Expert system approach to electroencephalogram signal processing

M T Hellyar, E C Ifeachor, D J Mapps, E M Allen* and N R Hudson*

The human electroencephalogram (EEG) is often corrupted by ocular artefacts (OAs) caused by the movement of the eyes and/or the eyelids, making the recognition of abnormal EEG signals more difficult. The removal of OAs using conventional signal processing is complicated by the similarity between abnormal EEGs and OAs, which can lead to corruption of the EEG signal. The paper describes the development of a novel approach that uses expert system techniques to differentiate OAs from genuine EEG signals, enabling OA removal to be applied only where appropriate, and ensuring that clinically relevant EEG information is left unaffected.

Keywords: electroencephalograms, ocular artefacts, intelligent signal processing

The conventional electroencephalogram (EEG) and more recent brain mapping techniques for monitoring cerebral activity produce multiple complex signals. These signals are regularly and significantly corrupted by noise or artefacts, which are caused by the electrical activity of sources internal and external to the body. The large amplitude of these artefacts obscures the true cerebral signals and the similarity of some artefacts to cerebral signals of interest can make interpretation, even by an EEG expert, more difficult. This is also true of systems for automated analysis of the EEG. The most significant type of artefactual signal present in the EEG is the ocular artefact (OA), caused by the movements of the eyes and/or the eyelids [1]. Attempts to remove OAs using conventional signal processing are complicated by both the spectral overlap of EEGs and OAs, and the similarity of their waveform shapes [2].

Current OA removal utilises adaptive filtering techniques [2][3], and it subtracts from the EEG a proportion of the signal measured close to the eyes (the electrooculogram (EOG)). This technique relies on a high correlation between the OA and the EOG, and a low correlation between genuine cerebral signals and the EOG. However, the abnormal EEG, caused by brain injury or illness, tends to manifest itself as a slowing of the usual cerebral rhythms and/or the presence of sudden EEG waveforms [4][5]. The slowing of the EEG rhythms to 0.5–4.0 Hz (the EEG delta frequency band) brings the waveforms into the same frequency band as normal eye movements; transient EEG waveforms can be similar to blink-type artefacts. In addition to this, the EOG often contains signals of cerebral origin because of the lack of electrical isolation between the electrode positions. Under these conditions, the correlation requirements above are not satisfied and the adaptive filter cannot perform correctly. *Figure 1* shows a portion of an EEG record containing an abnormal waveform in both the EEG and EOG channels, the so-called bilaterally synchronous delta (BSD) wave caused by a gross cortical atrophy or loss of brain surface material. Previous work [1] has shown that, when such abnormal waves are present, the artefact correction process leads to a reduction in amplitude or waveform distortion of the desired signal. A number of schemes [4] have been attempted for automated analysis of epileptic EEGs, typified by sudden EEG waveforms.

School of Electronic, Communication and Electrical Engineering, University of Plymouth, Plymouth PL4 8AA, UK
*Department of Clinical Neurophysiology, Derriford Hospital, Plymouth PL6 8DH, UK
M T Hellyar is also at Digital Signal Processing, Computing and Communication, Z-TEC Ltd, Devonia House, 85 North Hill, Plymouth, UK
Paper received 12 November 1993. Revised paper received 13 July 1994. Accepted 27 July 1994

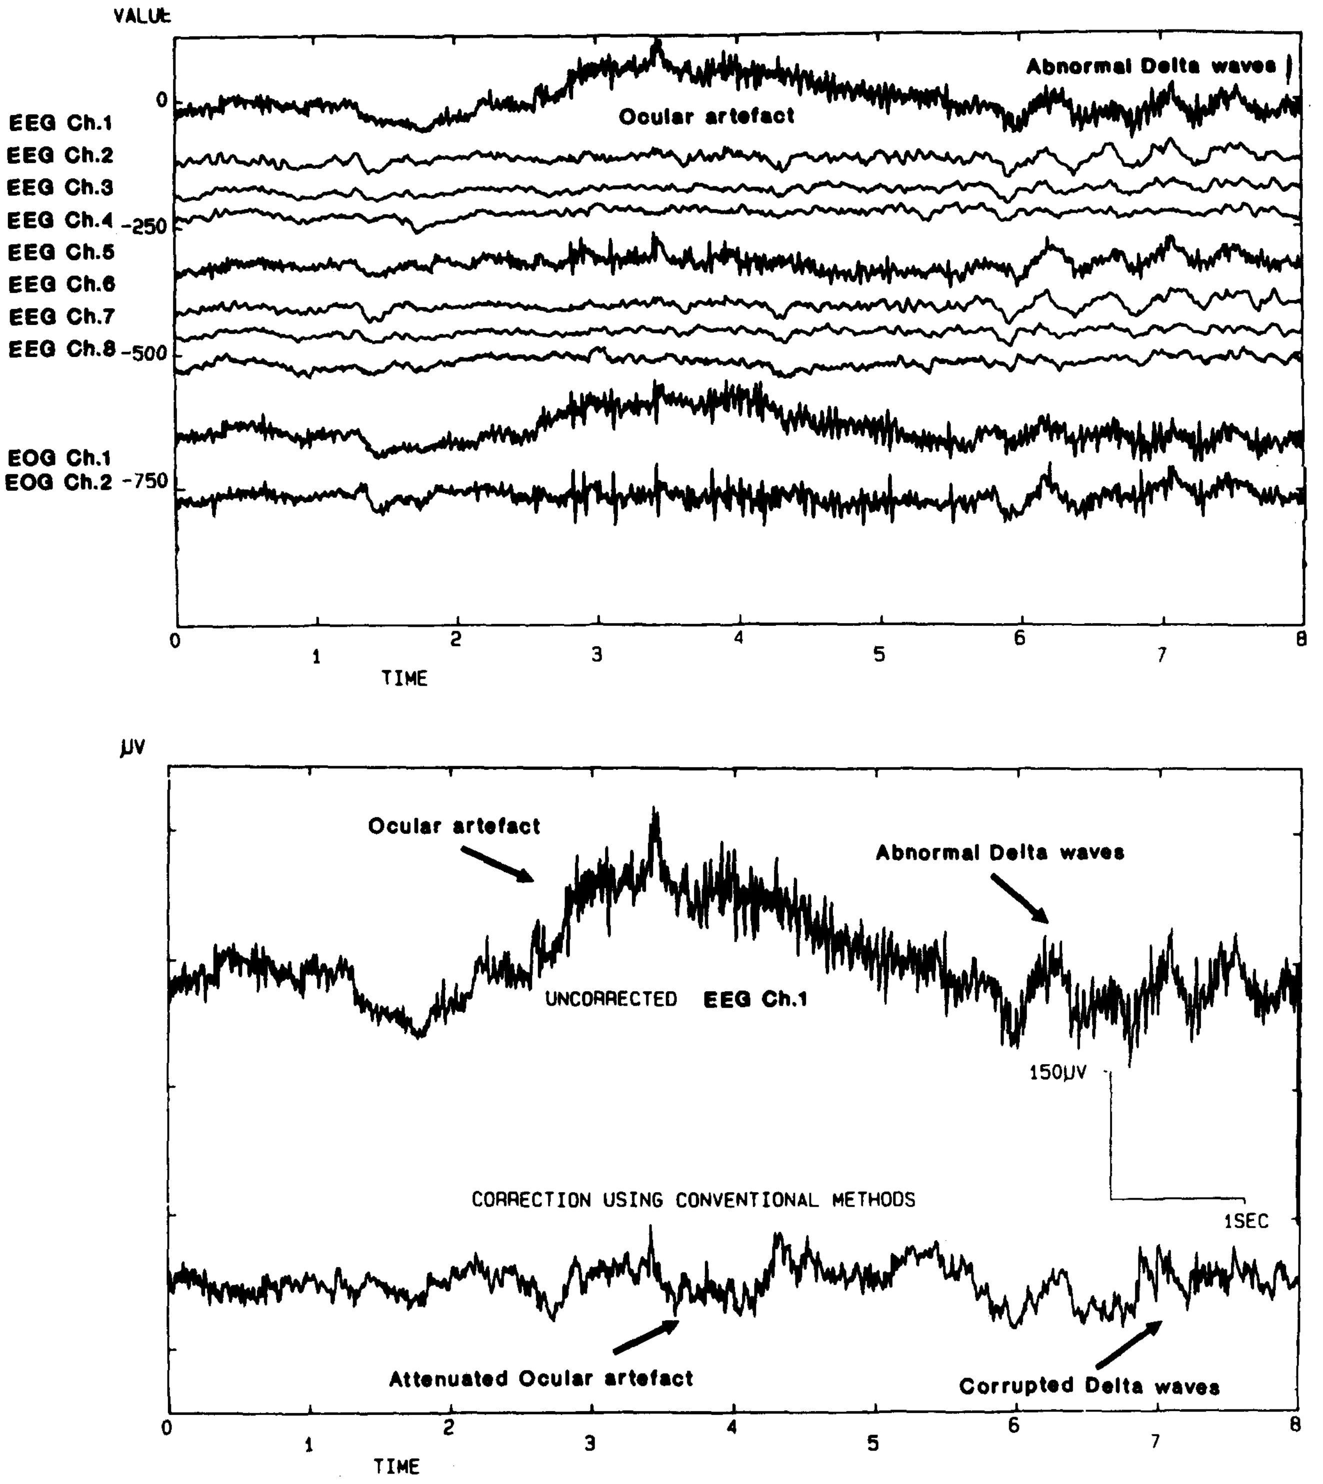

Figure 1 Removal of ocular artefacts; (a) EEG record with abnormal slow waves in both the EOG and EEG channels, (b) effects of removing OAs in presence of slow waves

In many cases, the value of such analysis is seriously impaired by false detections of epileptic activity due to blink-type artefacts.

OAs therefore present a significant hurdle in the move toward automated EEG analysis. The main limitation of the current OA removal techniques is the inability to differentiate between artefactual and abnormal waveforms [2]. An experienced electroencephalographer is often able to recognise OAs and can differentiate them from abnormal waveforms by examination of characteristic features of the signal. Thus, to overcome the deficiencies of the present methods, it is necessary to incorporate some intelligence into the correction algorithm to determine when and where to apply OA removal, and to distinguish between different artefacts, allowing a suitable adaptive filter to be applied [6][7].

This paper details the development of a novel approach which combines expert system techniques with conventional signal processing to produce an intelligent signal processing tool. This approach allows

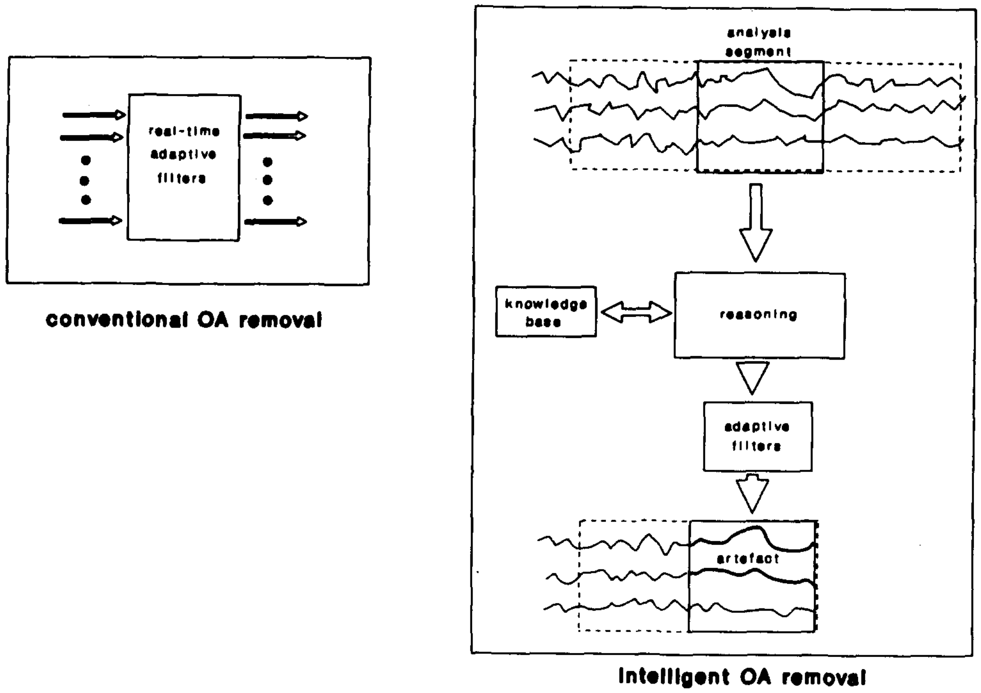

Figure 2 Conceptual comparison between conventional OA removal using adaptive filters and an intelligent signal processing approach

reasoning about signal qualities to control conventional signal processing, enabling OAs to be removed only where necessary, leaving clinically relevant EEG information unaffected.

Figure 2 illustrates conceptually the intelligent OA removal system (IOARS). Digital signal processing (DSP) modules are utilised to extract time and frequency domain features, necessary for OA identification, from 2 s blocks of 16-channel EEG/EOGs. Reasoning, using heuristic rules present in the expert system, and the extracted features allow blocks containing OAs to be identified. A positive identification of an OA allows an adaptive filter appropriate to the type of OA to be applied.

SYSTEM DEVELOPMENT

Knowledge acquisition

The rules for identifying OAs were represented as standard productions, that is as IF–THEN rules, which is

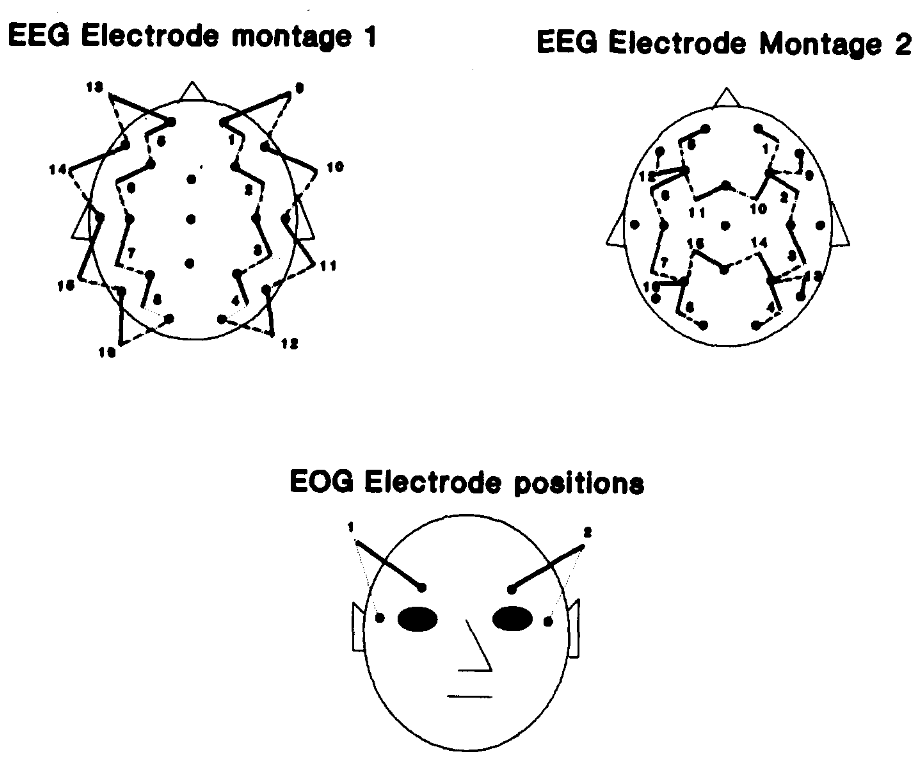

Figure 3 EEG and EOG signal derivations

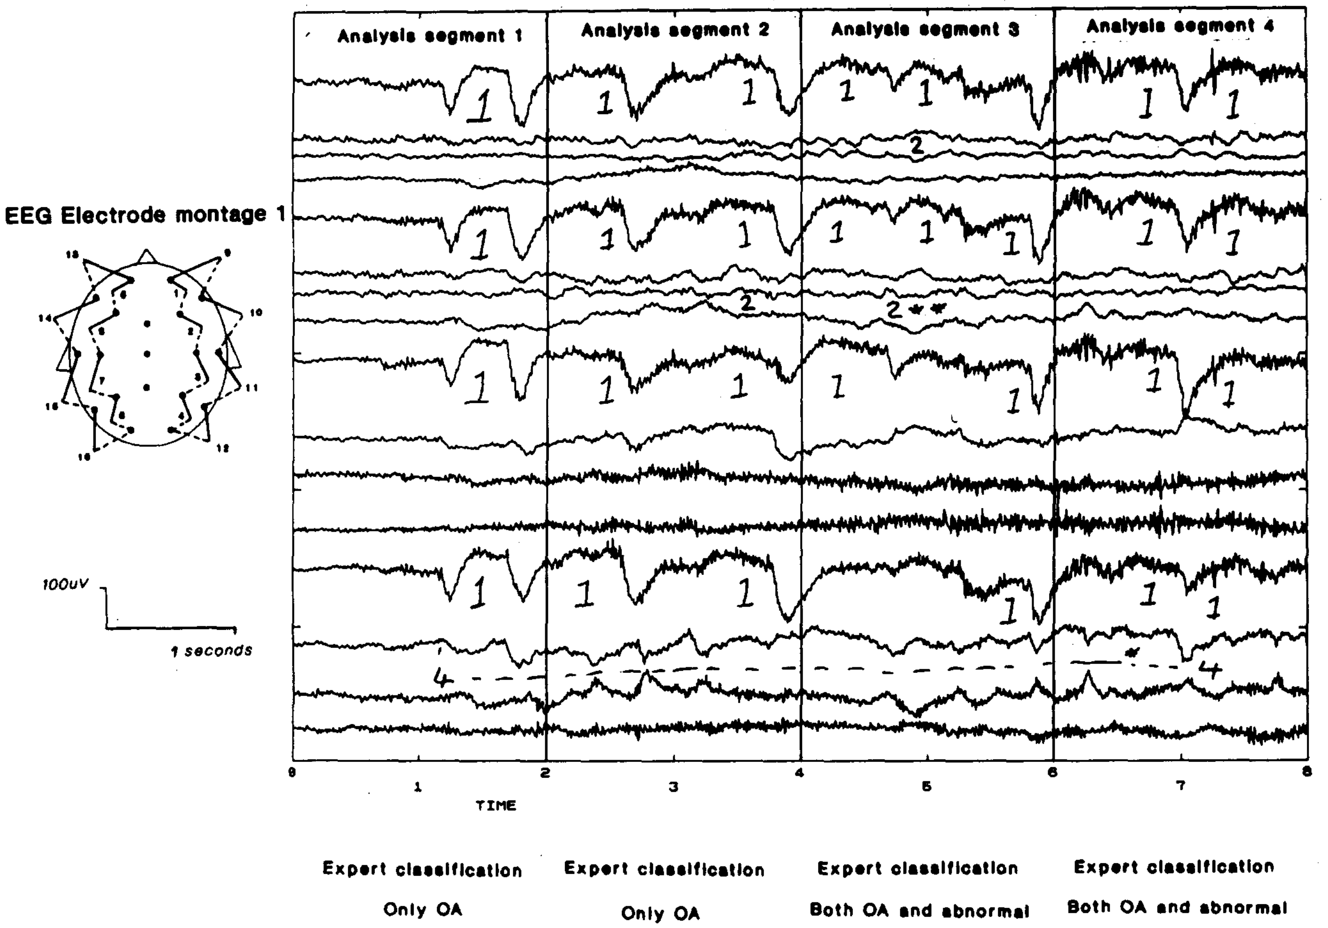

Figure 4 Example 8 s portion of 16-channel EEG/EOGs
[*: could all be electrode artefact; ** maybe some 2, but would probably have to ignore this.]

the most commonly used approach in medical expert systems [8][9] because of its simplicity and ease of understanding. The rules operate on explicit signal qualities, that is intrachannel and interchannel features, rather than on patient information such as patient age, sex etc. The EEG/EOG signals were recorded over a one year period from a variety of normal and abnormal patients using standard EEG electrode positions, as illustrated in *Figure 3*, and they were stored in an EEG database. From the acquired patient data, 64 2 s segments of data were selected for knowledge elicitation. Structured interviews with two experts (one consultant clinical neurophysiologist and one chief EEG technician) were undertaken in order to classify each of the segments into one of the following four categories:

- only artefacts present,
- only cerebral signals present,
- both artefact and cerebral signals present,
- neither artefact nor cerebral signals present.

Each classification was based on only the signal features and interchannel relationships to allow us to separate the rules based on measured data from those based on contextual information, such as patient history and age. During the analysis each expert was encouraged to describe any problems and to note any points of interest on the appropriate analysis segment. *Figure 4* shows a typical EEG trace showing four 2 s segments.

Waveforms of significance to the expert have been noted by either underlining or by writing the appropriate identification number next to it. The overall classification for each analysis segment is noted at the bottom of each segment and is dependent on the contents of the segment. In *Figure 4*, the EEG expert has identified all four segments as containing slow waves attributable to OAs, and segments 2, 3 and 4 as containing additional suspected abnormal slow waves. The abnormal slow waves can be seen clearly in channels 14 and 15 and are identified because of their measurement position on the scalp.

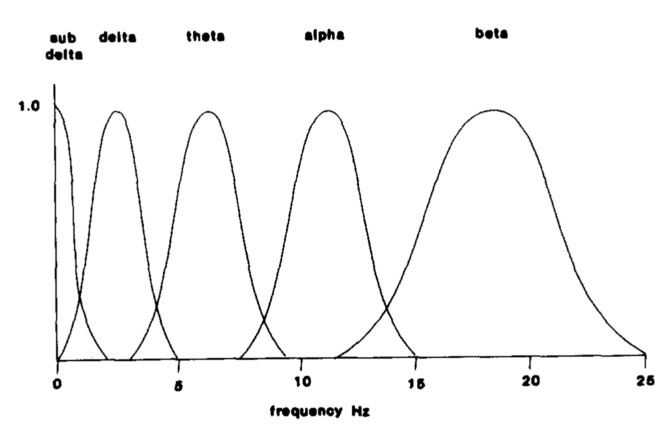

Figure 5 Frequency band fuzzy sets

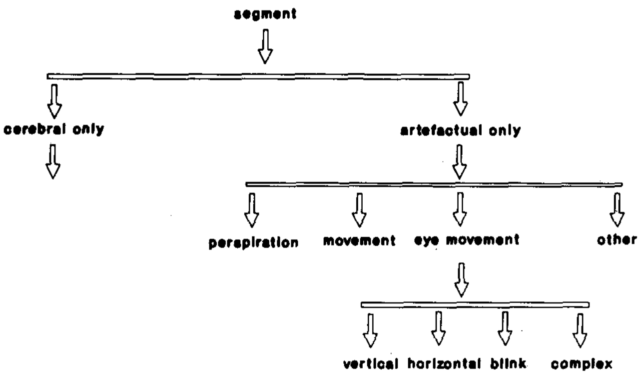

Figure 6 Simplified version of decision tree

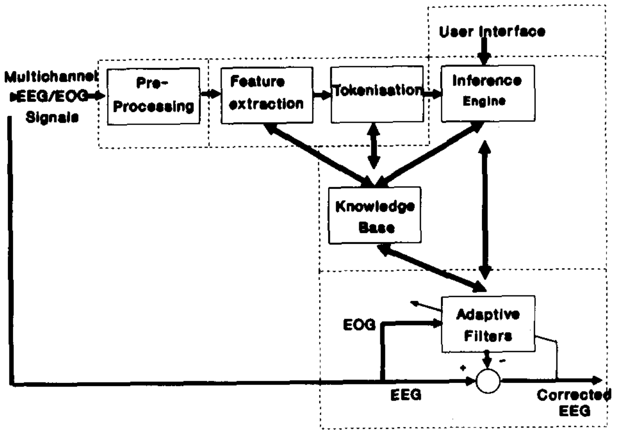

Figure 7 Intelligent ocular artefact remover

Two rules inferred from this data were the following:

Rule A:

IF slow waves are present on channels 1, 5, 9 and 13
THEN there is an increased reason to believe that OAs are present.

Rule B:

IF slow waves are present on channels 14 or 15
THEN there is an increased reason to believe that signals of nonocular origin are present.

In the first rule, channels 1, 5, 9 and 13 are the channels closest to the eyes, and therefore it is unlikely that OAs are present if there is no activity in any of these channels.

Uncertainty management

The uncertainty of each rule is represented by feature uncertainty and rule uncertainty. Feature uncertainty may be caused by a combination of measurement errors and natural biological variations in signals. The utilisation of fuzzy set theory [10] allows signal frequencies to be attributed to recognised EEG frequency bands (see *Figure 5*) with a certainty based on the frequency band set membership. A frequency transformation of the signals of *Figure 4* shows that the slow waves in channels 1, 5, 9 and 13 have a frequency of approximately 2.5 Hz. Referring to *Figure 5*, it can be seen that this frequency falls in the delta frequency set with a set membership of 1.0. Each clause in the IF statement of rule A will therefore be true with a certainty of 1.0.

Waveforms with a frequency near the boundaries of the conventional EEG frequency bands (e.g. delta band 0–4 Hz) will be interpreted as belonging to two EEG frequency bands. The membership of the sets will vary according to the frequency and a small variation in the frequency will only change the value of this membership. Similarly, signal amplitude fuzzy sets are utilised to represent signal strength.

Rule uncertainty was defined as the confidence an expert had in a particular rule. Initially, this was represented by a heuristic value between 0 and 1 elicited from the expert. Subsequent system evaluation provided a statistical value with which these were compared. The rule confidence has the effect of attenuating the combined certainty of the individual rule clauses. The final confidence *CF* in a segment classification *CL*, given the feature evidence *FE*, is calculated as

$$CF(CL{:}FE) = MB(CL{:}FE) - MD(CL{:}FE)$$

where the measure of belief *MB* and the measure of disbelief *MD* both range from 0 to 1. $MB(CL{:}FE)$ is updated by a new rule (RULE) and feature evidence as follows:

$$MB(CL{:}FE,\mathrm{Conf(RULE)}) = MB(CL{:}FE) + \{\mathrm{Conf(RULE)} \times [1 - MB(CL{:}FE)]\}$$

where the rule confidence Conf(RULE) is calculated as

$$\mathrm{Conf(RULE)} = \min\{\mathrm{Conf(FEATURE1)}, \mathrm{Conf(FEATURE2)}, \ldots \mathrm{Conf(FEATURE}n)\}$$

and the feature confidence Conf(FEATURE) is obtained from *Figure 5*.

It is assumed that each successful rule provides further evidence to support the belief in a classification. This method has the advantage that, unlike in other methods of combination, successive evidence increases the certainty asymptotically towards a certainty of 1, or complete belief. Also, the order of acquired evidence is unimportant.

Rules are executed using a combination of forward and backward chaining techniques to navigate a segment classification decision tree (see *Figure 6*). Depth first, backward chaining is employed as the principal rule search technique [11]. This allows the low level signal features to be operated on. However, the classification of previous segments allows a forward chaining approach to be taken to current segment classification. Contextual information enables the decision tree to be restricted, or pruned. The following rule is an example:

Rule C:

IF abnormal waveforms have been identified in previous segments
THEN there is an increased reason to believe that the present segment contains abnormal waveforms.

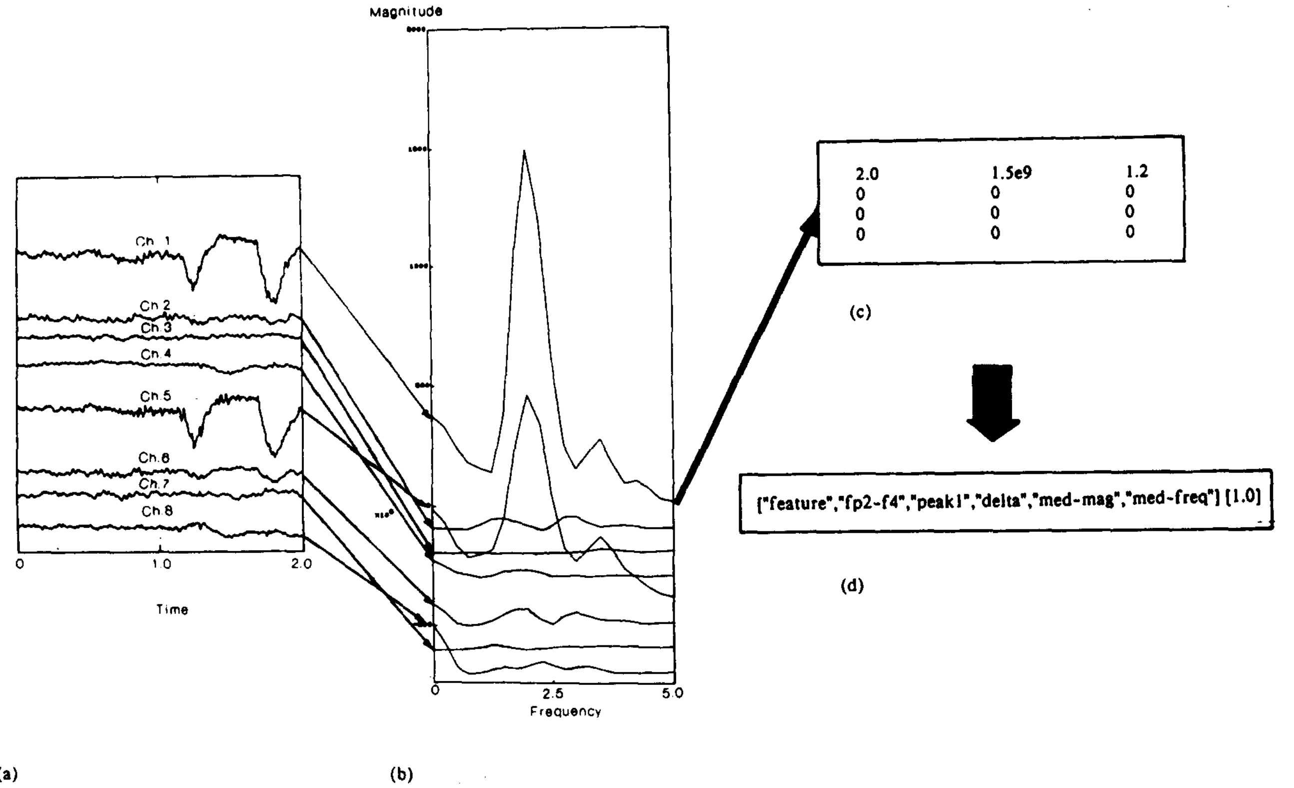

Figure 8 Spectral feature generation from power spectral density; (a) 8 channel analysis segment time series, (b) 8 channel power spectral density of *Figure 8a*, spectral peak feature matrix for channel 1, (d) asserted feature fact for channel 1

The addition of contextual information allows a mixed search strategy to be employed. Contextual data allows inference to a conclusion by forward chaining and then additional data can be sought to confirm the conclusion by backward chaining. This more closely matches the experts' method of analysis as well as improving the performance of the system.

IMPLEMENTATION

The intelligent OA removal system can be divided into four main functional blocks as illustrated in *Figure 7*. Each functional block has been implemented in software using a mixture of Prolog, C and Assembly languages on an IBM compatible personal computer.

Signal preprocessing

The EEG and EOG data are fetched in 8 s 16-channel blocks, and the mean of each signal is removed to minimise DC offset. EEG and EOG signals are derived from the data to provide a single left and right EOG signal and 16 EEG signals that reflect the activity in the left, right, front and back regions of the scalp. Analysis of the spectral characteristics of EEGs and OAs has shown that most of the energy lies in the regions 0–30 Hz and 0–5 Hz, respectively. For this reason, EOG signals are bandlimited to 0.5–5.0 Hz and EEG signals are bandlimited to 0.5–30.0 Hz. Finite impulse response (FIR) digital filters were used in both cases. Bandlimiting EOG signals has the advantages of attenuating very low frequency potentials caused by skin/electrode resistance variations and also attenuating higher frequency cerebral potentials which appear as artefacts in the EOG. Finally, each 8 s block is divided into four 2 s analysis segments in preparation for feature extraction.

Feature extraction

Signal features are extracted from three domains:

- the frequency domain,
- the time domain,
- the contextual domain.

Frequency domain features include the frequency, magnitude and phase of a wave. Time domain features include wave duration and shape. Contextual information includes patient information (age, sex), and interchannel differences, as well as immediate past, current and accumulated past features. Frequency and time domain features are extracted using Turbo C procedures. The power spectral density of each 2 s analysis segment is calculated for each channel using fast Fourier transforms (FFTs) for convenience. From this, spectral peaks are identified by comparison with a magnitude threshold related to the average power in the standard EEG frequency bands. The first four significant and distinct peaks in each channel are then used to compile a feature matrix that contains the magnitude, centre frequency and threshold width for each peak. *Figure 8* illustrates the process of spectral feature generation from a simplified analysis segment.

The feature matrix is read by Prolog routines which are used to convert the numerical features into symbolic tokens. The symbolic tokens are used to create individual spectral feature facts in a dynamic knowledge base that changes for each analysis segment. For example, the waveforms in channels 1 and 5 of *Figure 8a* are transformed to the clear spectral peaks at a frequency of 2 Hz in *Figure 8b*. The feature matrix of *Figure 8c* shows that, for channel 1, only one spectral peak greater than the threshold exists. This spectral peak has a centre frequency of 2.0 Hz, a magnitude of 1.5e9, and a threshold width of 1.2 Hz. *Figure 8d* shows the respective tokenised feature fact which will be stored in the dynamic knowledge base. The feature fact states that a spectral peak exists in channel *Fp2–F4* (channel 1), that it is the first spectral peak, and that it is attributable to the delta frequency band. The certainty of the fact is represented by the final figure in parenthesis, i.e. 1.0, and it is determined by the fuzzy set technique detailed above. Incorporation of fuzzy set theory allows a spectral peak which has a frequency close to one of the conventional EEG frequency band boundaries to be represented as two feature facts in the dynamic knowledge base. For example, if the spectral peak of *Figure 8c* had a centre frequency of 4.0 Hz, the respective feature facts would be as follows:

["feature", "*fp2-f4*", "peak1", "delta", "med-mag", "med-freq"] [0.5]

["feature", "*fp2-f4*", "peak1", "theta", "med-mag", "med-freq" [0.5]

where "theta" is the next adjacent EEG frequency band.

The last two elements in the first set of parentheses for each spectral feature fact relate to the magnitude and frequency of the respective spectral peak. These values are used for comparisons of the power distribution of the scalp.

Inference engine

The inference engine comprises the knowledge base and the reasoning mechanism. Both of these have been implemented using Prolog. The knowledge base consists of 40 rules and 54 conditions to those rules. A sample of these rules are as follows:

Rule 2:

Segment contains no slow waves (1.0)
IF no significant spectral activity exists in the delta band.

Rule 4:

Segment contains artefact only (0.93)
IF slow activity is maximum in the frontal channels
AND prefrontal channels are symmetrical

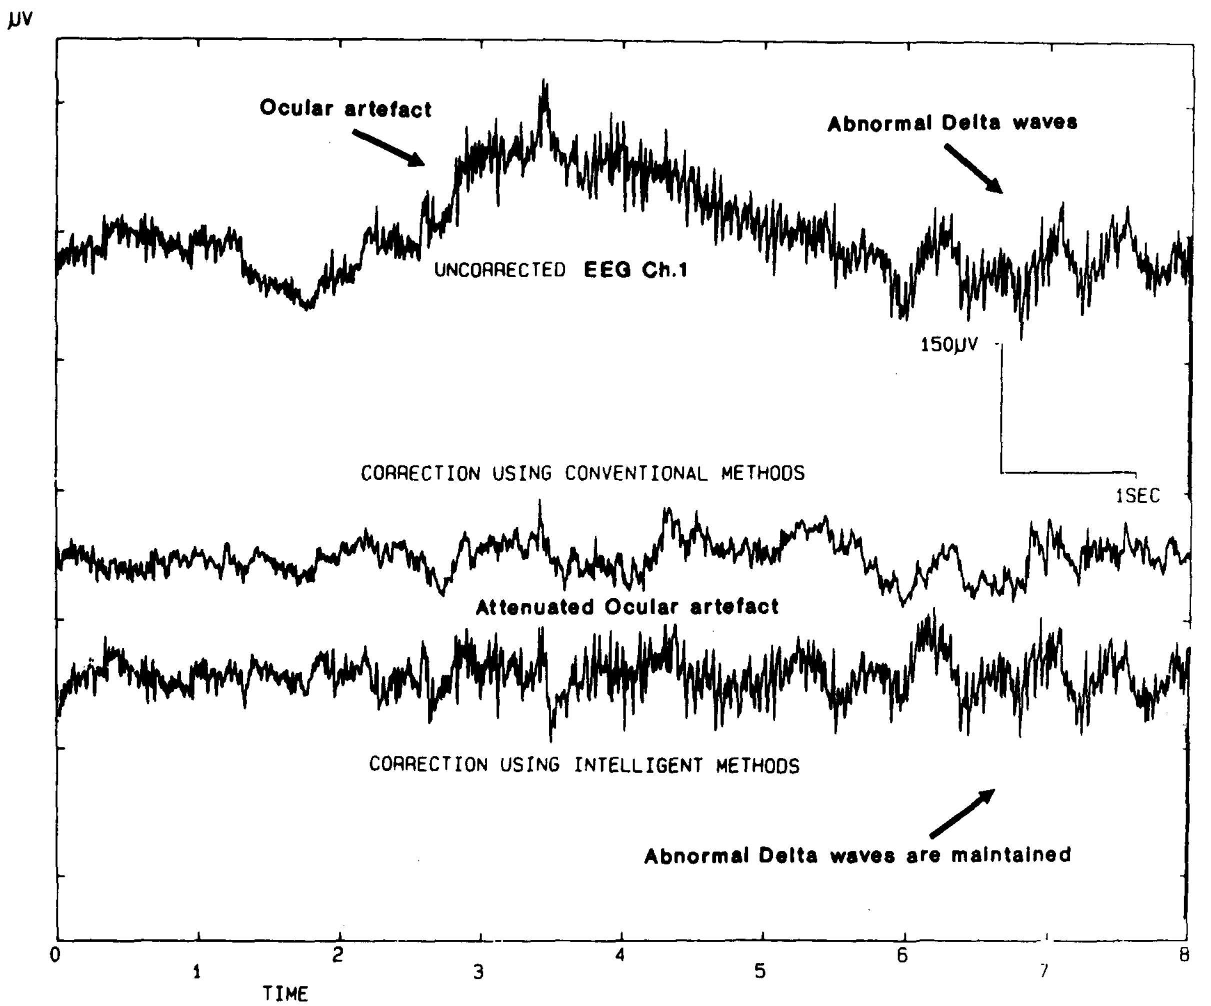

Figure 9 Comparison between conventional OA removal and selective and directed OA removal

AND EOG channels are symmetrical
AND slow waves only appear in frontal channels.

Rule 12:

Segment contains artefact only (0.4)
IF slow activity is maximum in the frontal channels
AND slow waves only appear in anterior channels
AND no slow waves exist in the EEG that are not present in the EOG.

Rule 20:

Segment contains both artefact and abnormal waves (0.8)
IF slow activity is not maximum in the frontal channels
AND prefrontal channels are symmetrical
AND slow waves exist in the EEG that are not present in the EOG.

Rule 23:

Artefact only contains eye movements (0.9)
IF EOG channels are symmetrical
AND slow waves are attributable to more than one electrode
AND slow activity is maximum in the EOG channels.

Rule 28:

Eye movements contains blink artefact (0.9)
IF slow waves are larger in prefrontal channels than in temporal
AND the slow waves are of short duration
AND the slow waves are in phase in temporal channels
AND the slow waves are phase reversed around the eyes.

The key features utilised by the inference engine in classifying each segment and identifying OAs are the position of the maximum delta band spectral magnitude, the similarity of the spectral content in the left and right EOG signals, the cerebral distribution of spectral peaks in the delta band, and the correlation between combinations of EOG and EEG signals.

Ocular artefact potential is strongest in the anterior regions of the scalp, around the eyes, and it is also symmetrical. Rules 4 and 12 identify a segment as being only artefactual if delta activity is maximum in the frontal channels. Either rule will be found to be true when all respective conditions are satisfied with some degree of certainty. Rule 12 is only tested when Rule 4 fails. The confidence in successive rules decreases as these conditions are relaxed. For example, Rule 4 has a greater number of signal symmetry conditions than Rule 12, and is therefore given a higher confidence than Rule 12 should it be satisfied. The waveforms in *Figure 8a* satisfy all the conditions of Rule 4.

Segments identified as containing slow waves only attributable to artefacts are further identified as containing only OAs by examination of the EOG channels. For OAs to be present, the delta waves must be maximum in the EOG channels and symmetrical. Rule 23 identifies a segment as containing only OAs by examining the spectral symmetry of the EOG channels. The EOG channels for the segment in *Figure 8* are identified as possessing symmetrical power spectra, and they are therefore further identified as containing only OAs.

For OA identification, the phase relationship between signals close to the eyes is also of importance. Vertical eye movement (VEM) and blink produce signals that are in phase on both sides of the scalp, while horizontal eye movement (HEM) will produce convergent signals on one side of the scalp and divergent signals on the other. The time courses of OAs are also different: blinks produce large, short duration, negative potentials above the eye, and small positive potentials below the eye. VEMs and HEMs produce longer duration potentials. Rule 28 examines the time domain features of the waveforms contained in the analysis segment to further identify the contents of the analysis segment. The waveforms of *Figure 8* are finally identified as being attributable to blink-type OAs with a sufficient measure of belief, and the OAs are removed by applying the signals to an adaptive filter with coefficients preset to the values that are most suitable for blink-type OAs.

The knowledge base is a separate ASCII file divided into two sections. The static, or long term, knowledge base contains the rules and signal information, i.e. channel numbers and their location. The dynamic, or short term, knowledge base contains the extracted segment features from the current analysis segment and the contextual features obtained from other segments. Separation of the knowledge in this manner allows easy modification.

The static knowledge can be classified as follows:

- textual rules,
- textual conditions,
- numerical conditions,
- feature demons,
- contextual facts.

An example of a simplified knowledge base is given below:

text-rule(3,1,"segment","no significant delta waves",[4])
text-rule(12,0.4,"segment","artefact only", [5,13,16])
text-cond(3,"any significant spectral activity exists in the delta band")
text-cond(4,"not any significant spectral activity exists in the delta band")
text-cond(5,"slow activity is maximum in the frontal channels")
text-cond(6,"not slow activity is maximum in the frontal channels")
text-cond(13,"slow waves only appear in anterior channels")
text-cond(14,"not slow waves only appear in anterior channels")
text-cond(15,"slow waves exist in the EEG that are not present in the EOG")
text-cond(16,"not slow waves exist in the EEG that are not present in the EOG")
num-cond(3,["feature","","","delta","","",])
num-cond(5,["max-front"])

```
num-cond(14,["posterior-features"])
num-cond(15,["non-eog"])
    cl(["non-eog"],[],[
                ["intpred","channel","X","eeg"],
                ["feature","X","","delta","", A"],
                ["intpred","notpeak","peak","A"],
                ["anyeog","A"],
                ["intpred","notpeak","peak","A"]])
    cl(["any-eog","A"], [], [
                ["intpred","channel","Y","EOG"],
                ["feature","Y","","delta","","A"].
                ["intpred","assertpeak","peak","A"],
                ["intpred","cutback","",""]])
    cl(["any-eog","A"],[],[ '
                ["intpred","true","",""]])
    cl(["anterior","fp2-f8"],[1],[])
    channel(0,"fp2-f8","EOG")
    freqbounds("delta",6,18)
```

In order to provide the explanation facilities crucial to this system, rules and conditions are stored in a textual format. This is augmented by the necessary numerical rules which operate on the features obtained from the data. For example, examination of the example knowledge base above shows the equivalence of textual condition 15 and numerical condition 15 which test for the existence of delta spectral peaks present in the EEG but not in the EOG. Condition 15 is the inverse of condition 16. Condition 15 is implemented by the use of the feature demon 'non-eog'. The feature demon is a small 'meta'level representation [12] of a Prolog clause that examines the actual feature facts extracted from the analysis segment. The metalevel representation allows the knowledge base to contain Prolog type code as conditions. Condition 15 will be found to be false if delta frequency waveforms are found in the EEG that are uncorrelated with those in the EOG, causing condition 16 to be found true with the appropriate certainty.

EVALUATION

In order to evaluate the performance of the system, 140 2 s 16-channel EEG/EOG segments were randomly selected from the patient database. These segments had not previously been used in the analysis. 40 segments were selected from normal volunteer data containing various types of OA, and 100 segments were selected from the patient data containing abnormal EEGs. The abnormal EEG segments contained a wide variety of both OAs and abnormal slow waveforms.

The selected segments were presented to an expert for each segment to be classified into one of the four categories detailed above in a similar manner to that used for knowledge elicitation. Abnormal classification by the expert was based on the slow wave or delta band contents of each segment, and not on any higher frequency abnormalities. Segments were also presented to the intelligent OA removal system which was to produce a similar classification. For the purposes of this initial evaluation, the knowledge base could only call upon features extracted from the time and frequency domain.

The operation of the present system is limited to segment identification only. OA removal is achieved by implementing a realtime adaptive filtering algorithm for those segments identified as containing OAs only. The OA-dependent starting coefficients for the filter are derived from offline analysis using multiple regression of segments containing the respective OA. These coefficients preload the adaptive filter and are then adjusted in realtime, for the duration of the segment only, using the recursive least squares algorithm [3]. The final coefficients are then stored and act as the starting coefficients for the next segment containing that particular type of OA. *Figure 9* shows an example of the output obtained from the IOARS. This data is the same as that in *Figure 1*. Segments 1, 2 and 3 are classified as OA only. Segment 4 is classified as containing abnormal slow waves because of the spread and correlation of the slow waves. The adaptive filter was only applied to segments 1, 2 and 3. Segment 4 was left unaltered. This has the effect of avoiding the corruption of the abnormal slow waves caused by the conventional OA removal.

DISCUSSION

Overall, there was an agreement of 94% between the EEG expert and the intelligent OA removal system in identifying segments containing OAs only, and an agreement of 84% for segments containing mixtures of OAs and abnormal slow waveforms. It is certain that the use of more than one expert would create differences of opinion for a small number of segments, but overall the results are very encouraging considering the restricted number of features that the system is able to operate on. The greatest source of error can be seen to be caused by the system identifying segments that contain only abnormal slow waves as containing both OA and abnormal slow waves. This is caused by the limitation of a restricted feature set. The segments that were wrongly classified all contained slow waves that were maximum in the frontal and EOG channels and were distinguishable in phase over the entire scalp. The rule set correctly classified these segments as possibly containing both OAs and abnormal waves with a low degree of certainty. Given the fact that the waveforms were maximum in the frontal channels, symmetrical in both hemispheres, and present in posterior electrodes, this was the safest and only conclusion that could be drawn under the circumstances. It was considered better to err on the side of caution rather than to remove any EEG information by unnecessary OA removal. In order to overcome this discrepancy, it will be necessary to incorporate a number of interchannel correlation features. These, together with more important contextual features, will considerably improve the results. The authors are currently investigating the use of neural network techniques at the feature extraction stage to improve the identification and location of individual artefactual waveforms. The expert used in the preliminary evaluation of the system contributed to the development of the rule base, although the bulk of the knowledge was provided by the other expert. At the next stage of our work, the system will be evaluated with the help of experts at other centres after appropriate enhancements.

Realtime operation is presently hindered by the variable time taken for inferences to be made and by the necessity to perform multiple FFT calculations using the single PC processor during the feature extraction stage. This can be overcome by using multiple dedicated signal processing hardware and faster PC hardware. This, together with a refined inference mechanism and careful knowledge base organisation, will permit real-time intelligent OA removal.

The advantages of this directed and selective approach to intelligent noise removal is clear. Firstly, segments containing genuine cerebral signals of possible diagnostic importance are left unaffected, and, secondly, individual artefacts within a segment are removed using the most appropriate filter on the basis of the removal of previous similar artefacts.

The methods detailed in this paper have been shown to overcome some of the main limitations encountered with the conventional method for realtime OA removal from EEGs. The current system is being incorporated into a self contained decision support tool for use in the clinical environment. Developmental aspects of the intelligent OA removal system not covered in this paper include the user interface and library facilities which will form a crucial part of the system. Current work is concerned with extending the techniques used in the intelligent OA removal system to general automated EEG analysis to improve the reliability of any decisions made. Automated EEG analysis systems aim to alleviate the burden of routine EEG analysis for the clinician.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support from the EEG Trust for the work described in this paper. The EEG data were obtained in collaboration with the Department of Clinical Neurophysiology, Derriford Hospital, Plymouth, UK, and the authors would like to extend their gratitude to the technical staff in this department and to the many volunteers and patients involved in obtaining the data. Dr M T Hellyar would also like to thank the University of Plymouth for financial support.

REFERENCES

1 Jervis, B W, Ifeachor, E C and Allen, E M 'The removal of ocular artefacts from the electroencephalogram: a review' *Med. & Biol. Eng. & Comput.* Vol 26 (1988) pp 2–12

2 Ifeachor, E C, Hellyar, M T, Mapps, D J and Allen, E M 'Knowledge based enhancement of human EEG signals' *Proc. IEE* Vol 137 No 5 (1990) pp 302–310

3 Ifeachor, E C, Jervis, B W, Allen, E M and Hudson, N R 'A new microcomputer-based online ocular artefact removal (OAR) system' *Proc. IEE* Vol 133 (1986) pp 291–300

4 Ktonas, P Y 'Automated analysis of abnormal electroencephalograms' *CRC Critical Reviews in Biomedical Engineering* Vol 9 No 1 (1983) pp 39–97

5 Muras, J S and Binnie, C D 'The recognition of frontal slow activity in the presence of eye movements' *Proc. Electrophysiological Technologists Assoc.* Vol 17 (1970) pp 131–143

6 Sharman, K C, Chambers, C and Durrani, T S 'Rule driven adaptive signal processing' *IEE Colloquium on the Application of AI Techniques to Sensor Systems* London, UK (Apr 1987)

7 Ferrara, E R and Widrow, B 'Time sequenced adaptive filter' *IEEE Trans. Acoust. Speech & Signal Processing* Vol 29 (1981) pp 679–683

8 Baas, L and Bourne, J R 'A rule base microcomputer system for electroencephalogram evaluation' *IEEE Trans. Biomed. Eng.* Vol 31 No 10 (1984) pp 660–664

9 Jagannathan, V, Bourne, J R, Jansen B H and Ward, J W 'Artificial intelligence methods in quantitative electroencephalogram analysis' *Comput. Programs in Biomed.* Vol 15 (1982) pp 249–258

10 Zadeh, L A 'Fuzzy sets' *Information & Control* Vol 8 (1965) pp 338–353

11 Winston, P H *Artificial Intelligence* (3rd Ed.) Addison-Wesley, UK (1992)

12 Hashim, S H and Seyer, P *Turbo Prolog Advanced Programming Techniques* Tab Books, USA (1988)