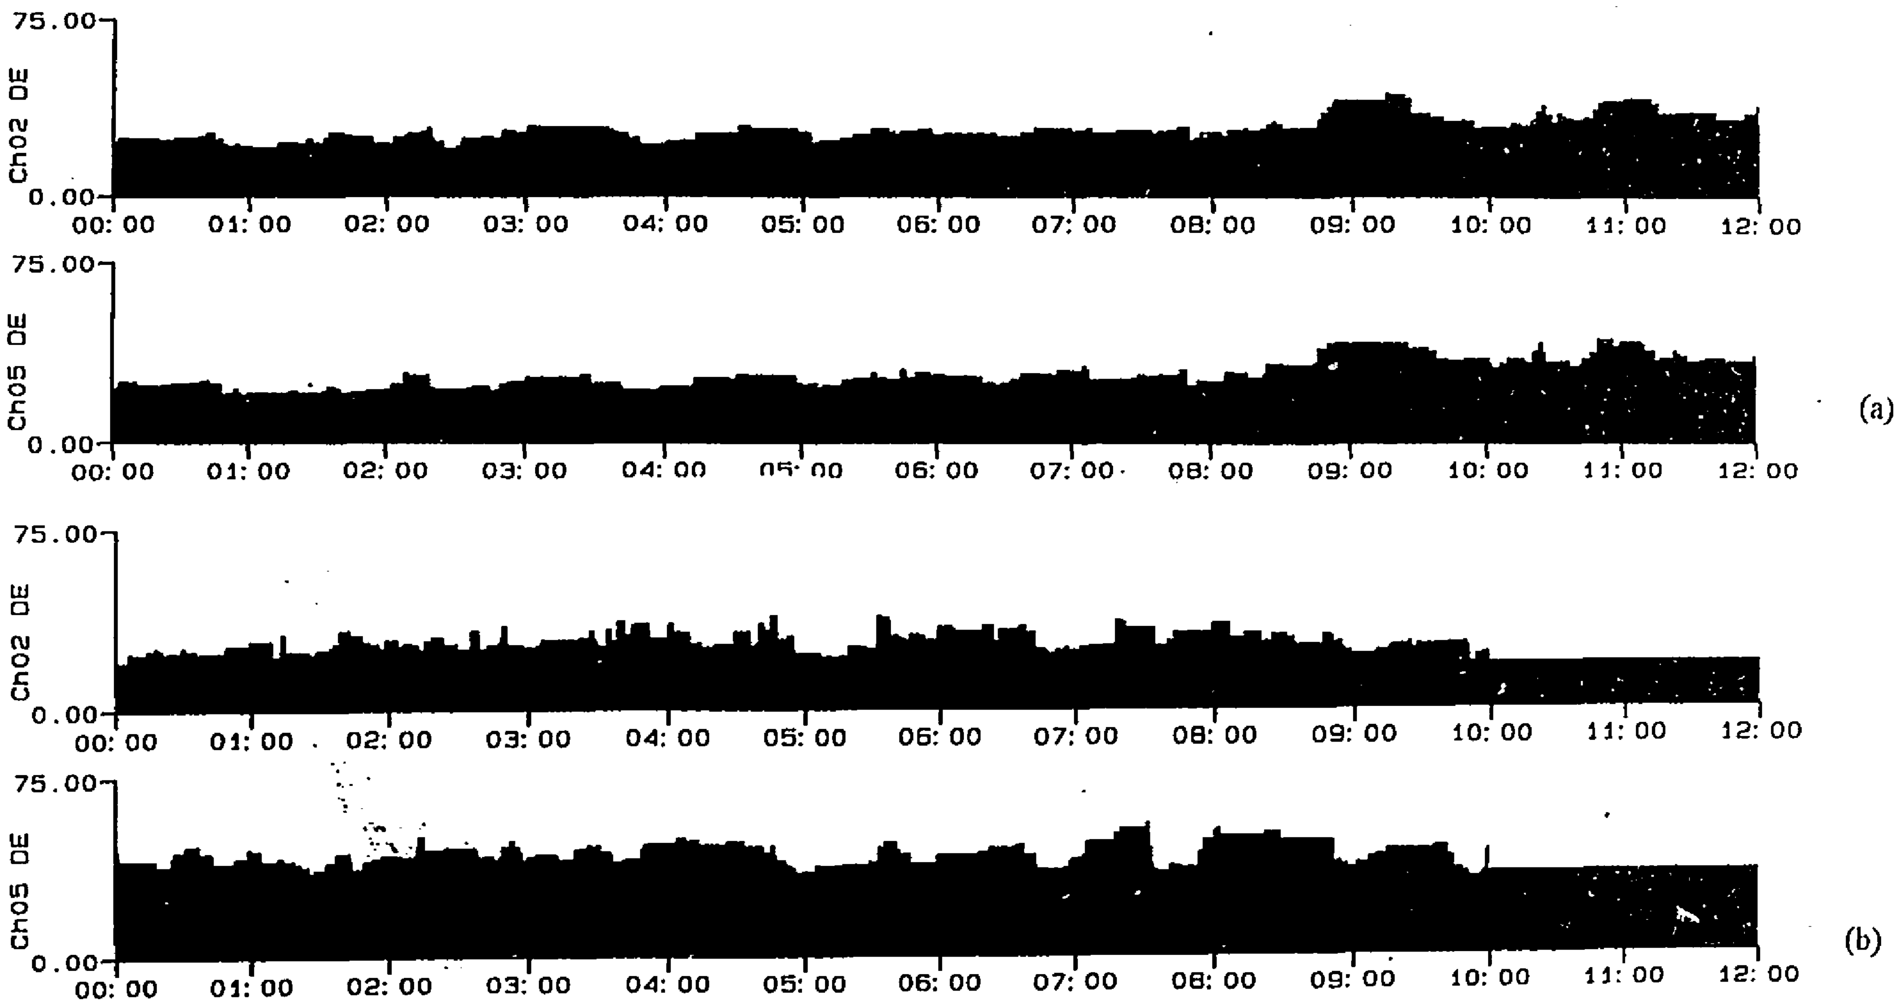

Figure 4.2 - Delta band activities of the central -parietal channels of both hemispheres of a normal record (a) and an abnormal record (b). The amplitudes of the activity of two hemispheres appear equal in 4.2a demonstrating a symmetry in activity. The activity of the left hemisphere (ch02) in 4.2b is relatively depressed as compared to the right hemisphere (ch05) demonstrating an assymmetry in activity.

Average activity here represents average delta band activity over a 5 - minute period which corresponds to an average of 10 points of the delta band activity. Symmetry can also be measured by comparing total band activities between hemispheres instead of delta band activities. However, the total band activity is susceptible to noise over a larger frequency range and could show asymmetries due to non-pathological reasons as well. On the other hand, there is a high probability that an asymmetry in activity due to pathology would manifest itself in the delta band in addition to any other frequency ranges and hence the delta band is preferred.

Four Left / Right symmetry measures corresponding to the four pairs of channels symmetrically located in the two hemispheres are extracted every 5 minutes of an EEG record. Each channel pair monitors activity in one of the following four head regions -- anterior, central-parietal, posterior and the central-temporal and so the measures extracted monitor symmetry in activity in the corresponding head region. A quantified 6 hour EEG record is thus associated with 72 left / right symmetry values for each of the four channel pairs.

4.3 Measure of Front / Back Differentiation

As explained in chapter 2, a normal EEG record is associated with a gradient in amplitude of activity in the low frequencies with amplitude decreasing in the posterior to anterior direction. This is referred to as Front / Back differentiation. Figure 4.3 illustrates the delta band activities from the anterior and posterior channels of the right hemisphere of two patients.

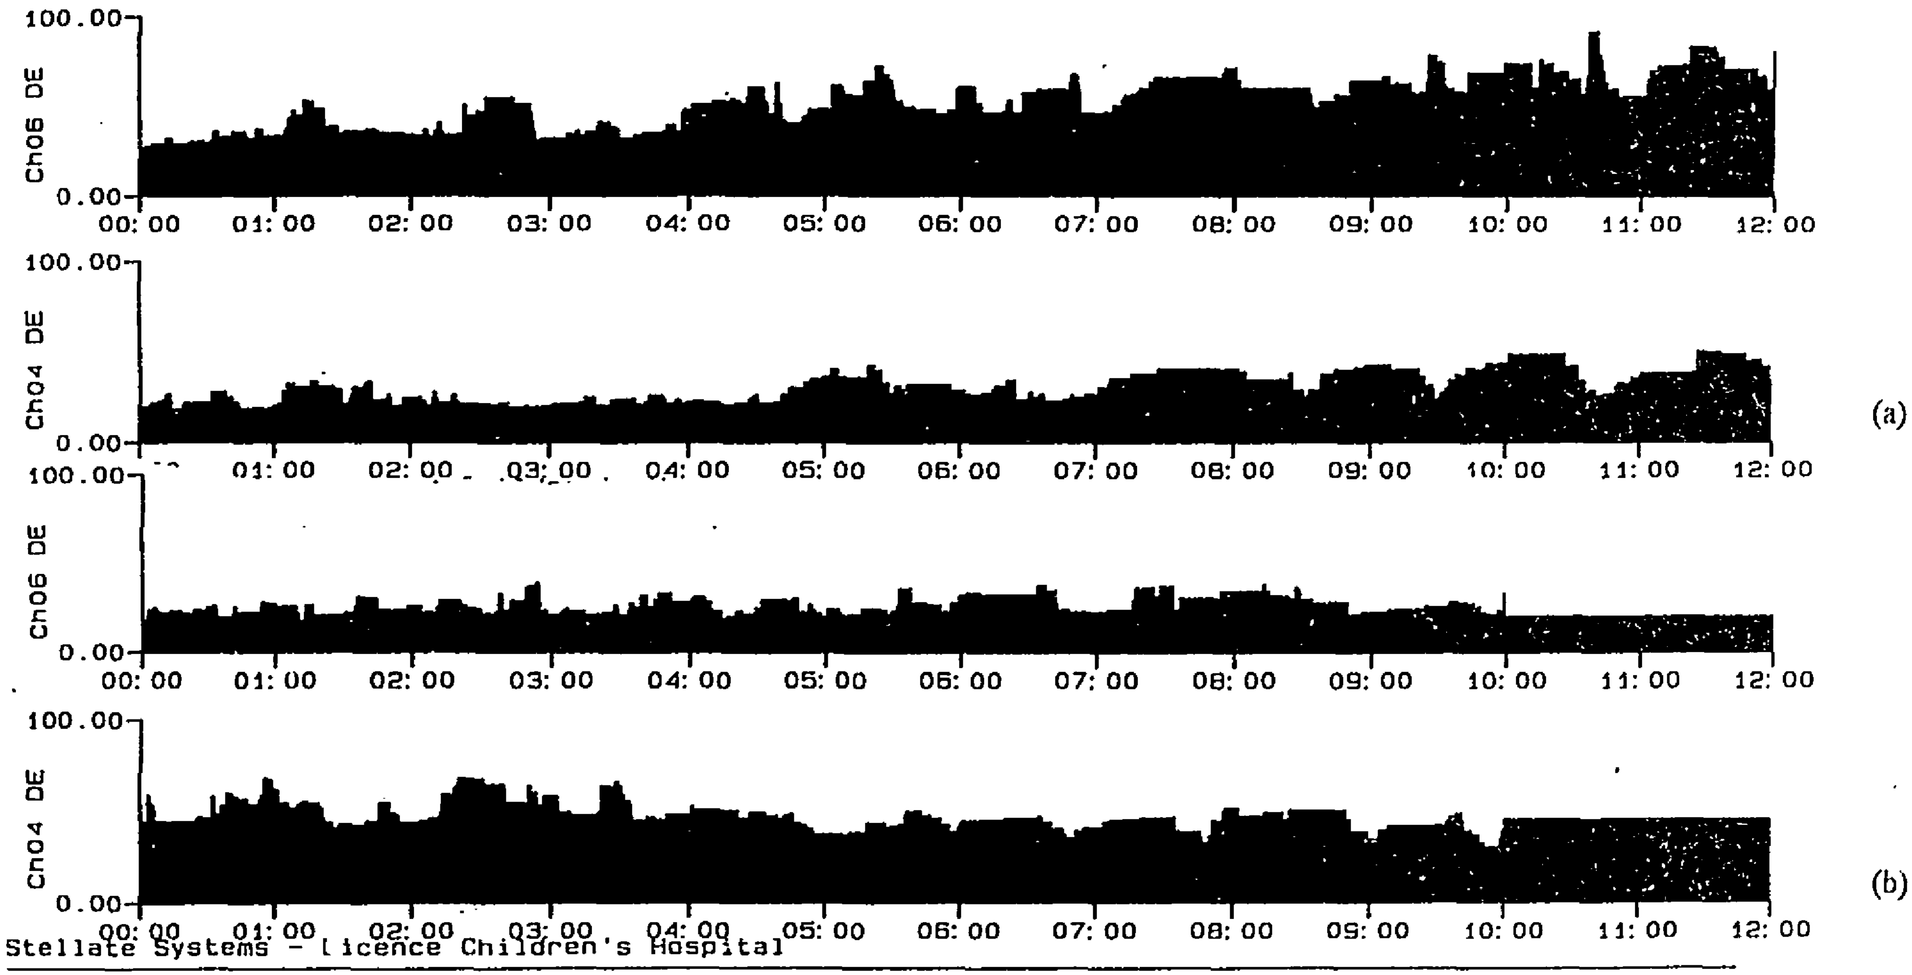

Figure 4.3 - Delta band activities of the posterior (ch06) and anterior (ch04) channels of the right hemisphere of a normal record (a) and an abnormal record (b). A clear anterior-posterior gradient is evident in (a) with the amplitude of the anterior less than that of the posterior - a characteristic of a normal record. In 4.2b the amplitude of the anterior is greater than that of the posterior indicating an abnormality.

The EEG of the patient in figure 4.3a shows a clear anterior posterior gradient while that of the patient in figure 4.3b shows no such differentiation as confirmed by the neurologist. A ratio of the delta band activity of the posterior channel to that of the anterior channel of the same hemisphere reflects effectively the extent of their differentiation.

As explained earlier, four channels record activity from each hemisphere. The channels that monitor the anterior and the central-parietal regions share a common electrode. Similarly, the central-parietal and posterior channels share a common electrode. As a result, large enough gradients do not exist between activities of these channel pairs even for normal recordings and detection of abnormalities on their basis may not feasible. Therefore anterior to posterior gradient in activity is measured by comparing the activities recorded by the anterior and posterior channels. Two such measures are extracted for every EEG record corresponding to the two hemispheres. The parameter is derived using the formula given below.

$$r = \text{logarithm}\left[\frac{\text{average delta activity of the posterior head region}}{\text{average delta activity of anterior head region of the same hemisphere}}\right]$$

Average delta activity in the formula above corresponds to an average over a 5 - minute period of the delta band activity. A posterior predominance of delta activity during a 5 - minute period is reflected by a ratio value greater than 1 while an anterior predominance is reflected by a ratio value less than 1. The boundary between normal and abnormal gradient values is not clear-cut and this is further discussed in the section on Statistical analysis. The front / back differentiation of every EEG record is described by 144 values, 72 corresponding to the left hemisphere and another 72 corresponding to the right hemisphere.

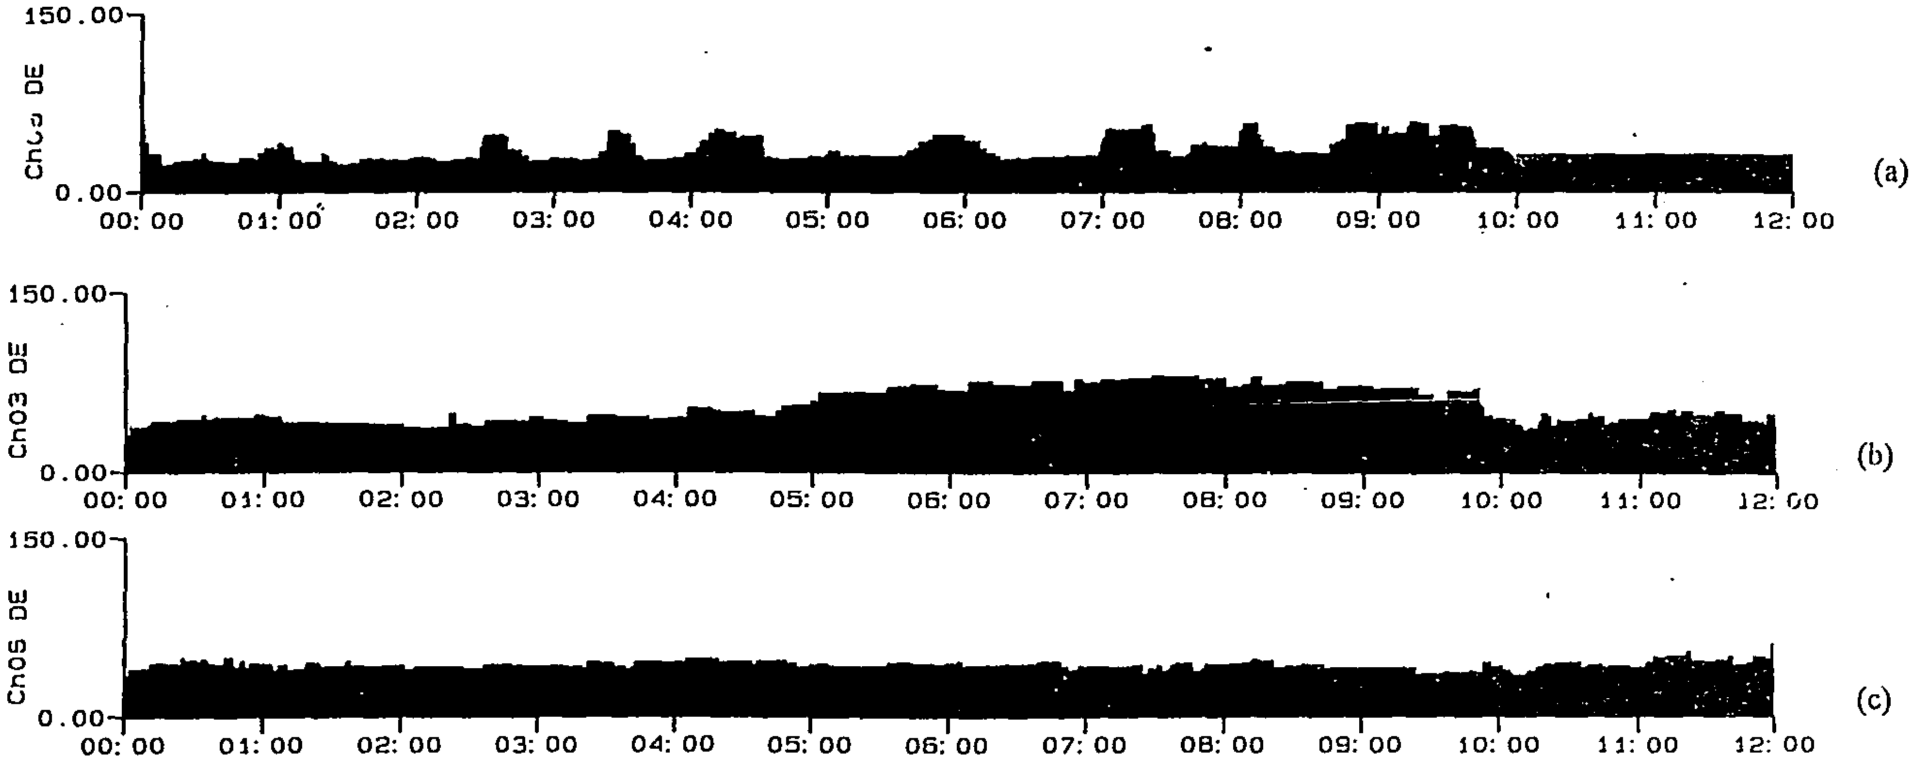

Figure 4.4(a) - Delta band activity of the posterior channel depicting normal EEG variability in the form on a good number of broad humps. 4.4b - One large hump is observable between 5 and 10 hrs. This does not characterize cycling. 4.4c - represents an absolutley flat delta band activity indicating an abnormality.

4.4 Measure for Variability

This measure assesses the extent of spontaneous cycling in an EEG record. The total and delta bands of the posterior channels of the two hemispheres are most reflective of a cycling EEG. The total band however, is susceptible to noise of a wider frequency range and it is therefore reasonable to measure variability using the delta band activity.

Figure 4.4 illustrates three posterior delta band activities subsequent to artifact rejection, one with spontaneous cycling figure 4.4a and two others figure 4.4b and figure 4.4c, without. Figure 4.4c shows an absolutely flat band activity with no signs of variability. A solitary large hump such as the one present in Figure 4.4b is not representative of a spontaneously varying EEG either. Both these examples were classified as monotonous, and hence abnormal EEGs by the neurologist. The example in Figure 4.4a shows several humps in the band activity each of which extend over a substantial period of time. Spontaneous alterations in the EEG appear in the band array as humps that are significantly higher in amplitude than the valleys that interrupt them. Each hump extends over a substantial period of time and a 6 hour recording is usually characterized by several such humps. Therefore to quantify variability, the number, duration and height of the humps are to be quantified.

Figures 4.4a, 4.4b and 4.4c can be treated as alternating signals, with components at several frequencies, superimposed on a DC signal, the DC component being the average delta band activity over the entire duration. The alternating signal of Figure (4.4c) would be characterized by very low amplitude but high frequency components (20 to 30 cycles/6 hours) since it is a very flat time series. Figure (4.4b) on the other hand, will be associated with high amplitude low frequency components (1 to 2 cycles/6 hours) as it has one extended hump. A band array representing genuine cycling, as in Figure (4.4a) would be associated with high amplitude components with 5 to 8 cycles for the entire 6 hour period. It is therefore essential to estimate roughly, both the range of frequencies and the amplitudes of the alternating components of the delta band arrays to quantify the variability. This is accomplished by computing the zero-crossing rate and the energy of the alternating signal as described below.

The DC component, which is a simple average amplitude of the delta band activity, is subtracted from the band array values of the same channel and what remains is the alternating signal. As explained earlier, frequency alternations of about 5 to 8 cycles/6 hours are to be identified. Fluctuations faster than 8 cycles/6 hours are removed by linear filtering as they do not provide information about variability of the EEG that is of prognostic significance. The band arrays of the two posterior channels from which the respective DC components have been removed, are filtered with 5 point linear averaging filters This filter replaces every point in the band array, by the average of activity within a 5-point window centered about it. This filtering operation smoothes out the humps in activity that last for less than a couple of minutes, thereby removing their influence on the zero-crossing rate.

As the name implies, the zero-crossing rate is a count of the number of times the signal crosses the time axis. This is a rough estimate of the frequency of the alternating signal. A low zero-crossing rate implies either a monotonous record or very few humps which may be artifactual. Both these cases do not represent normally cycling EEGs. A high zero-crossing rate corresponds to a band array depicting several humps. However, these humps may extend over a very short period of time or may have a very low maximum amplitude and in neither case corresponds to a normally cycling EEG. The extent of the humps is ascertained by computing the energy of change of the alternating signal using the formula below.

$$e = \log \left[\frac{\sum_{i=0}^{i=120} x_i^2}{(\text{average delta band activity})^2} \right]$$

where "x_i" denotes the amplitude of each point of the band array after the DC component has been removed from it. The energy of change is measured as a fraction of the average delta band energy. The boundary for absolute fluctuations between normal and abnormal EEGs is a function of the background amplitude of the EEG itself. Relative fluctuations on the other hand, are patient independent. This normalization facilitates measurement of relative fluctuation of the EEG with respect to its average energy rather than absolute fluctuations. The energy of change is a function of the extent and height of the humps. A low energy of change corresponds to a short hump, short either in duration or in extent or both. These do not represent cycling EEGs. A high energy of change could correspond either to a single large hump or several humps, the former will have a low zero-crossing rate while the latter a large zero-crossing rate.

From the discussions above it is evident that in order to quantify variability of an EEG, both zero-crossing rate and energy of change are essential. A zero-crossing rate and an energy of change measure are extracted from the posterior channel of each of the two hemispheres for every 6 hour EEG record and used for further interpretation.

Eighteen quantitative measures describe 6 hours of raw EEG data after feature extraction. Further interpretation is carried out by statistical methods discussed in chapter 5.

5. STATISTICAL ANALYSIS

Upon feature extraction, the 500 KB activity file is reduced to a set of 1012 data values, 576 of which describe its amplitude, 288 others describe the symmetry in activity between the left and right hemispheres of the brain, 144 describe the anterior to posterior gradient in activity and 4 describe the variability of the EEG. In the previous section on feature extraction, several phrases such as "low amplitude", "high frequency", "values close to 1", etc. were used without qualification. Where does the amplitude threshold lie between a normal and depressed EEG? Below what value is the measure for anterior posterior differentiation considered to represent an abnormal EEG? Where exactly does the boundary between normal and abnormal categories lie for all of the measures discussed in the previous section? These problems must be solved to proceed further with the task of interpreting the EEG.

The expert is unable to help with the definition of the quantitative boundaries for the various measures extracted as he does not relate to these quantitative quantities. He does not use them to visually interpret the EEG and no known boundaries or threshold values are available for these features from other sources either.

Interpretation of quantitative EEG measures, as explained earlier, is usually performed by statistical analyses based on a normal control population. The selection of a representative normal control population is crucial for good system performance. For this project, of the EEGs recorded, a group of eight post-cardiac surgery patients had normal post-operative EEG recordings and normal short-term neurologic outcome and are chosen to comprise the control population. Long-term EEG recordings of these patients lasting between 18 and 24 hours in duration evaluated as being normal by the neurologist serve as the control population. The measures describing amplitude, symmetry, anterior / posterior differentiation and variability are extracted from each of these recordings. Upon extraction of these measures it was observed that three of these patients showed signs of asymmetry in activity between hemispheres and one exhibited minimal anterior-posterior differentiation. The abnormalities were very minimal, however. This is confirmed by the neurologist on re-evaluation of the EEG visually. All other measures of these records were found to be normal. Abnormal features of records were dropped from the corresponding control population. Therefore, the control populations of each of the amplitude and variability measures consists of eight patients, control populations of the measures of symmetry consist of five patients and anterior / posterior differentiation consist of seven patients.

Cumulative frequency distributions of the control population are constructed for each of the eighteen measures . The characteristic measures of a new patient whose EEG is to be interpreted can then be statistically compared with the distributions of the control population. A measure of similarity between them will indicate the level of normality of the new patient. The construction of the distributions and the statistical tests are discussed below.

5.1 Population Distributions

As seen in the previous chapter, all features extracted from the raw EEG are defined as the logarithm of a certain estimate rather than the estimate itself. For instance, symmetry measures are defined as a logarithm of the ratio rather than as the ratio itself, amplitude measures are defined as the logarithm of average amplitudes rather than as the average amplitudes themselves. The distributions of the estimates such as average amplitude do not follow the characteristics of a normal distribution. This may either be due to the biological mechanism generating the EEG or due to rigid boundaries associated with the estimates themselves. To facilitate transformation of these distributions toward the normal distribution, the logarithm of the estimate is used as the parametric definition (Gasser et.al., 1982). The construction of the distributions for all of the 18 parameters is discussed below.

5.1.1 Amplitude Normality

Data values extracted from the EEGs of the eight "normal" patients extracted at 5 minute intervals is used to construct frequency distributions of each of the eight measures of amplitude corresponding to the eight channels recorded. Each of these patients contribute between 216 and 288 data points to the distributions (recordings last between 18 and 24 hrs) amounting to a total of 1738 points. The distributions appear in Figure 5.1. Values of skew and normal tests lie between -1 and +1 confirming that all amplitude distributions are within acceptable limits of the normal distribution. Averages and standard deviations of the various distributions appear in Table 5.1.

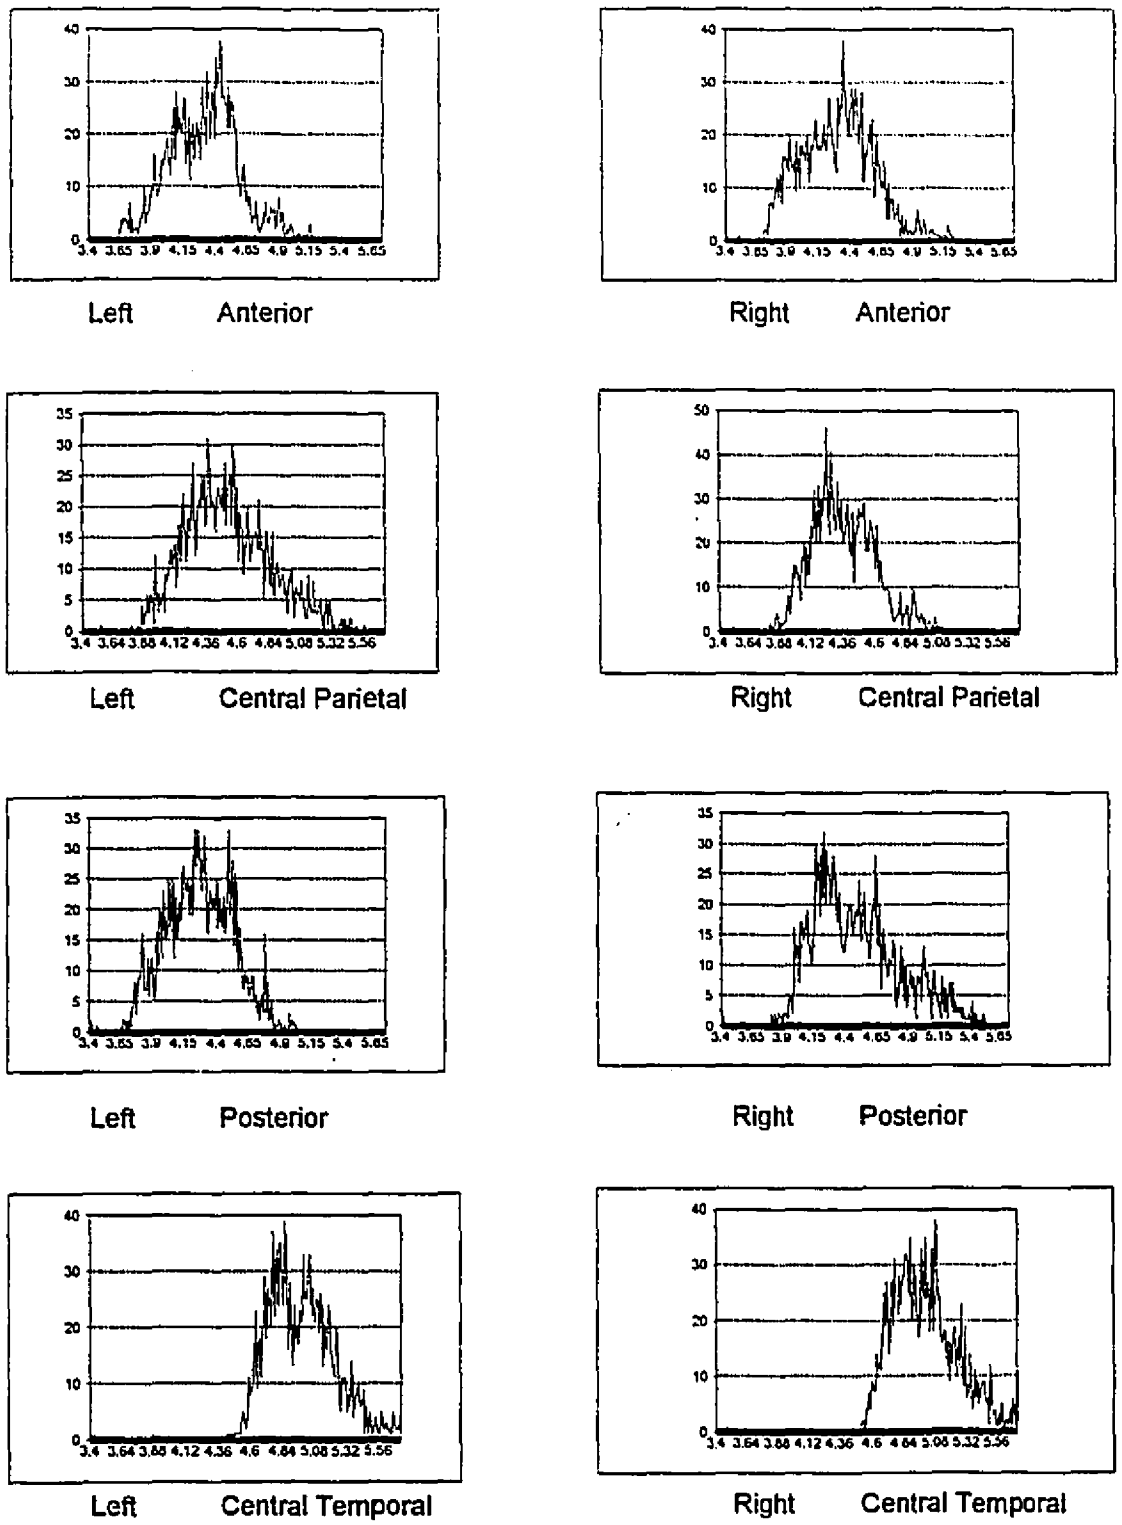

Figure 5.1 - Distributions of the amplitude measures for the control population.

Head region (Channel)	Average	Std. Deviation
Left Anterior	4.30	0.28
Left Central Parietal	4.28	0.29
Left Posterior	4.50	0.33
Right Anterior	4.34	0.24
Right Central Parietal	4.25	0.26
Right Posterior	4.48	0.35
Left Central Temporal	4.97	0.25
Right Central Temporal	4.99	0.26

Table 5.1

It was stated earlier that EEGs of normal children exhibit a posterior anterior gradient in the amplitude of activity with decreasing values in the anterior direction. The average values of activity presented in the table exhibit this posterior anterior gradient in the two hemispheres demonstrating this aspect. The last two amplitude measures represent the central temporal head regions. The inter-electrode distance for these two channels is twice that of the other channels and hence the amplitude averages are higher than those of the other channels.

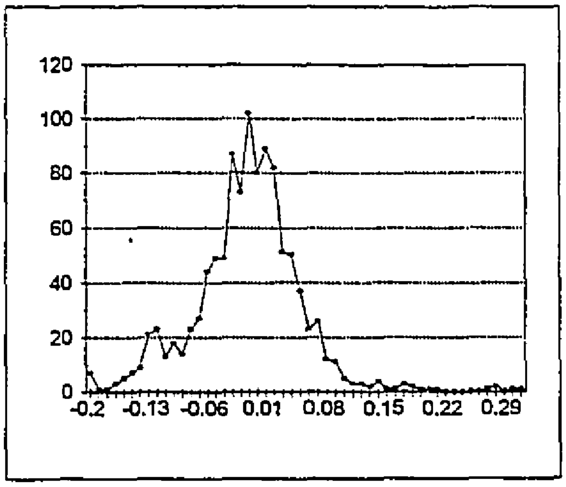

Symmetry measure for anterior head regions

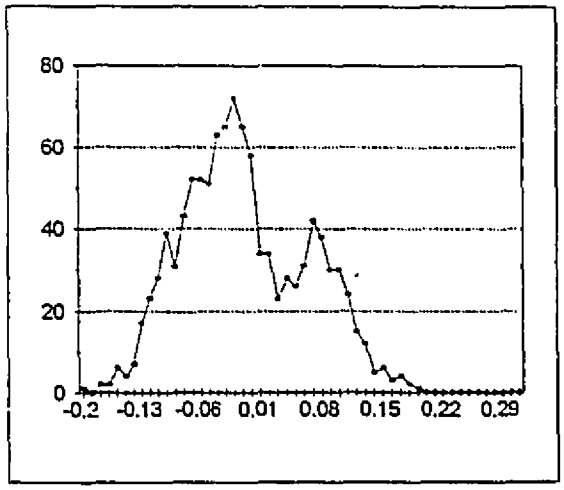

Symmetry measure for central parietal head regions

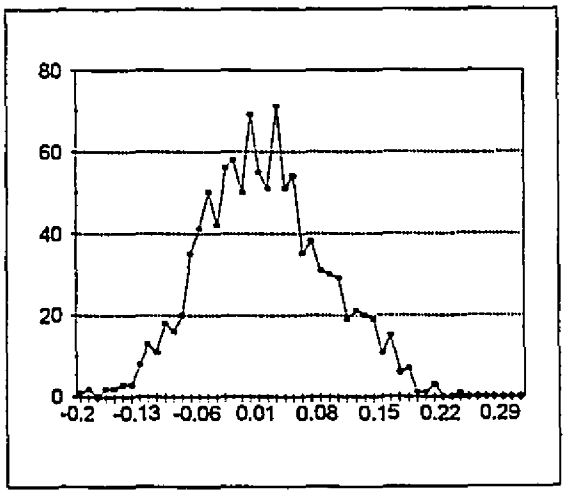

Symmetry measure for the posterior head regions

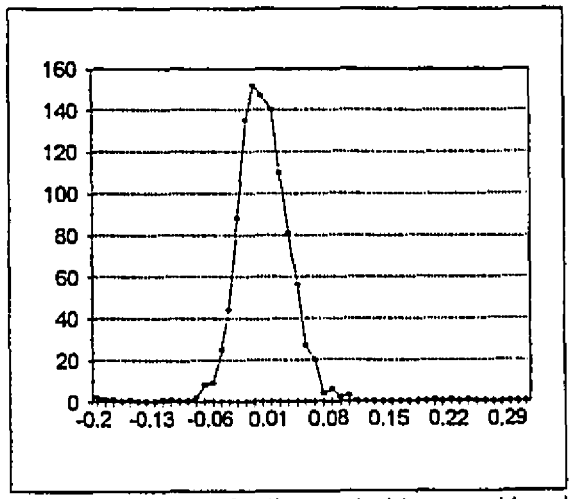

Symmetry measure for the central temporal head regions

Figure 5.2 - Distributions for the symmetry measures corresponding to the various head regions.

5.1.2 Symmetry

Data values from five patients evaluated normal by the neurologist are used to construct the four frequency distributions for the left / right symmetry measures. These distributions too are within acceptable limits of normal distributions. These distributions consisted of 1085 points and are illustrated in Figure 5.2 above. The averages and standard deviations of the four distributions are presented in Table 5.2.

Head region	Average	Std. Deviation
Anterior	0.99	0.07
Central Parietal	0.99	0.08
Posterior	1.01	0.07
Central Temporal	1.00	0.03

Table 5.2

As is evident from Figure 5.2 and the standard deviations listed in the table, the distribution from the channel pair that monitors the central-temporal head region has a much smaller range as compared to the other distributions. The two channels of this pair have a common electrode i.e. the channel on the left hemisphere measures voltage between electrodes T_3 and C_z while the channel on the right hemisphere measures voltage between electrodes T_4 and C_z. Hence the difference in activity between hemispheres is much less as compared to the other channel pairs.

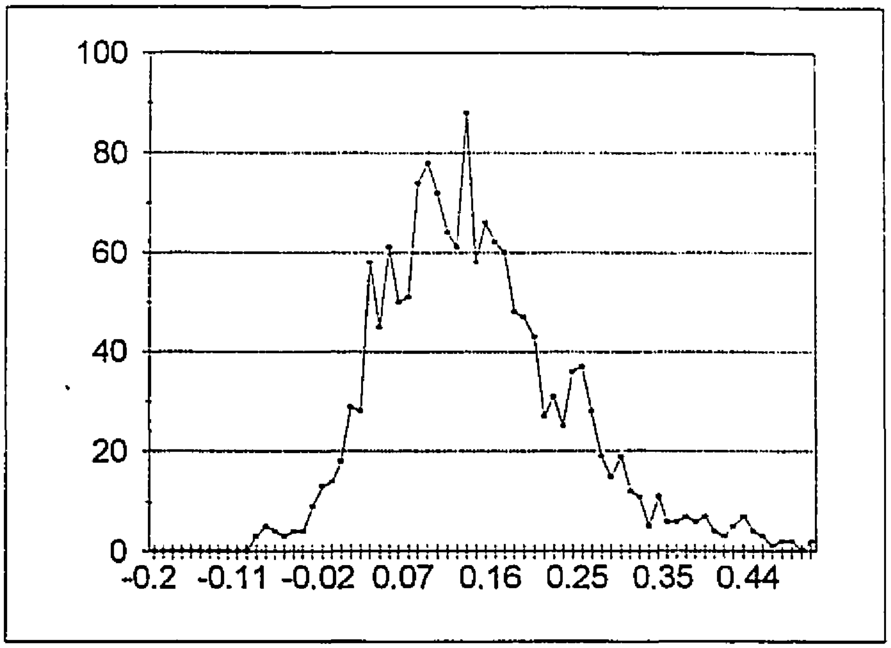

Left Hemisphere

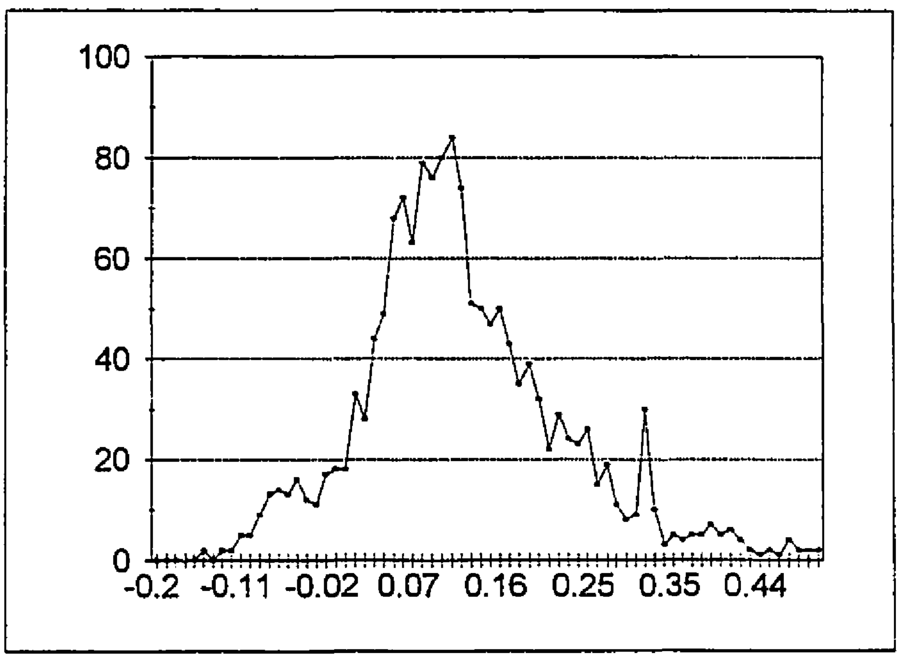

Right Hemisphere

Figure 5.3 - Distributions of the front/back differentiation measures for the control population.

5.1.3 Front / Back Differentiation

Seven patients comprise the frequency distributions characterizing front / back differentiation of the control population. Skew and nomal tests verify that these distributions consisting of 1569 data values are within acceptable bounds of the normal distribution. The averages and standard deviations of the two curves shown in Figure 5.3 corresponding to the two hemispheres is presented in Table 5.3. The posterior to anterior activity ratios in young children may reach values as high as 4:1 (Slater and Torres, 1979).This explains the large standard deviations of these distributions as compared to those representing the symmetry in activity..

Hemisphere	Average	Std. Deviation
Left	1.14	0.10
Right	1.13	0.12

Table 5.3

5.1.4 Variability

Twenty one data points from eight patients constitute the four frequency distributions - two for zero-crossing rate and two for energy change, corresponding to the two hemispheres. Unlike all other parameter values that are extracted once every five minutes, these parameters are extracted only once every six hours and hence the small sample set. The averages and standard deviations of the four distributions represented in Figure 5.4 is presented in Table 5.4.

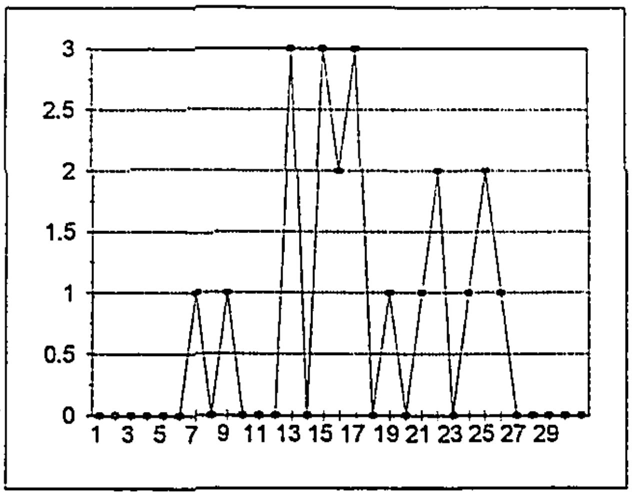

Zero-crossing rate - Left Hemisphere

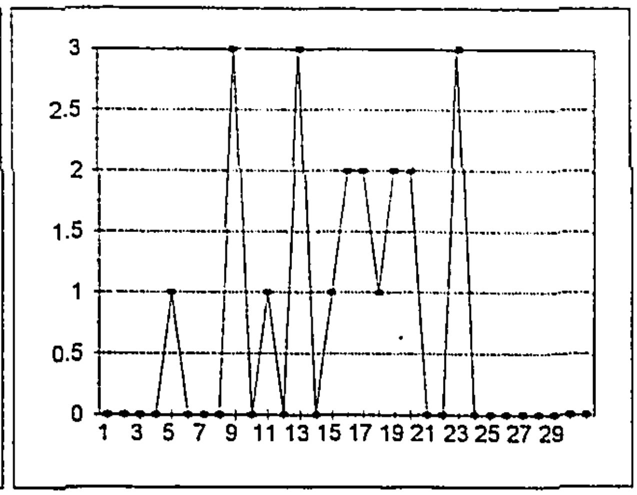

Zero-crossing rate - Right Hemisphere

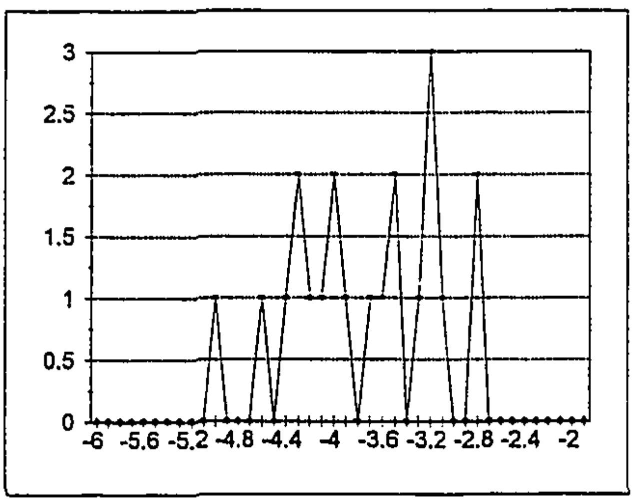

Logarithm of energy change - Left Hemisphere

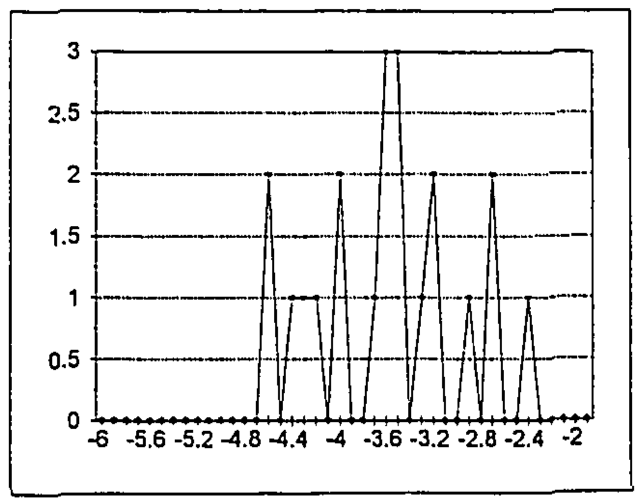

Logarithm of energy change - Right Hemisphere

Figure 5.4 - Distributions of measures of variability of the control population

Parameter	Average	Std. Deviation
Zero-crossing rate - Left	17.48	1.12
Zero-crossing rate - Right	15.62	1.09
Log energy of change - Left	-3.80	0.13
Log energy of change - Right	-3.64	0.13

Table 5.4

5.2 Statistical tests

Software analyzes six hours of EEG data and extracts eighteen features from it which amounts to 1012 data points. The aim at this point is to arrive at the level of normality of each of these features by analyzing the corresponding sample sets of data. Each of the eight amplitude measures, four symmetry measures and two front / back differentiation measures are associated with 72 data values each. The amplitude and ratio distributions of a new "normal" patient could be expected to be quite similar to the corresponding distributions of the control population. An abnormal EEG, on the other hand, would have distributions quite different from that of the control group. This is evident from Figure 5.5 which illustrates the symmetry measure that compares the posterior head regions in the three cases, (a) control population distribution, (b) distribution of a normal patient and (c) distribution of an abnormal patient.

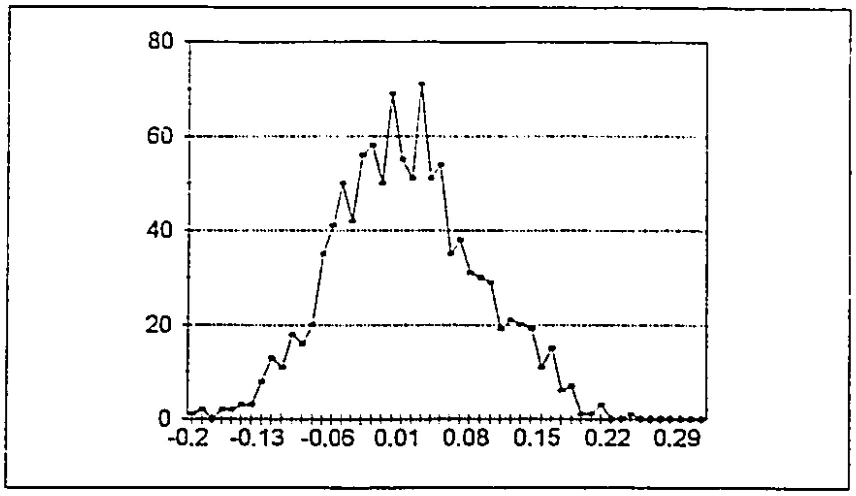

(a)

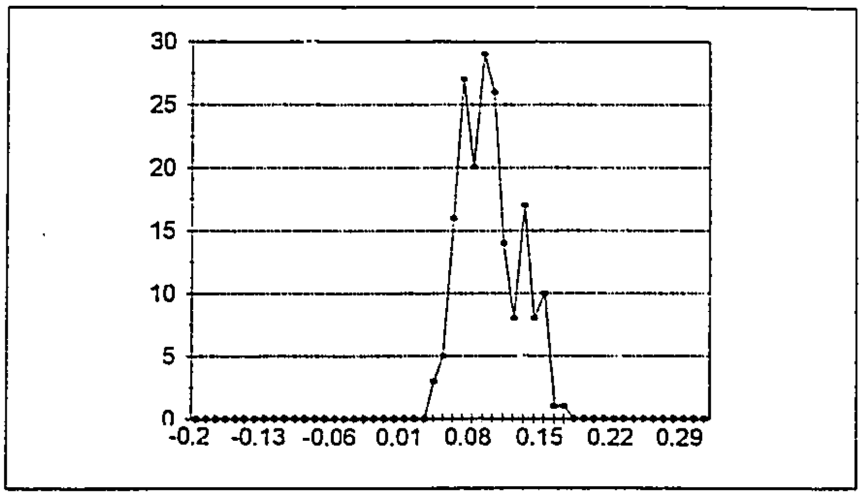

(b)

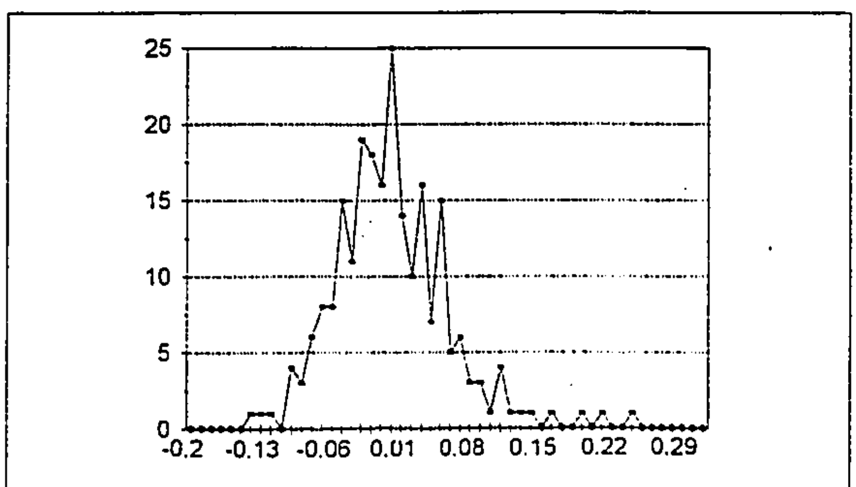

(c)

Figure 5.5 - Symmetry distributions of the posterior head regions of (a) control popula
(b) an abnormal patient with relative-depression on the right hemisphere a
(c) a normal patient. The means of the distributions are (a) 0.01, (b) 0.10
(c) 0.00.

The distributions in (a) and (b) are similar in terms of the value corresponding to the peak, the range of the distribution, etc. while (c) is quite different from both these distributions.

A measure of the degree of similarity between the distributions of the control population and the EEG being analyzed would give an estimate of the normality of the feature concerned. An appropriate definition of the "measure of similarity" is crucial to facilitate extraction of the required information from the sample sets. For all of the fourteen features i.e. eight amplitude and six ratio features, that provide such a sample set of data, an overall measure of normality that ignores individual fluctuations of the data values is desirable. A comparison of the means of the data set and the corresponding control distribution would serve as a suitable measure of similarity. The basic premise is that similar distributions have similar means and arithmetic means provide a good estimate of overall data trend. This is clearly evident in Figure 5.5, where the means of (a) and (c) are at 0.01 and 0.00 respectively, while the mean of (c) is far away at 0.10 indicating an abnormality. A t-statistic measures the level of similarity between distributions by comparing their means. This could be used to compare distributions of new patients with the control population to extract information about their normality. This is further discussed in section 5.2.1 below.

5.2.1 T - Statistic

Very often inferential statistics is sought in order to make decisions about the value of a parameter such as a population mean or population proportion. Often cross-correlation techniques are used to compare populations. However, when the similarity of the means of distributions is to be compared rather than the forms of the distributions themselves hypothesis testing is preferred. Hypothesis testing first involves making a suitable hypothesis. Tests are then carried out on the population and the hypothesis is either accepted or rejected on the basis of the values obtained for the test-statistic applied. One test-statistic often used for hypothesis testing is the t-statistic.

Suppose that independent random samples of sizes n_1 and n_2 are taken from two normally distributed populations with means μ_1 and μ_2 respectively. Let x_1 and s_1 represent the sample mean and standard deviation of the sample from population with mean μ_1 and x_2 and s_2 the sample mean and standard deviation of the sample from the population with mean μ_2. Then the random variable

$$t = \frac{(x_1 - x_2) - (\mu_1 - \mu_2)}{\sqrt{(s_1^2/n_1 + s_2^2/n_2)}} \qquad .....5.1$$

has approximately the t-distribution with degrees of freedom given by

$$df = \frac{[(s_1^2/n_1 + s_2^2/n_2)]^2}{\dfrac{(s_1^2/n_1)^2}{n_1 - 1} + \dfrac{(s_2^2/n_2)^2}{n_2 - 1}} \qquad .....5.2$$

rounded to the nearest integer (Weiss and Hassett, 1991). The t curve is symmetric about '0' and extends indefinitely in both directions. As with the normal curve, the area under a t-curve is equal to 1 and it approaches a normal curve as the number of degrees of freedom get larger.

For the test of a hypothesis $H : \mu_1 = \mu_2$, therefore the random variable,

$$t = \frac{(x1 - x2)}{\sqrt{(s_1^2/n_1 + s_2^2/n_2)}} \qquad .....5.3$$

may be used as the test statistic and is called a t-statistic. A t-statistic value of 0 implies that the hypothesis is absolutely true. If the t-statistic is > 0 then the area under the t-curve to the right of it gives the probability of the truth of the hypothesis. Similarly, the area under the t-curve to the left of a t-statistic value less than 0 gives the probability of the truth of the corresponding hypothesis. The t-statistic essentially measures the normalized distance between the means of the two distributions and hence greater its absolute value lesser is the probability of the two means being equal.

The probability value associated with the t-statistic value corresponding to any two distributions provides a measure of similarity between the arithmetic means of the two distributions being compared. This is exactly what is desired while comparing parameter distributions of the control population and a new patient, and hence the t-statistic maybe used to perform the task. The t-test maybe applied only to distributions that satisfy two conditions: Independent Samples and Normal Populations. Two samples are independent if the samples selected from one of the populations has no effect on those selected from the other. For the problem at hand, the two samples are absolutely independent of each other since they are from different patients and hence do not influence each other. The distributions of the control populations themselves are within limits of a normal distribution as discussed in section 5.1. The various distributions of most patients have also been found to conform with the characteristics of a normal distribution. The t-test statistic may therefore be applied to obtain information about similarity of these distributions.

Subsequent to feature extraction from an EEG to be interpreted, the means and standard deviations of the 14 parameters: 8 amplitudes, 4 symmetries and 2 front back differentiations, are calculated. The t-statistic is then calculated as below.

$$t_{stat} = (x_1 - x_2) / \sqrt{(\sigma_1^2/n_1 + \sigma_2^2/n_2)} \qquad .....5.4$$

where,

x_1 = mean of a parameter distribution for the EEG being evaluated

x_2 = mean of the control population distribution for the same parameter

σ_1 = standard deviation of the parameter distribution for the EEG being evaluated

σ_2 = standard deviation of the control population distribution for the same parameter

n_1 = number of samples for the EEG evaluated

n_2 = number of samples in the control population

The number of degrees of freedom of the t-curve associated with the random variable is defined above in Equation 5.2. Since n_1 is very large (> 1000), $df \approx (n_2 - 1) = 71$. Beyond df = 30, the t-curve converges to the normal distribution (Weiss and Hassett, 1991) and therefore the probability values corresponding to the various t-statistic values are read off the normal distribution tables. If the t-statistic > 0 then the mean of the control population is less than that of the EEG being analyzed for the particular parameter. Similarly, if the t-statistic < 0 the control population mean is greater than that of the EEG being evaluated.

The range of t-statistic values extends from $-\infty$ to $+\infty$. Its value represents the level of normality of the measure associated with it. For amplitude measures, depression in EEG activity is an abnormality and amplitudes greater than or equal to that of the control population represent normal EEG. Therefore a t-statistic value greater than or equal to zero irrespective of its magnitude implies absolute amplitude normality for the corresponding amplitude parameter. Therefore the probability of normality is assigned a value of 1.0. A t-statistic < 0 corresponds to a depression in activity of the EEG being evaluated and the level of depression is a function of the absolute value of the t-statistic. It has been observed that a t-statistic value < -50 indicates a very severe depression and is therefore assigned a 0.0 probability of amplitude normality. Therefore the t-statistic value range from -50 to 0 is mapped linearly onto a probability range from 0.0 to 1.0. The probability of normality of an amplitude parameter is thus defined as,

$$\text{Prob. of amp. normality} = \begin{cases} 1.0, & \text{t-value} \geq 0.0 \\ (50 + \text{t-value})/\,50, & -50 < \text{t-value} < 0 \\ 0.0, & \text{t-value} \leq -50 \end{cases} \qquad \ldots\ldots 5.5$$

The symmetry parameters compare activity of the left hemisphere to that of the right hemisphere. For normal EEGs the activity between hemispheres is symmetrical as is evident from the means of the control population distributions for these parameters. A t-statistic value > 0 for a symmetry parameter therefore implies that the mean of the distribution for the EEG being analyzed is $>$ mean of the control population which is ≈ 1.0. This in turn implies that the right hemisphere is relatively depressed as compared to the left hemisphere. A t statistic < 0 on the other hand implies a relative left hemispheric depression. Unlike in the case of parameters for amplitude, for symmetry, a t-value different from 0.0 represents an abnormality and the level of asymmetry is a function of the magnitude of the t-statistic. It has been observed that a t-statistic value with magnitude > 50 indicates a very severe asymmetry and is therefore assigned a probability of 0.0. A t-statistic value equal to 0.0 indicates absolute symmetry and is hence associated with a probability of normality of 1.0. Thus each of the t-value ranges from -50 to 0 and from 0 to 50 is linearly mapped on to a probability range of 0.0 to 1.0.

$$\text{Prob. of norm. of symmetry} = \begin{cases} 0.0, & \text{t-value} \leq -50 \\ (50 - |\text{t-value}|)/50, & -50 < \text{t-value} < 50 \\ 0.0, & \text{t-value} \geq 50 \end{cases} \qquad \text{.....5.6}$$

In the case of front / back differentiation measures, a front back ratio less than the mean of the control population distribution indicates an insufficiency in differentiation in activity between posterior and anterior head regions and the level of abnormality itself is a function of the magnitude of the mean of the parameter of the EEG being analyzed. A parameter value greater than the mean of the control distribution indicates a greater differentiation than average and this does not represent an abnormality. In terms of t-statistic values therefore, t = or > 0 is considered normal and the probability of normality associated with it is set to 1.0. Once again it was observed that a t-value < -50 indicated a very severe lack of differentiation and the probability of normality associated with it is set to

0.0. The probability of normality for a t-value between -50 and 0.0 is linearly mapped on to a 0.0 to 1.0 probability range as below.

$$\text{Prob. of normality of diff.} = \begin{cases} 1.0, & \text{t-value} \geq 0.0 \\ (50 + \text{t-value})/\,50, & -50 < \text{t-value} < 0 \\ 0.0, & \text{t-value} \leq -50 \end{cases} \qquad \ldots.. \; 5.7$$

A t-statistic is calculated for each of the eight amplitude parameters, four left / right symmetry parameters and two front / back differentiation parameters. The various parameter values are thus, mapped on to a 0.0 to 1.0 probability range, each reflecting the probability of normality of the feature of the EEG that they measure, i.e. amplitude of a particular channel or left/right symmetry of certain region of the brain, or front to back differentiation.

5.2.2 Assessment of Variability

Unlike the amplitude and ratio characteristics which are summarized by data distributions for each new incoming patient, variability is characterized by four values each one summarizing a different aspect of variability. As it is not a distribution, the t-test may not be used to assess it. The variability of every six hour EEG record analyzed is characterized by two sets of measures corresponding to the two hemispheres, each set consisting of a zero-crossing rate and an energy of change measure. As discussed earlier, large enough values for both measures corresponding to a particular hemisphere indicates presence of variability in activity originating from that hemisphere. The data values themselves can be assessed for normality by comparing them with the corresponding frequency distributions of the control population. A measure of the relative standing of the data values within the distributions of the control population may be treated as an estimate of the level of normality of the measure itself. For instance, a zero-crossing rate greater than all sample points of the control population implies a 100% normality in the zero-crossing rate. A value less than all sample points of the control population would imply an absolute abnormality and one equal to the mean would indicate a 50% probability of normality of the zero-crossing rate. Such a measure of relative standing could be arrived at by computing the z-score of the parameter value as discussed in the section below.

Any normal curve is defined by its two parameters : μ, the mean about which it is symmetrical and σ its standard deviation. The total area under such a curve is equal to 1 and most of it lies between $\mu - 3\sigma$ and $\mu + 3\sigma$. The area under a normal curve with parameters μ and σ that lies between $x = a$ and $x = b$ is equal to the area under the standard normal curve that lies between

$$z = \frac{a - \mu}{\sigma} \quad \text{and} \quad z = \frac{b - \mu}{\sigma} \qquad \text{.....5.8}$$

A standard normal curve has a mean of 0 and standard deviation equal to 1. The area under a normal curve with parameters μ and σ that lies to the left of a given value x can be found in a similar way by first computing the corresponding z-score as

$$z = \frac{x - \mu}{\sigma} \qquad \text{.....5.9}$$

The area to the left of the z-score under the standard normal curve is equal to the required area and can be read from standard tables.

While assessing the zero-crossing rate and energy of change parameters, the relative position of these with respect to the control population data set is desired. This amounts to finding the area to the left of the parameter value under the corresponding control population distribution. Since all four distributions of the control population, 2 zero-crossing rates and 2 energies of change, fall within limits of a normal curve the z-score method described above maybe used to evaluate the areas required.

The z-score of the zero-crossing rate and energy of change parameters of an EEG being analyzed are computed using Equation 5.6. The area under the standard normal curve to the left of z-score gives a measure of normality of the corresponding parameter. Therefore, subsequent to computation of the z-score, the two zero-crossing rate parameters and the two energy of change parameters are reduced to four probabilities of normality each denoting the level of normality of the corresponding characteristic of the EEG.

5.3 Data Reduction

Subsequent to statistical analysis, an EEG is characterized by 18 probability measures lying between 0.0 and 1.0. From the information provided by these measures the level of normality of the EEG is to be deciphered. This can be performed by an automated learning machine which is trained by a representative set of examples. However, the efficiency of such a system is dependent on the number of variables provided to it and is better if fewer variables are input to the system. Therefore it would be useful to minimize the number of variables provided to such a learning device. With the automated EEG monitor, the eighteen measures may be reduced to four by integrating the information from the eight amplitude measures as one, the four symmetry measures as another, the two differentiation measures as a third feature and the variability measures as a fourth feature.

Of the eighteen measures generated from an EEG to be evaluated, eight characterize the amplitude of the EEG activity corresponding to the eight head regions. To obtain information about the overall normality of EEG amplitude, the information from these eight parameters should be suitably aggregated. Depression in several of the eight head regions is indicative of a worse prognosis than depression in just one region. To incorporate this gradation into the overall amplitude normality parameter, the eight amplitude normality parameters are averaged to provide one parameter of overall amplitude normality. Similarly, the four symmetry parameters that characterize symmetry in activity between the two hemispheres in the various head regions are averaged to provide a parameter of left / right symmetry. The front / back differentiation parameters corresponding to the two hemispheres are averaged to provide the front / back differentiation parameter. An EEG is said to be variable if it is characterized by a good zero-crossing rate AND a good-energy of change. A high value for one of them with a low value for another is as bad a prognosis as a low value for both parameters. The normality of variability is therefore indicated by the magnitude of the lower of the two parameter values. Therefore the variability parameter can be defined as

variability = min { probability of normality of zero-crossing rate, probability of normality of energy of change}

..... 5.10

The two variability parameters corresponding to the two hemispheres are then averaged to provide an overall measure of variability of the EEG.

Subsequent to statistical analysis and data reduction therefore, the EEG is characterized by four probability measures that speak about the amplitude, the left / right symmetry, the front / back differentiation and the variability of the EEG. This information is to be further interpreted to classify the EEG as being normal, mildly abnormal, moderately abnormal or severely abnormal. This is done by automated machine learning from prior examples of such classification done by the neurologist. This is discussed in the following chapter.

6. KNOWLEDGE-BASED EXPERT SYSTEM

The aim at this stage is to develop a system which accepts the four indices, i.e. amplitude, symmetry, front / back differentiation and variability measures of an EEG as inputs and classifies the EEG as one of the four categories namely, normal, mildly abnormal, moderately abnormal or severely abnormal (Figure 6.1).

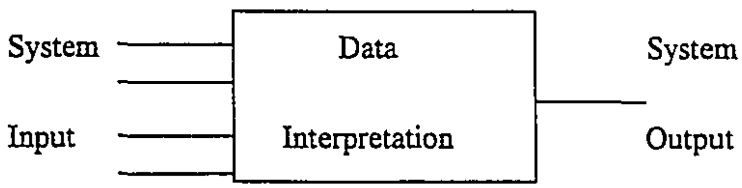

Figure 6.1

Such a system should mimic the neurologist's criteria for classification. The data interpretation in the figure above should correlate with the neurologist's decision-making about the level of EEG abnormality based on its various characteristics. However, since the neurologist interprets these EEGs visually and the characteristics themselves are not quantified, knowledge acquisition by direct consultation with the neurologist is not feasible. The system could acquire knowledge about the data interpretation from examples classified by the neurologist. The EEGs classified by the neurologist can be quantified and characterized by the four indices which would form the inputs to the system. The neurologists classification would serve as the corresponding output. Data interpretation is intrinsically codified in the input-output relationship of these examples and suitable automated techniques can be designed to extract the required knowledge. This chapter discusses the concept of automated machine-learning from examples and describes the knowledge-based expert system built to perform the task of EEG classification.

6.1. Machine learning from examples

One way to teach a system how to perform a task is by presenting it with examples of how it should behave. The system treats the examples as highly specific pieces of information which are then transformed into more general pieces of knowledge that maybe used effectively by the performance element. Simon and Lea (1974) call the space of possible training instances the instance space and the space of possible general rules the rule space. Then the system should search the rule space to come up with rules that would describe the behavior of instances in the instance space. An intelligent learning system would select its own instances to resolve ambiguities about rules in the rule space. Therefore, if the program were unsure whether all dogs have four legs, it might search the instance space to spot animals with different numbers of legs to see which ones are dogs.

Any system that learns from experience should address three component problems :

- *Aggregation*. The learner should identify the basic objects that constitute the instances from which he will learn. He should first and foremost separate signal from noise.
- *Clustering*. The learner must identify which objects or events should be grouped together into a class. He should develop extensional definitions for concepts based on those data.
- *Characterization*. The learner must formulate some general description or hypothesis that characterizes instances of the concept. In other words, he must generate an intensional definition of the concept.

The task of learning from examples maybe viewed as a degenerate case of the general learning task as the tutor solves the problems of aggregation and clustering by providing the learner with positive and negative examples of the concept to be learned. Therefore the task of learning from examples maybe viewed as simplified characterization, since this is the only component of learning that must be addressed. This simplification has proved quite useful to learning researchers, and many of the characterization methods that were initially developed for the task of learning from examples have been successfully transferred to more complex problems.

6.2. Concept Learning

A great many programs have been developed that learn one or a few concepts from instances. A concept maybe called a predicate described in some language, which when applied to a positive instance is TRUE and when applied to a negative instance is FALSE. A concept partitions the instance space into positive and negative subsets. Thus, given a representative language for concepts and a set of positive and negative training instances, the concept learning problem itself is to find a unique description in the rule-space that would encompass all the positive instances and none of the negative instances. Once the concept is learned, the system is ready to classify new unknown instances as positive or negative instances of the concept.

6.2.1. Concept Learning by Generalization

Concept learning from instances has primarily incorporated techniques of generalization. In this paradigm, the system initially assumes that all aspects of the first positive instance are relevant to the concept and systematically removes conditions as they fail to occur in new instances. The basic premise is that one can arrive at the definition of a concept by determining those features that are held in common by a set of positive examples. The performance of such a system is primarily dependent on two features - the manner in which the rule space is searched and the way in which negative instances are made use of.

Many of the early systems used a depth-first search and often did not use the negative instances at all. Bruner, et. al.(1956), developed a system to learn concepts that could be represented as attribute-value pairs. This approach started with a positive instance and initially all attributes were assumed relevant to the concept. Those attributes whose alteration led to a negative instance were retained and all others were eliminated. Although this strategy works well with conjunctive concepts connected by an AND, it cannot be used to learn disjunctive concepts connected by OR. Winston (1970), extended this approach further to learn more complex representations. In this case, with the introduction of new positive instances, in some cases, more than one generalization was possible. Since the system performed a depth-first search through the rule space, it needed the ability to backtrack and this is when negative instances come into play. A misclassification of a negative instance on the part of the system initiates backtracking to a prior more specific definition.

An advantage of breadth-first search strategies is that they need not retain positive instances of a concept since they need never backtrack. Hayes and McDermott (1977), have incorporated this in their system. In this technique, if a generalization was formed that covered non examples in addition to positive examples of the concept, it was considered overly general and that hypothesis was dropped from consideration. This requires that negative instances be stored to prevent over - generalizations.

The Version Space technique developed by Mitchell (1977) dispenses with such a need too. In this technique, in addition to maintaining a set of generalizations or maximally specific versions (MSVs) of a concept, a set of maximally general versions (MGVs) are also maintained. The MGV starts off with the most general point in the rule space, the null description, which places no constraints on the training instances and thus describes anything. The MSV starts off with the first positive instance itself represented in the same language as the rule space. New positive instances lead to a more general MSV due to fewer conditions which correlates with the generalization in the previous methods. New negative instances lead to more specific MGVs with additional conditions. An MSV hypothesis that successfully matches a negative instance is removed. Similarly, an MGV hypothesis that fails to match a positive instance is removed. This algorithm known as the candidate - elimination algorithm is a least-commitment algorithm since it does not modify the version space, a set of all plausible hypotheses, until it is forced to do so by the training instances. Positive instances force the program to generalize - thus, very specific concept descriptions are removed from the version space. Negative instances force the system to specialize, so very general concept descriptions are removed from the version space. Thus, the version space gradually shrinks until only the desired concept description remains.

The basic approach of learning by generalization has a number of drawbacks that limit its value as a path to knowledge acquisition.

1. Since the method examines features that are held in common by positive instances, it tends to over generalize when confronted with examples of disjunctive rules.
2. Generalization based learning systems have difficulty handling erroneous data. If even one of these examples is faulty, the entire learning sequence is thrown into confusion. Therefore, efficient error-handling procedures are to be built in.
3. Programs that learn through generalizations have difficulty responding to an environment in which the conditions predicting an event actually change over time.

Recently, more work has been done in this field to alleviate the above-mentioned shortcomings. Michalski and Larson (1978), have suggested the use of only the most-representative training instances and a more robust learning algorithm. Leng and Buchanan (1992), have improved the performance of their inductive inferencing by letting the system generate new terms e.g. ordering (> <) concepts. Clymer and others (1992), have built a system that is context sensitive by including a measure of effectiveness for the various parameters.

6.3. Neural Networks

Neural Networks are another form of automated machine learning from examples. Here, a functional form is assumed for the unknown system. This functional form has a vector of parameters, **w**, that must be determined from the instances. The most popular form assumed for the unknown system is a linear function. With more complex systems multiple layers are incorporated and the system creates new features in the hidden layers.

If the rule space is considered to comprise all of the possible parameter vectors, then finding these parameter vectors is a task of searching the rule space to describe the behavior of the training instances. Inductive learning machines work exactly on this principle and hence neural networks maybe likened to them. Unlike neural nets, constructive induction methods require either some prior knowledge of potentially useful features or ways to build them. This could be a liability in a truly knowledge-free domain but an advantage when a little is available since it can be directly encoded.

6.4 Knowledge-Based Expert System for EEG Classification

6.4.1 Neural Network or Induction by Generalization?

As discussed earlier, a Knowledge-Based Expert System (KBES) is to be designed to perform automatic interpretation of the EEG. Such a system should aggregate information provided by the four characteristic parameters of the EEG and classify it amongst one of four categories. The knowledge about how the information is to be aggregated could be acquired from prior examples in two different ways. This could be treated as a neural network with the four indices as the inputs and the four classifications may be quantified and used as the output of the system. The network could then be trained with a large set of examples. An alternative approach would be inductive concept learning by generalization. The system would be provided with the instance space consisting of the various examples and a rule space with an operative language that would be used by the system to describe the rules that it would generate. The system would create the rules by making generalizations based on its positive and negative examples.

As explained in section 6.3 a neural network is advantageous when no information is available about the relationship between the input and the output of the system. If however, a little information is available, an inductive mechanism is advantageous as this information can be incorporated into the operative language of the rule space. In this case, where the EEG is to be classified, one does have access to a little information about the relationship between the inputs and the output of the system. A patient p_1 with parameter indices greater than or equal to another patient p_2 will belong to a grade better than that of p_2 : if p_1 were moderately abnormal then p_2 could never be severely abnormal. This is so because all parameter values are probabilities of normality. A higher value for a parameter implies a greater probability of the parameter, and hence the EEG, being normal. While training a neural network this information must be learnt by the system from the examples given to it. However, it could be passed on to an inductive learning mechanism by incorporating $<$ and $>$ signs in the operative language of the rule space. This is further discussed in the section 6.4.2. Thus, it will be advantageous to solve this problem by designing an inductive tool that learns by generalization. However, the various shortcomings of the technique of generalization discussed in section 6.2.1 are to be remedied for acceptable system performance. The system design is discussed in the following section.

6.4.2 Defining the boundary

The problem at hand is to divide a four-dimensional space, each dimension corresponding to a parameter, into four regions -- normal, mild, moderate and severe categories of EEG classification. A new EEG to be classified would then be assigned the region it falls in. The boundary between any two regions is a monotonous surface in four-dimensions. Such a surface ensures that an EEG with all parameter values greater than or equal to that of another EEG would be associated with a region indicating as much or a lesser degree of abnormality. Figure 6.2 below illustrates the problem in two-dimensions. The two axes correspond to two parameters based on the values of which the region encompassed by them is divided. Here the regions are demarcated by monotonic lines. The problem then is to define the three lines between the four regions of EEG classification.

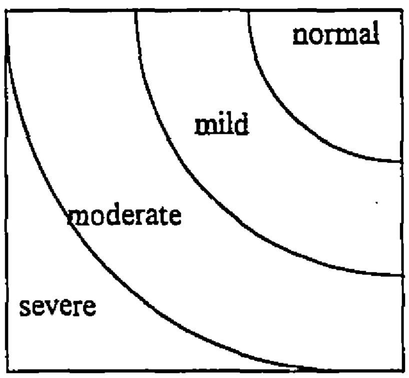

Figure 6.2

The automated EEG monitor learns concepts from examples by the technique of generalization. Three boundaries between the four categories of EEG classification are to be learnt. Each boundary may be treated as a concept and learnt independently. Every one of these boundaries divides the decision space into two regions. For instance, the boundary between the mild and moderate regions divides the decision space into two regions, the one to the right of the boundary consisting of the mild and normal regions and the other to the left of it, consisting of the moderate and severe regions. To define any one of the boundaries therefore, the exampies corresponding to the categories that are expected to be to the right of the boundary are called positive instances and those that are expected to be to the left are treated as negative examples. The aim then is to come up with a boundary definition that runs between these positive and negative instances. Such a definition of positive and negative instances feeds the system information about the ascendancy of the categories i.e. greater the parameter value, greater the level of normality.

The Most General Version (MGV) of any boundary starts at the origin of the four dimensional space, thus declaring the entire space as the region of interest and progressively cordons off regions with incoming negative instances to make the space more specific. The MGV interprets negative instances with the $<$ operator. A negative instance teaches the system that the region of the decision space with parameter indices $<$ those of the negative instance would lie to the left of the boundary. Each negative instance that alters the MGV is represented as a point with four coordinates. Therefore the definition of the MGV is a set of points. The negative instances that lie to the left of the existing boundary do not affect the MGV. Progressively, the definition of the MGV moves to the right with new incoming instances. The Most Specific Version (MSV) starts with the first positive instance of that region. With incoming positive instances the MSV relaxes its limits and makes the space more generalized. The MSV uses the greater than operator to generalize, i.e. if a positive instance comes in, then the system learns that parameter indices greater than or equal to that of the positive instance would lie to the right of the boundary being defined. The MSV boundary progressively moves to the right with incoming positive instances. Positive instances that lie to the right of the current definition of the boundary do not alter the MSV. Thus such an operative language in the rule-space, consisting of > and < signs also teaches the system about the ascendancy of the categories. At any time during training, the definition of the boundary could look like Figure 6.3 below. The +'s in the figure represent positive instances of the boundary to be defined and x's represent negative instances of the boundary. The dashed line denotes the MSV definition of the boundary while the solid line denotes the MGV definition.

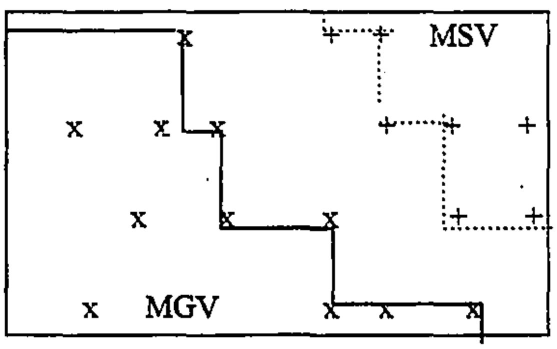

Figure 6.3

A good representative sample-set would ensure the convergence of the MSV and the MGV and would therefore define the boundary uniquely.

Each boundary limit is stored as a point with four co-ordinates. For instance P_a (p_{a1}, p_{a2}, p_{a3}, p_{a4}), P_b (p_{b1}, p_{b2}, p_{b3}, p_{b4}) and P_c(p_{c1}, p_{c2}, p_{c3}, p_{c4}) could be three points on the boundary between the normal and mild categories. A new EEG P_n (p_{n1}, p_{n2}, p_{n3}, p_{n4}) would be called normal if,

$\{p_{n1} \geq p_{a1}$ AND $p_{n2} \geq p_{a2}$ AND $p_{n3} \geq p_{a3}$ AND $p_{n4} \geq p_{a4}\}$ OR

$\{p_{n1} \geq p_{b1}$ AND $p_{n2} \geq p_{b2}$ AND $p_{n3} \geq p_{b3}$ AND $p_{n4} \geq p_{b1}\}$ OR

$\{p_{n1} \geq p_{c1}$ AND $p_{n2} \geq p_{c2}$ AND $p_{n3} \geq p_{c3}$ AND $p_{n4} \geq p_{c4}\}$

This design facilitates learning of both conjunctive concepts connected by AND and disjunctive concepts connected by OR. Learning of disjunctive concepts is not possible with most conventional generalized inductive learning techniques. The method developed here overcomes this shortcoming.

6.4.3. Dealing with bad instances

A negative instance does not alter the MSV as long as it falls to the left of boundary defined by it. However, if it falls to the right of the boundary, then a conflict between two instances, one positive and another negative is taking place. The conflict occurs because a negative instance has parameter indices greater than that of the positive instance and it is necessary to identify the bad instance. To begin with, both instances are treated as erratic data. The negative instance is dropped and does not alter the MGV. The MSV boundary is pruned, i.e. it is made less generalized by dropping the bound set by the conflicting positive instance. If a similar positive instance comes in a second time, then the system is reassured that the positive instance was indeed genuine and the boundary is reinstated. If on the other hand, a similar negative instance comes back, then no conflicting bounds exist in the MSV and hence the instance goes on to alter the boundary of the MGV.

A positive instance that falls to the right of the boundary defined by the MGV does not affect its definition. If on the other hand it falls to the left of it, it is treated as above and the MGV boundary is pruned by dropping the bound set by the conflicting negative instance. The conflicting positive instance is dropped as well and it does not alter the MSV. If however, a similar negative instance comes in again as a member of the training set, the system resets the MGV boundary as with the case for the MSV.

Thus, erratic instances are dropped and the boundary that best suits the majority of the data is taken as the definition. Such a technique for erroneous data management makes the system robust and insensitive to errors in data. The knowledge-based expert system described above is trained and tested. The results are presented and discussed in Chapter 7.

7. RESULTS AND CONCLUSION

7.1 Results of statistical analysis

After artifacts are rejected from the EEG, quantitative features are extracted from it. These features are then compared with a normal population to estimate their level of normality. As explained earlier, the t and z-statistic values estimate the levels of normality. The measures extracted were found to quantify the level of abnormality quite well as is evident from a few examples presented below.

Figure 7.1a depicts an EEG with normal amplitude and figure 7.1b illustrates a depressed EEG. The depression in activity is present in all channels and the neurologist graded it as a moderate level of depression. The distributions of the amplitude measure for the left frontal channel is presented in figure 7.2. Figures 7.2a and b corresponds to 6 hours of the EEG in figures 7.1a and b respectively. The t-values associated with these distributions are +10.12 and -28.35 respectively. As explained earlier, a t-value greater than zero implies an absolute amplitude normality and is associated with a probability of normality 1.0. The EEG in figure 7.1b is associated with a probability of normality 0.43 according to equation 5.5. A probability of normality is arrived at for each of the 8-channels recorded and these are averaged to obtain the overall probability of normality of amplitude. The probability of amplitude normality for the EEG in figure 7.1a was found to be 1.0 since all channels had a t-value greater than 0.0. The amplitude normality probability for the EEG in figure 7.1b was found to be 0.30, in keeping with the neurologist's classification.

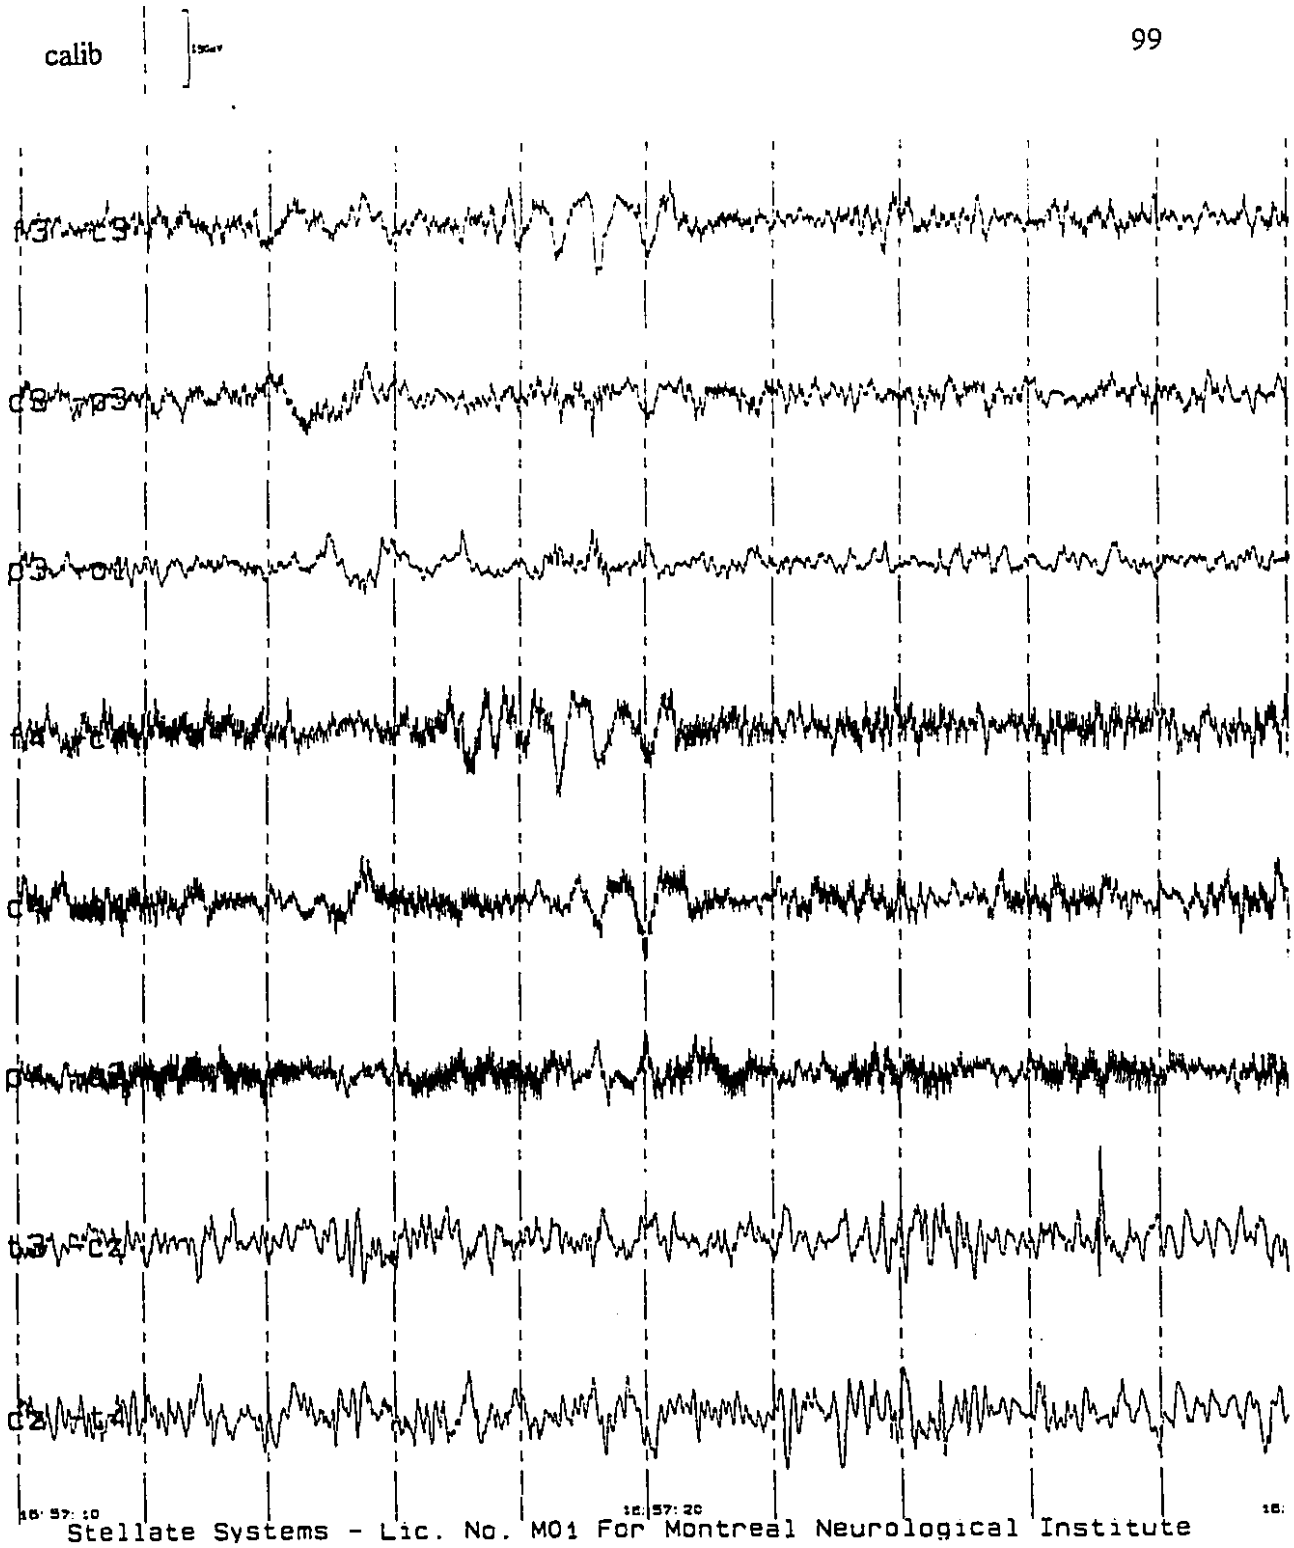

Figure 7.1(a) - 20 secs of an 8-channel EEG of a patient with normal amplitude of activity.

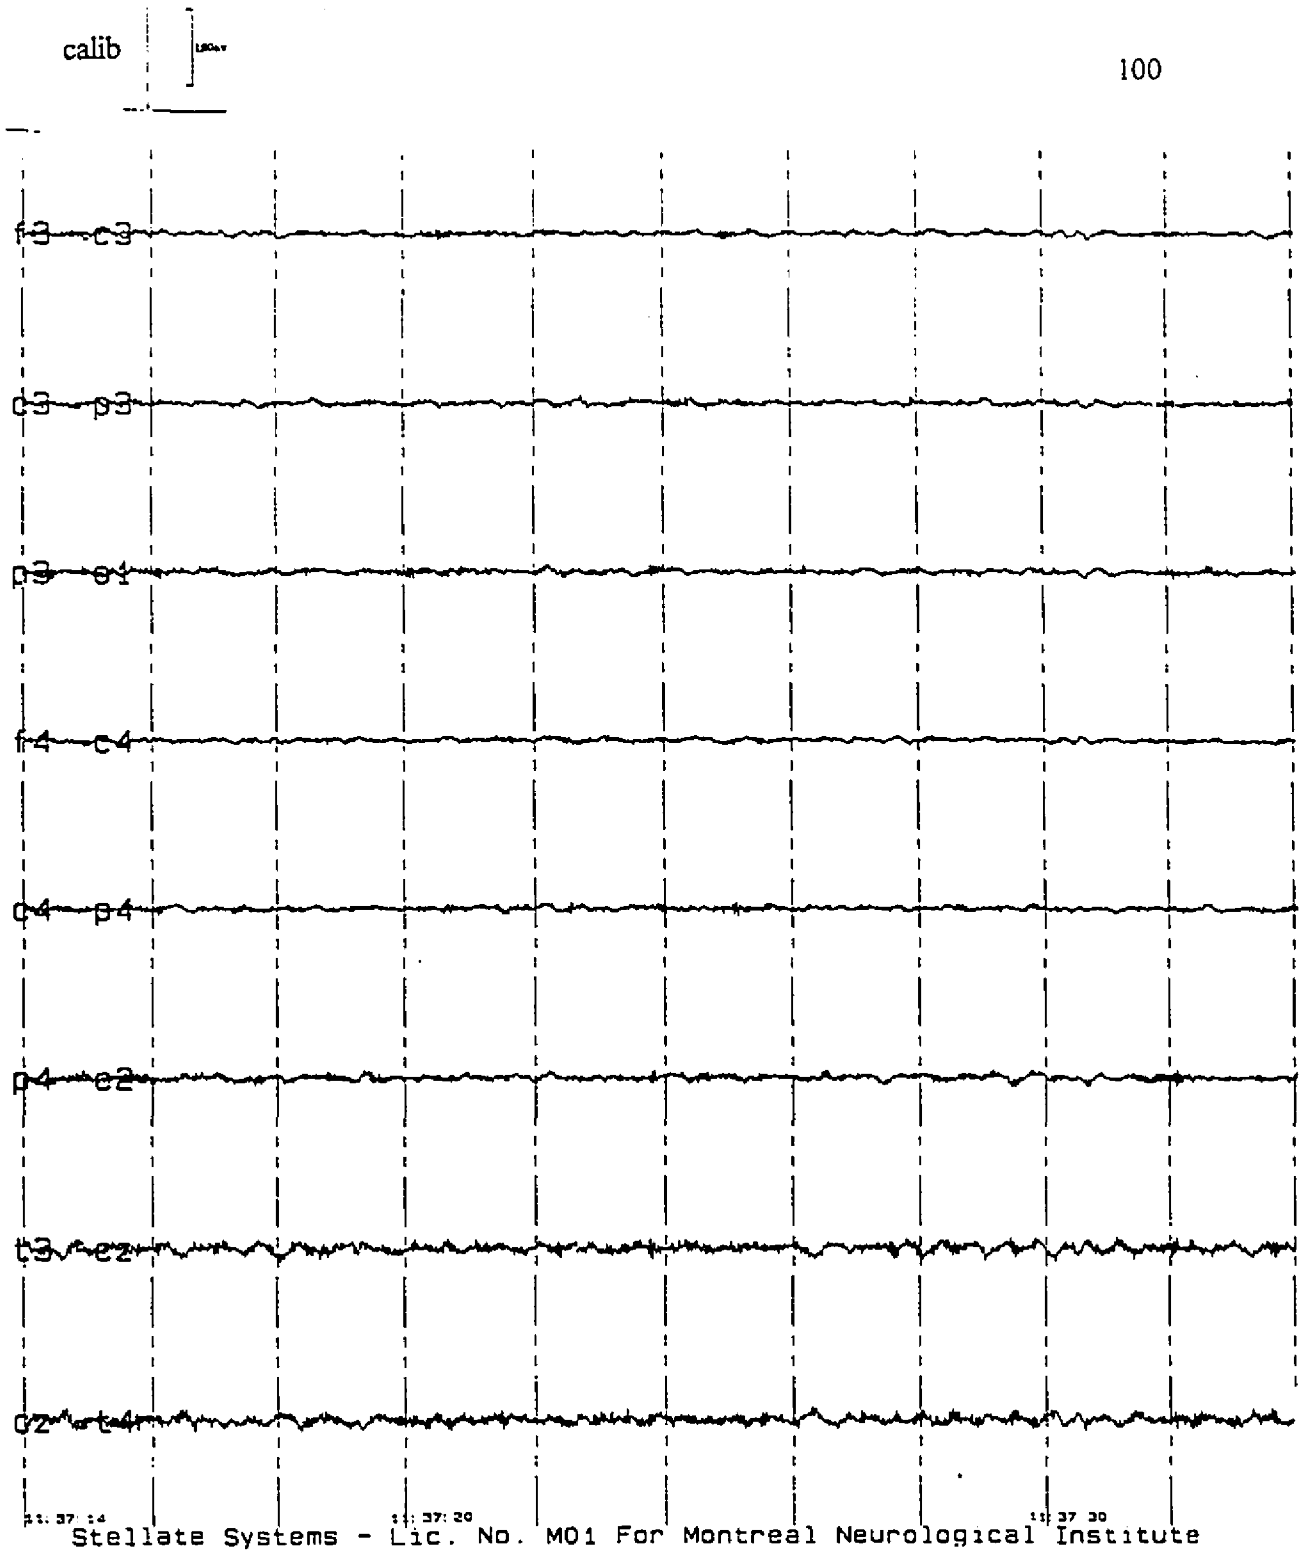

Figure 7.1(b) - 20 secs of an 8-channel EEG of a patient with depressed cerebral activity.

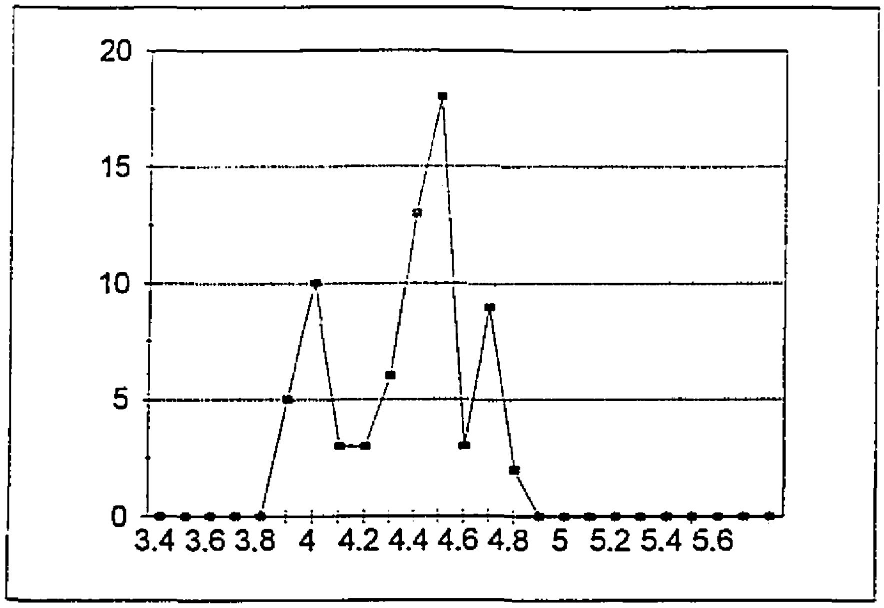

(a)

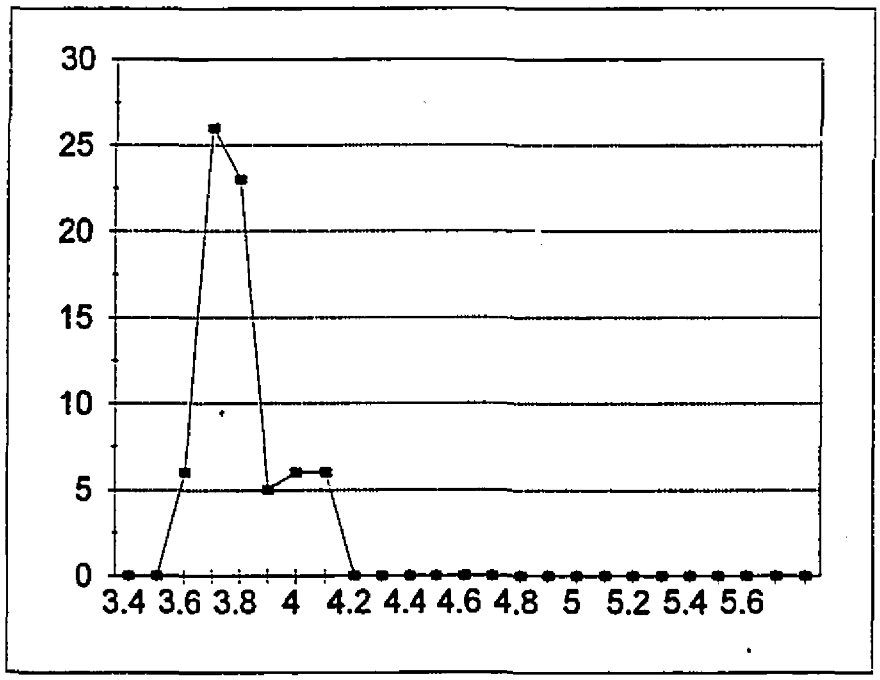

(b)

Figure 7.2 - Amplitude measure distributions of the left frontal channel for (a) a patient with normal amplitude of activity whose EEG appears in figure 7.1a and (b) a patient with depressed brain activity whose

Figure 7.3 shown on the following page depicts an EEG with an asymmetry in activity. Upon visual examination, the asymmetry in activity is found to be most pronounced in the frontal head regions (f3-c3 vs f4-c4) with a right-sided relative depression in activity. The central-parietal channels (c3-p3 vs c4-p4) also demonstrate an asymmetry in activity though to a much lesser degree. The activity in the posterior appears quite symmetrical. Figure 7.4 illustrates the symmetry measure distributions for the three channel pairs. Figure 7.4a corresponding to the frontal pair shows a clear right shift from 0.0, indicating a left-sided predominance in activity. Figure 7.4b also shows a very slight right shift while figure 7.4c is quite well-centered about 0.0. The t-values associated with each of these distributions capture well the observations made above. The frontal pair has a t-value +40.32 indicating a severe asymmetry, the central pair has a t-value +7.21 indicating a mild asymmetry and the posterior pair is associated with a t-value 0.94 indicating a symmetric activity. The probabilities of normality associated with each of these is calculated using the formula in equation 5.6 and they are 0.19, 0.86 and 0.98 respectively. The overall probability of symmetry of the EEG in figure 7.3 is found to be 0.68, obtained by averaging the three measures of symmetry. The probability of symmetry associated with this EEG does not indicate a very severe asymmetry since it is pronounced only in one of the three channel pairs.

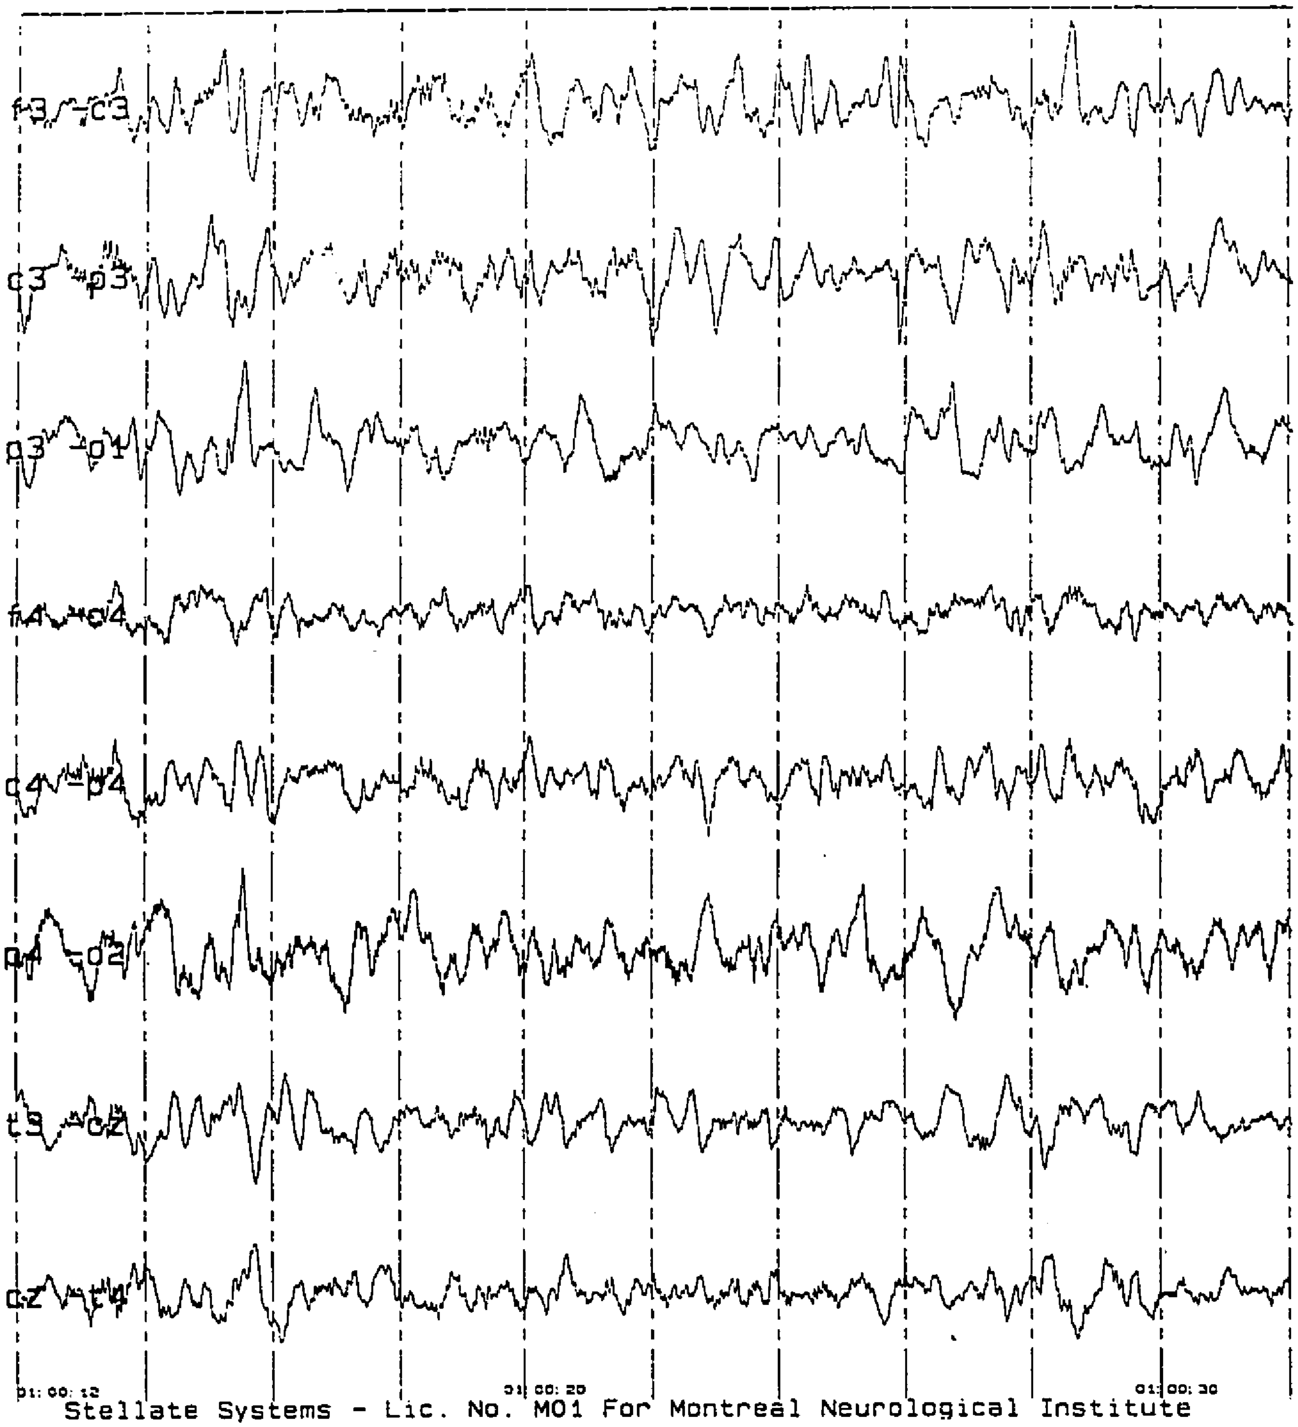

Figure 7.3 - 20 secs of an 8-channel EEG of a patient with assymmetric cerebral activity. A relative depression on the right side is evident.

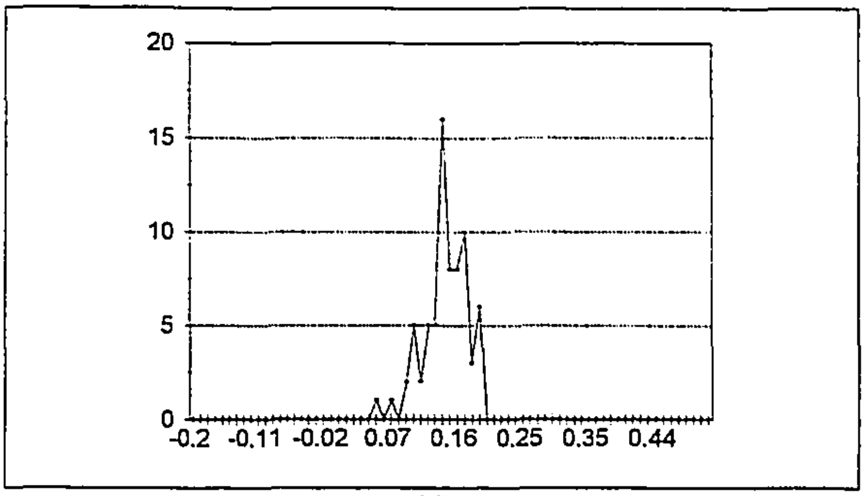

(a)

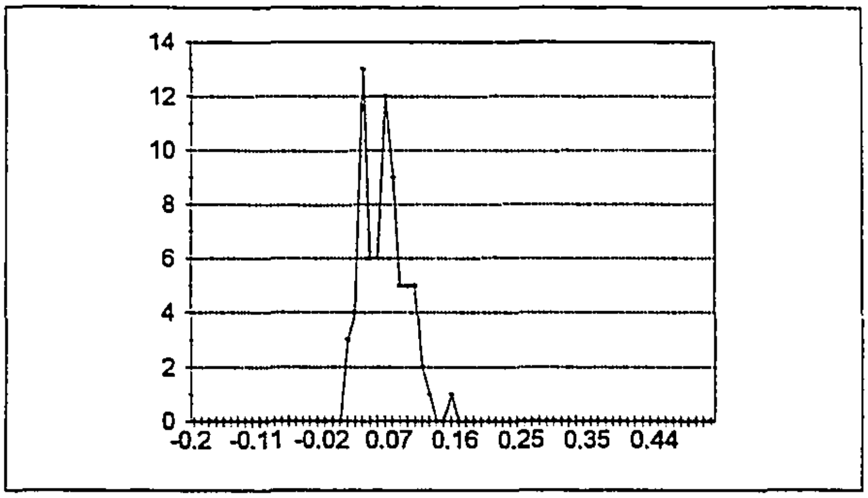

(b)

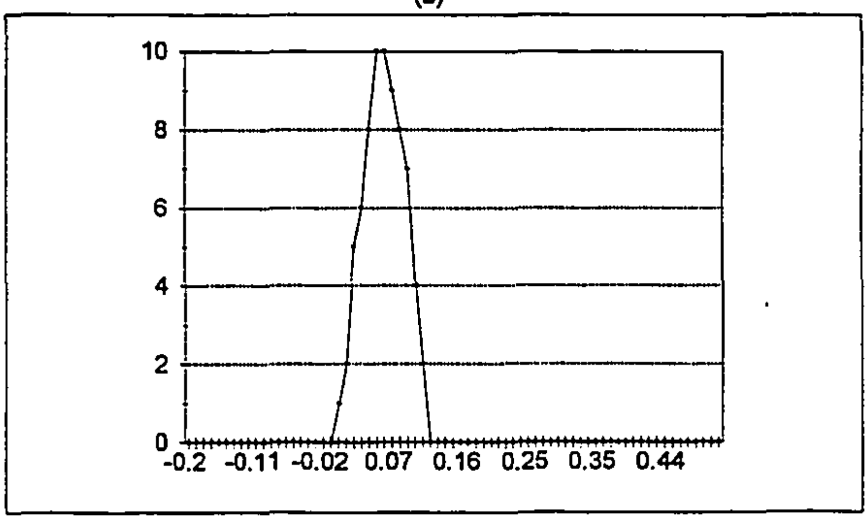

(c)

Figure 7.4 - Distributions of the measures of symmetry for the (a) frontal pair (b) central-parietal pair and (c) posterior pair for the EEG in fig. 7.3.

Variability of EEG is quantified by the zero-crossing rate and energy of change measures. Figure 4.4 illustrates the delta band activity of three EEGs - figure 4.4a with good spontaneous cycling and figures 4.4b and c without. The values of the measures and the corresponding z-statistics for the three examples are listed below.

Fig. No.	Zero-crossing rate		Log of energy of change	
	measure	z-value	measure	z-value
4.4a	20	1.85	- 3.46	2.56
4.4b	6	-10.19	- 2.74	8.14
4.4c	26	7.85	-4.52	-5.56

The EEG corresponding to figure 4.4a is associated with high values for both the energy of change and zero-crossing measures. These in turn correspond to positive z-statistic values. Figure 4.4b is associated with a high value for the energy of change alone while figure 4.4c is associated with a high value for the zero-crossing measure alone. In these cases too, high values for the measure are associated with positive z-values and low values are associated with negative z-values. A positive z-value implies that at least 50% of the control population lies below the patient being analyzed. The z-values obtained for the three examples discussed above is in keeping with the discussions of chapter 4.

The probability associated with variability of the EEG is obtained using equation 5.10. EEG corresponding to figure 4.4a has a probability of having a normal level of variability of 0.75 while figures 4.4b and c are associated with a probability of 0.0. The probability measures for variability associated with the examples discussed above clearly indicate that for good spontaneous cycling high values for both the zero-crossing rate and the energy of change are required as explained in chapter 4.

Qualitative assessment of results obtained after statistical analysis indicates a good level of correlation between the t and z-statistic values and visual interpretation. The feature extraction techniques and the statistical analysis succeed well in characterizing the EEG in terms of a few quantitative measures. The performance of the expert classifier is discussed in section 7.3. The following section discusses the data acquired and the possible and optimal ways of utilizing the data for training and testing of the expert system.

7.2 Training and Testing Data

Twenty two EEGs from as many patients recorded at the Montreal Children's Hospital are used as the training and testing data for the expert monitoring system. Each of these recordings lasts between 18 and 22 hours in duration and a majority are recordings after corrective cardiac surgery. The recordings are evaluated by a neurologist and every 6 hour section is graded as normal, mildly abnormal, moderately abnormal or severely abnormal based on visual interpretation. Every one of the twenty two EEGs yields two or three 6 hours sections depending upon the extent of artifact. In all, a set of 60 sections is available from the twenty two patients, for training and testing the expert system. There are several ways in which a data set can be divided into training and testing subsets. The merits and demerits of the various schemes are discussed below.

A popular technique used with perceptron-like learning algorithms is the resubstitution method. Here, the entire data-set is used to train the system and the same data-set is used to test it as well. One essential property of the perceptron algorithm and also the inductive learning technique used for this project is that they converge to one hyperplane that correctly classifies all the training examples. However, with small sample-sizes the hyperplanes tend to be multiple and the performance results on the design sets are not replicated by independent test sets. For this project where the sample-size is quite small the technique of resubstitution would be quite inappropriate.

The most obvious alternative to the resubstitution method is to partition the data into two mutually exclusive subsets and to use one for testing and the other for training the expert system. This scheme known as the holdout method makes poor use of the data since a learning machine trained on a larger data set will, in general, perform better than one that is trained on a smaller data. When the sample-size is small the performance of a system that is designed with only part of the data would suffer remarkably due to a non-representative data-set. The holdout method is therefore, uneconomical in its way of using the data and gives pessimistic error estimates.

The third method called the leave-one-out method goes a long way towards making efficient use of the available data and reducing the bias of the error estimate. By this technique, if the sample size is n, then, the system is trained with (n-1) samples and tested with the other. This is carried out n times until all the samples have been used for testing. Here for each run almost the entire sample-set is used for training and ultimately all samples are used in the tests, though each run consists of independent training and testing sets. The leave-one-out method has been found experimentally to be approximately unbiased whatever be the classifier used. However, the extensive computation involved as n training sessions are required, is a big drawback of this technique.

The rotation method is a compromise between the holdout and leave-one-out methods. For this method, the n samples are divided into r sets with n/r samples each. In each run one of the r sets serves as the testing set while all others are used to train the learning machine. The performance of the system is then arrived at by calculating its average performance for the r runs. The rotation method reduces both the bias inherent to the holdout method and the computational complexity associated with the leave-one-out method (Devijver, 1982).

For the purpose of this project, the rotation method of performance estimation is used. The 60 sections are divided into six subsets with 10 sections each. Six training and testing runs are carried out and the performance is the average of the six runs.

7.3 Results of the classifier

Features are extracted from each of the 60 six hour sections. For each run, the learning machine learns from the inputs and outputs of the 50 training examples and is then tested on the 10 testing examples. The cumulative result for the six runs is presented in a matrix form in table 7.1 below.

Table 7.1. Comparison of classification results from the automatic method and the human expert.

Expert's Grade / Automated grade	NORMAL	MILD	MODERATE	SEVERE
NORMAL	2	2	0	0
MILD	5	2	3	1
MODERATE	3	2	4	3
SEVERE	1	16	10	6

In the 4X4 matrix presented, rows represent system classification and the columns correspond to classification by the expert. The main diagonal represents the concordance between the two methods. A clear bias in classification toward the left of the main diagonal is evident. This implies a conservative performance by the monitoring device, i.e. the monitor in most cases, assigns to the EEG either the same grade as the neurologist or a grade of greater abnormality. This may be due to the presence of an undefined gray area between boundaries or due to an insufficiency in knowledge engineering.

According to the technique of generalization for inductive learning, the general and specific versions of the boundary should converge and provide a unique solution (Mitchell, 1977). However, this requires a large enough sample-set. If the training set is not representative enough, a gray area is present between the general and specific versions of any boundary(Figure 6.3). Such a system will perform conservatively if it uses the specific version of the boundary as the gray area would then be included in the negative instance domain. The performance of the system would be generous if the general version of the boundary were used instead. Since the technique of inductive learning is used for this monitoring unit, its conservative performance maybe due to a large gray area.

To ascertain the extent of the gray area, the system performance using the specific and general versions of the boundaries maybe compared. The performance of the monitor using the general version of the boundaries is summarized in table 7.2 below. Results appearing in tables 7.1 and 7.2 are quite similar. This indicates that the two versions of all boundaries (one between normal and mild, a second between mild and moderate and a third between moderate and severe) are quite close to each other and the intervening gray area is quite negligible. Therefore, the conservative behavior of the system is not due to insufficiency in the training set. The possibility of insufficiency in knowledge engineering is discussed in the section below.

Table 7.2 - Comparison of performance using the general version of all boundaries.

Monitor's grade \ Expert's Grade	NORMAL	MILD	MODERATE	SEVERE
NORMAL	2	4	1	0
MILD	5	2	5	2
MODERATE	4	2	3	2
SEVERE	0	14	8	5

7.4 Discussion

The aim of the project has been to develop a monitoring system for automated interpretation of long-term EEG in the pediatric ICU. Almost none of the systems that exist today are capable of complete automatic interpretation. The cerebral function monitor (CFM), extracts quantitative features from the EEG and these values are to be interpret by the user. The compressed spectral array is used extensively by several systems. Here again, the user is expected to interpret the spectral array by himself. The system developed at the Montreal Children's Hospital tries to automate the entire process of EEG interpretation so that no expertise is expected from the user.

The first step in this process is artifact rejection. Most other monitoring units with a similar aim accomplish this task in the time domain by identifying and rejecting specific artifacts usually with a single amplitude threshold (Bickford 1950, Pronk 1987, Prior 1987). Artifact rejection of this monitoring unit performed in the frequency domain, has been more efficient due to the following reasons : in the ICU it is quite difficult to anticipate the waveforms of the various artifacts and also it is computationally economical to reject artifacts in the frequency domain. By using a median filter and a hard-limiter thresholder, artifacts are not separated from the underlying signal. The artifactual page is merely replaced by the average activity of the preceding few hours. This simply ensures that long-term interpretation does not suffer unduly from to this artifactual section. However, if the entire recording is artifact-ridden, the output subsequent to artifact rejection ceases to reflect true brain activity. Therefore for accurate EEG interpretation it is important to ensure a good quality of recording. Most recordings performed in the pediatric ICU at the Montreal Children's Hospital, except an initial few, were found to have fewer artifacts than was anticipated facilitating further interpretation.

Very few monitoring units developed so far make an attempt to mimic the neurologist while performing quantitative analysis. All measures used for EEG interpretation by this monitoring unit are quantitative representations of qualitative EEG features used by the neurologist. These measures quantify the EEG quite effectively and their values demonstrate abnormalities observed by the expert on visual analysis. As observed by Chiappa (1979), these quantitative measures can be very sensitive and on several occasions picked up abnormalities missed by the expert on visual interpretation. On review, the presence of abnormalities were confirmed by the expert.

All of the four measures used by this monitoring unit fail to identify two very important EEG patterns that reflect abnormality in brain function. One such pattern is the burst suppression which heralds an unfavorable outcome reaching high statistical predictability as observed by Lombroso (1985). Burst suppression is characterized by periods of inactive background interrupted by synchronous or asynchronous bursts of activity. This pattern may be reported as generalized depression by the amplitude measure if the interrupting bursts are also low in amplitude. This is quite rare and in most cases the burst suppression pattern is reported as being normal by the monitor. The second pattern that the monitor fails to report is generalized high amplitude slow -wave activity. Chiappa (1979) reports that generalized slowing is one of the first few signs of ischemic brain damage. However, this is not identified by any of the four quantitative measures used. It is very crucial to add a few other measures that could identify the patterns described above.

Statistical analysis compares the attributes of new EEGs to be classified with those of a group of "normal controls" and grades them based on the analysis. It is important to choose the control population appropriately. The EEG of most patients in the ICU would be abnormal, if their characteristics were compared with those of an absolutely normal child due to the effects of anesthetics, medication, etc. Normal EEGs were therefore chosen from the recordings of ICU patients by visual analysis. Quantitative measures were then extracted from these EEGs and if they indicated an abnormality the EEGs were reviewed to re-confirm their classification. Thus the control population consists of patients in the ICU whose EEGs are normal.

When a t-test is used to compare distributions, it is general practice to quantify similarities of the distributions by estimating the probability associated with the test-statistic on the t-distribution. In this case the t-curve probability associated with the test statistic estimates the probability of absolute normality of the corresponding feature. It does not however, quantify the level of abnormality since the control population consists only of "absolutely normal" patients. According to the t-distribution, any t-statistic value > 3.08 is associated with a 0.0 probability of 100% normality. The t-statistic value itself on the other hand, quantifies the normalized distance between the means of the two distributions compared. Since the degree of abnormality is a function of the distance between the means, the t-statistic value measures the level of abnormality quite effectively. It is therefore more meaningful to use the t-statistic value rather than its probability value from the t-distribution. The expert system built uses the t-statistic value to quantify abnormalities. The literature shows instances of t-statistic mapping to localize normal and abnormal functions of the brain (Duffy et. al., 1981).

The EEG of children varies extensively from birth up to about 6 years of age, after which the variation is much less pronounced. An EEG pattern considered normal for a 3 month-old may be an abnormal pattern for a 3 year old child. The neurologist takes the age of the patient into account while classifying the corresponding EEG as normal or otherwise. For this monitoring unit it would be ideal to define age intervals such as < 1 year, 1 to 3 years, 3 to 6 years, 6 to 10 years and > 10 years within which the characteristics of a normal EEG recording are not expected to vary extensively. Independent "normal" control populations for the various age intervals could then be created and training and testing could be done for each of these intervals. This would probably improve the accuracy of classification of the EEGs. However, due to non-availability of representative populations for the various age intervals such a scheme could not be implemented.

The control population used for this monitoring device includes patients of all ages up to 12 years. As mentioned earlier, the neurologist takes the age of the patient into consideration while classifying the EEGs. Therefore, patients with similar quantitative measures may be associated with different degrees of abnormality due to their age and this would result in conflicting training examples when used to train the monitor. To make the system robust and insensitive to erratic training examples a scheme for rejection of conflicting training examples has been incorporated as explained earlier. This mechanism combined with the conflicting examples due to age could render the system conservative as was observed in the results obtained.

In addition to age, the neurologist's classification could be influenced by other factors such as the coma scale and patient's drug or anesthetic levels. The performance of the system would mimic that of the neurologist's better if the influence of these factors is investigated and incorporated in the monitoring unit.

The inductive learning technique by generalization used to train the expert system is found to be superior to a neural network in terms of scope for system manipulation and design. Ascendancy of grades of EEG abnormality (normal, mild, moderate and severe) could be coded right into the decision space whereas this would be impossible with a neural network. However, the expert system built thus, does not lend itself to empirical analyses of performance. It is quite difficult to get a feeling for the extent of influence of each of the measures on classification. In other words, it is quite difficult to comprehend the threshold limits physically. To further improve the performance of the monitor, it is very important to be able to gauge the importance of each of the measures used and consider alternative ways of handling them. A feel for the importance of each of the measures could probably be got by training and testing the system with three measures, dropping one each time. The nature of the results thus obtained should summarize the importance of the measure that was not used.

The system was trained with a set of fifty patients. It is possible that the training set was not large enough for the system to learn the boundaries accurately. Training and testing of the system could be performed with a set of simulated data. A large data set with various different combinations of the parameter values could then be simulated and the system performance could be evaluated better.

The overall system performance of the EEG monitor built is not very impressive. However, efficient methods of artifact rejection and feature extraction have been devised. Collection of more data for training the system could help the system learn the boundaries accurately. As explained earlier using small age intervals should go a long way in improving system performance. Inclusion of few more quantitative measures suggested above and better techniques to interpret the information provided by the measures would also enhance system performance.

8. REFERENCES

Astrup J, Simon L, Siesjo BK. Thresholds in cerebral ischemias -- the ischemic penumbra. Editorial. *Stroke* 1981; 12:723-5.

Arroyo S, Lesser RP, Gillinov AM, Redmond M, Zehr KJ, Troncoso JC, Jackson D, Baumgartner WA, Cameron DE. EEG and prognosis of neurologic recovery of dogs under profound hypothermic circulatory arrest. *Electroencephalography and Clinical Neurophysiology* 1993; 87:242-249.

Bergamasco B, Bergamini L, Dorriguzzi T, et al. EEG sleep patterns as a prognostic criterion in post-traumatic coma. *Electroencephalography and Clinical Neurophysiology* 1968;24:374-377.

Bickford RG. Neurophysiological applications of an automatic anesthetic regulator controlled by brain potentials. *Electroencephalography and Clinical Neurophysiology* 1950; 2: 93.

Bickford RG, Fleming NI, Dillinger TW. Compression of EEG data by isometric power spectral plots. *Electroencephalography and Clinical Neurophysiology* 1971; 31:632.

Bickford RG, Billinger TW, Hoffman R, Sims J. EEG monitoring techniques in the diagnosis of irreversible coma. *Electroencephalography and Clinical Neurophysiology* 1971;31:295.

Bickford RG. Automation of electroencephalography. *Electroencephalography and Clinical Neurophysiology* 1973; 35: 410.

Bricolo A, Turella G, Ore GD, Terzian H. A proposal for the EEG evaluation of acute traumatic coma in neurosurgical practice. *Electroencephalography and Clincal Neurophysiology* 1973; 34:789.

Bricolo A, Turrazzi S, Faccioli F, Odonizzi F, Sciarretta G, Erculiani P. Clinical application of compressed spectral array in long-term EEG monitoring of comatose patients. *Electroencephalography and Clinical Neurophysiology* 1978;45:211-225.

Bruner JS, Goodnow JJ, Austin GA. *A Study of Thinking*, Wiley, New York: 1956.

Cant BR, Shaw NA. Monitoring by compressed spectral array in prolonged coma. *Neurology* 1984; 35-9.

Chiappa KH, Burke SR, Young RR. Results of electroencephalographic monitoring during 367 endarterectomies. Use of a dedicated mini computer. *Stroke* 1979;10:381-8.

Chiappa KH, Hoch DB. Electrophysiological Monitoring. In: *Neurological and Neurosurgical Intensive Care*. Ropper AH, ed. Raven Press, New York 1993. 147-183.

Cho I, Smullens SN, Streletz LJ, Fariello RG. The value of intraoperative EEG monitoringduring carotid endarterectomy. *Annals of Neurology* 1986; 20:508-12.

Clymer JR, Cheng DJ, Hernandez D. Induction of Decision-Making Rules for Context Sensitive Systems. *Simulaton*. 1992.vol 59, no.3; 198-206.

Devijver PA. *Pattern Recognition : A statistical approach.* Englewood Cliffs, NJ : Prentice Hall International, 1982.

Duffy FH, Bartels PH, Burchfiel JL. Significance probability mapping : An aid in the topographic analysis of brain electrical activity. *Electroencephalography and Clinical Neurophysiology*. 1981, 51 : 455-462.

Emmerson RG, Chiappa KH. Electrophysiologic Monitoring. In: Ropper AH, Kennedy SK, ZervasNT, eds. *Neurological and Neurosurgical Intensive Care*, 2nd ed. Rockville, MD: Aspen, 1988: 13.

Engel J. Seizures and epilepsy. Philadelphia: F. A. Davis, 1989:112-34. (Contemporary Neurology Series).

Gasser T, Kneip A, Verleger R. Transformations towards the normal distribution of broad band spectral parameters of the EEG. *Electroencephalography and Clinical Neurophysiology* 1982; 53(1) : 119-24.

Greeley W, Ungerleider RM, Kern FH, et.al. Effects of cardiopulmonary bypass on cerebral blood flow in neonates, infants and children. *Circulation* 1989; 80:I-209-I-215.

Gross CE, Adams HP, Sokoll MD. Use of anticoagulants, electroencephalographic monitoring, and barbiturate cerebral protection in carotid endarterectomy. *Neurosurgery* 1981; 9:1-5.

Hayes-Roth F, McDermott J. Knowledge Acquisition from Structural Descriptions. *Fifth International Conference on Artificial Intelligence*, Cambridge, Mass. 1977; 356-362.

Homan RW, Herman J, Purdy P. Cerebral location of International 10-20 System electrode placement. *Electroencephalography and Clinical Neurophysiology* 1987; 66:376-82.

Jasper HH. The ten-twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology* 1958; 10:371-375.

Jordon K. Continuous EEG monitoring in the neurological intensive care unit. *Neurology* 1990;40(Suppl 1):180.

Jordon KG, Stringer WA. Correlative xenon-enhanced CT cerebral flow (XeCTCBF) and EEG to functionally stratify acute cerebral infarction. *Neurology* 1991; 41(Suppl 1):336.

Jordon KG. Nonconvulsive status epilepticus in the neuro-ICU detected by continuous EEG monitoring. *Neurology* 1992;42 (Suppl 1):194.

Kayser-Gatchalian MC, Neundorfer B. The prognostic value of EEG in ischaemic cerebral insults. *Electroencephalography and Clinical Neurophysiology* 1980;49:608-617.

Leng B, Buchanan BG. Extending Inductive Methods to create Practical Tools for building Expert Systems. *Proceedigs of the Fourth International Conference on Tools with Artificial Intelligence, TAI '92* 1992; xviii + 475, 442-5.

Lombroso CT. Neonatal Polygraphy in Full-Term and Premature Infants: A Review of Normal and Abnormal Findings. *Journal of Clinical Neurophysiology* 1985; vol 2:no 2: 105-55.

MacGillvray G, Kennedy J. An automatic EEG monitor using analogue techniques. *Electroencephalography and Clinical Neurophysiology* 1968;25:507.

Maulsby RL. Toward an EEG screening test : A simple system for analysis and display of clinical EEG data. *Electroencephalography and Clinical Neurophysiology* 1973;35:410.

Maynard DE, Jenkinsin JL. The cerebral function analysing monitor. Initial clinical experience, application and further development. *Anaesthesia* 1984; 39:678-690.

Michalski RS, Larson JB. Selection of Most Representative Training Examples and Incremental Generation of VL1 Hypotheses : The Underlying Methodology and Description of ESEL and AQ11. *Report 867*, University of Illinois, 1978.

Miller JD, Bobo H, Kapp JP. Inaccurate pressure readings from subarachnoid bolts. *Neurosurgery*. 1986; 19: 253-255.

Mitchell TM. Version Spaces : A Candidate Elimination Approach to Rule Learning. *Fifth International Conference on Artificial Intelligence*, Cambridge, Mass. 1977; 305-310.

Prior PF. *The EEG in acute cerebral anoxia*. 1973

Prior PF. Polygraphic monitoring in prevention of ischaemic brain damage during surgery and intensive care. *Electroencephalography and Clinical Neurophysiology*. 1987;39(Suppl) 221-226.

Pronk RAF. Data processing for monitoring brain function during anesthesia and surgery. *Electroencephalography and Clinical Neurophysiology* 1987;39 (Suppl): 215-220.

Ropper A, Rockoff MA. Treatment of intracranial hypertension . In: Ropper AH, Kennedy SF, eds. *Neurological and Neurosurgical Intensive Care*. Baltimore: University Park Press, 1983: 138.

Rumpl E, Lorenzi E, Hackl J, GerstanbrandF, Hengl W. The EEG at different stages of acute secondary traumatic midbrain and bulbar brain syndromes. *Electroencephalography and Clinical Neurophysiology*. 1979;46:487-497.

Salerno TA, Lince DP, White DN, Lynn RB, Charrette EJ. Monitoring of electroencephalogam during open-heart surgery : a prospective analysis of 118 cases. *Journal of Thoracic and Cardiovascular Surgery* 1978; 76: 97-100.

Siegel G, Agranoff B, Albers RW, Motinoff P, eds. *Basic Neurochemistry*, 4th ed. New York: Raven Press, 1989.

Simon HA, Lea G. Problem Solving and Rule Induction : A Unified View. *Knowledge and Cognition,* Gregg LW (ed). Lawrence Erlbaum Associates. Maryland : 1974; 105-127.

Slater GE, Torres F. Frequency - Amplitude Gradient : A New Parameter for Interpreting Paediatric Sleep EEGs. *Archives of Neurology* 1979;36:465-470.

Thomas WF, Cole PV, Etherington NJ, Prior PF, Stefanson SB. Electrical activity of the cerebral cortex during induced hypotension in man: a comparison of sodium nitroprusside and trimetaphan. *British Journal of Anaesthesia* 1985; 57: 134-141.

Trojaborg W, Boysen G. Relation between EEG, regional cerebral blood flow and internal carotid artery pressure during carotid endarterectomy. *Electroencephalography and Clinical Neurophysiology* 1973; 34:61-9.

Verzeano M. Servo-monitor integration if the electrical activity of the brain and its applications to the automatic control of narcosis. *Electroencephalography and Clinical Neurophysiology* 1951;3:25.

Watanabe K, Miyazaki S, Hara K, Hakamada S. Behavioural state cycles, background EEGs and prognosis of newborns with perinatal hypoxia. *Electroencephalography and Clinical Neurophysiology* 1980;49: 618-625.

Weiss NA, Hassett MJ. *Introductory Statistics* 3rd ed. Addison Wesley 1991; 475-536.

Winston PH. Learning Structural Description from Examples. *The Psychology of Computer Vison.* Winston PH (ed). McGraw Hill, New York : 1970.

Wood JH, Polyzoidis KS, Epstein CM, Gibby GL, Tindall GT. Quantitative EEG alterations after isovolemic hemodilution augmentation of cerebral perfusion in stroke patients. *Neurology* 1984;34:764-8.