EEG-based communication: a pattern recognition approach

William D. Penny[a], Stephen J. Roberts[a] and Maria J. Stokes[b]

[a]Department of Electrical Engineering
Imperial College,
London SW7 2BT, UK

[b]Research Department,
Royal Hospital for Neuro-disability,
London SW15 3SW, UK.

May 17, 1999

Abstract

The overall aim of this research is to develop an EEG-based computer interface for use by people with severe physical disabilities. The work comprises an 'offline' study and an 'online' study, the offline study establishing principles of interface design and the online study putting those principles into practice. The work focuses on using EEG signals to drive one-dimensional cursor movements on a computer screen and our approach is characterised by our emphasis on pattern recognition methods rather than on biofeedback training. Two key technical features further define our approach: firstly, we use dynamic rather than static pattern recognition algorithms and, secondly, we infer not just the parameters of our classifier but also the uncertainty on those parameters. Both of these features result in more robust cursor control.

1 Introduction

The ultimate aim of this research is to develop an EEG-based computer interface for use by people with severe physical disabilities. This would, for example, facilitate interaction with a wordprocessor package or manipulation of various environmental controls.

Our work is inspired by the programme of research at the Wadsworth centre. In particular, the work reported by Wolpaw *et al.* [18] who used fixed features of the EEG (8-13Hz activity) to drive one-dimensional cursor movements on a computer screen. Subjects learn to drive cursor movements via a biofeedback process. Though successful, the process is rather long, taking up to several weeks before users can achieve the accuracy required for a practical communication device.

Our approach relies less on biofeedback training and more on the use of pattern recognition methods, the idea being that the burden of communication be met by the user adapting to the computer *and* the computer adapting to the user. In this type of scheme, cursor movements are generated by the output of a pattern classifier such as a neural network.

The approach is somewhat similar to that used in the Graz BCI [14] but is different in two important technical respects. Firstly, we infer not just the parameters of our classifiers (eg. weights in a neural net) but also the uncertainty on those parameters. This allows us to estimate the uncertainty associated with each subsequent classification. If the cursor is then allowed to move only for high confidence classifications the system has some ability to perform automatic rejection of muscle artifacts and automatic rejection of trials containing irrelevant cognitive components (eg. during lapses of concentration). Secondly, we use dynamic classifiers such that the cursor movement at a given time step is dependent on cursor movements at previous time steps. Both of these features lead to more robust cursor control [13, 16].

A further aspect of our work is an exploration of the cognitive tasks used to provide a starting point for communication. To date, we have looked at motor imagery and mental arithmetic tasks. This aspect of our work is similar to that of the Colorado group [6, 1] but is different both in the pattern recognition approach and in the particular choice of tasks.

2 Offline studies

Our research into EEG-based communication began in 1996. At that time, whilst there was some anecdotal evidence from biofeedback experiments [18, 14] to suggest that motor imagery can be identified from the background EEG, there were no formal experiments to suggest that this is indeed the case. Or indeed, any information on what proportion of subjects these patterns could be detected in or with what accuracy.

To clarify the situation we recorded EEG from seven subjects performing cued imagined hand movements [13]. Control recordings were also made to ensure we were not picking up stimulus-related activity. The EEG was recorded from a single reference electrode and two 11-electrode arrays placed over left and right sensorimotor cortex (a total of 23 electrodes).

Laplacian operators were applied to estimate local activity at three sites over each sensorimotor cortex. Analysis of mu-rhythm power in the resulting signals showed that imagined hand movements could be identified in six out of seven subjects with a typical accuracy of 70%. The most discriminative electrode positions were found to be 3cm posterior to C3 and C4. Extraction of complexity features [15] showed that, in four out of seven subjects, imagined hand movements could be discriminated from background EEG activity with a typical accuracy of 80%. A comparison of classification accuracy using neural network versus logistic classifiers showed no benefit in using neural nets; logistic regression was sufficient [13].

This research was useful in concretely establishing that motor imagery signals could be picked up by spectral features and that, in principle, they could be used to drive cursor movements. It also identified the best position to place a smaller number of electrodes. Similar findings have also been made in recent research by McFarland *et al.* [8].

3 Online system

Whilst the above research established possible principles of EEG-based communication a number of practical issues remained. Firstly, the pattern recognition must be implemented online. But due to the yearly doubling of computer speed this is one of the lesser problems. Secondly, we would like to use a small number of electrodes and thirdly we would like to have a free-running communication protocol (self-paced movements) rather than communicating in response to a cue (although this last point is something of an open issue).

These factors have influenced the design of our EEG-based interface. It uses only three electrodes, a single isolation amplifier and a 266Mhz PC. The electrodes are placed at C3'-C4' (3cm behind C3 and C4 in the 10/20 system) and a reference electrode is placed over the right mastoid.

In experiments with the interface the communication protocol is as follows. Subjects move a cursor up or down a computer screen in order to hit targets that appear either at the top or at the bottom. The number of discrete cursor positions and other details are identical to that described in [18]. These details are not important, however, as to date, we choose to analyse the data on a segment-by-segment basis (see later).

Subjects move the cursor by performing different cognitive tasks and are given a maximum of ten seconds to hit each target. We have also carried out experiments where the cursor does not move. We have tried two different pairs of cognitive tasks; (i) motor imagery versus a baseline task and (ii) motor imagery versus a maths task. For the motor imagery tasks subjects were asked to imagine opening and closing their hand (right or left according to handedness), and for the maths tasks subjects were asked to serially subtract seven from a large number. Further experimental details are available in [12].

Cursor movements were generated by extracting autoregressive (AR) features from the EEG and classifying them using a logistic regression model. Specifically, a 'lagged-autoregressive' (LAR) model was applied to short-overlapping windows of data. LAR models can pick up relevant changes in EEG signals in whatever frequency band they occur and, in our experience, are superior to AR models in being less sensitive to noise [12].

3.1 Handling uncertainty

The LAR features are classified using a logistic regression model trained using the Bayesian evidence framework [7]. This procedure estimates both the classifier weights and the distribution of those weights. The distribution captures the fact that the classifier is not entirely certain as to how to classify some inputs. If this uncertainty is taken into account when making a new prediction (as it should be) then the correct predictive output to use is the 'moderated' output. Moderation in a two-class problem changes the output to a value nearer to 0.5 (the class prior) by an amount which is proportional to the uncertainty on the weights. Moderated outputs are typically better than unmoderated outputs in terms of the likelihood of predictions [7]. It is possible to make even more robust cursor movements by choosing to not move the cursor if the classifier output is not sufficiently different to 0.5. This is known as 'rejection'.

3.2 Dynamic models

For EEG data, averaging classifier outputs over a number of consecutive data segments is known to significantly increase classification accuracy [1]. This is known as 'temporal smoothing'. One approach to temporal smoothing is to average classifier responses not in the output space, but in the space of activations, or the 'latent' space. This scheme arises from considerations of how to make optimal decisions in a 'committee' of classifiers [9] where each committee member makes a prediction from a different time point. Optimal smoothing periods can be established via cross-validation; a typical period is two seconds.

Temporal smoothing can also be achieved with a Hidden Markov Model (HMM) [10]. HMMs have a number of discrete states (eg. one for cursor up and one for cursor down) each of which is associated with particular characteristics of the data. In our case, these would be spectral characteristics as captured by an LAR model. Transition from the current state, i, to a new state, j is determined by (i) the state transition probability p_{ij} and (ii) the characteristics of the new data point (eg. LAR vector). The amount of temporal smoothing is determined implicitly by the matrix of state transition probabilities and, importantly, these can be learnt from the data set.

3.3 Results

We report results from online experiments by analysing the EEG data on a segment-by-segment basis (we could measure the proportion of targets hit but this would tie the results rather strongly to the details of the particular communication protocol which is yet to be optimised). Figure 1 shows classification accuracies on seven subjects for (a) stationary cursor recordings and (b) moving cursor recordings for two pairs of cognitive tasks. In the moving cursor trials four out of seven subjects achieved at least 75% accuracy.

Once we have trained a classifier to discriminate between two different cognitive tasks on the basis of LAR features it is interesting to then go back and look at what

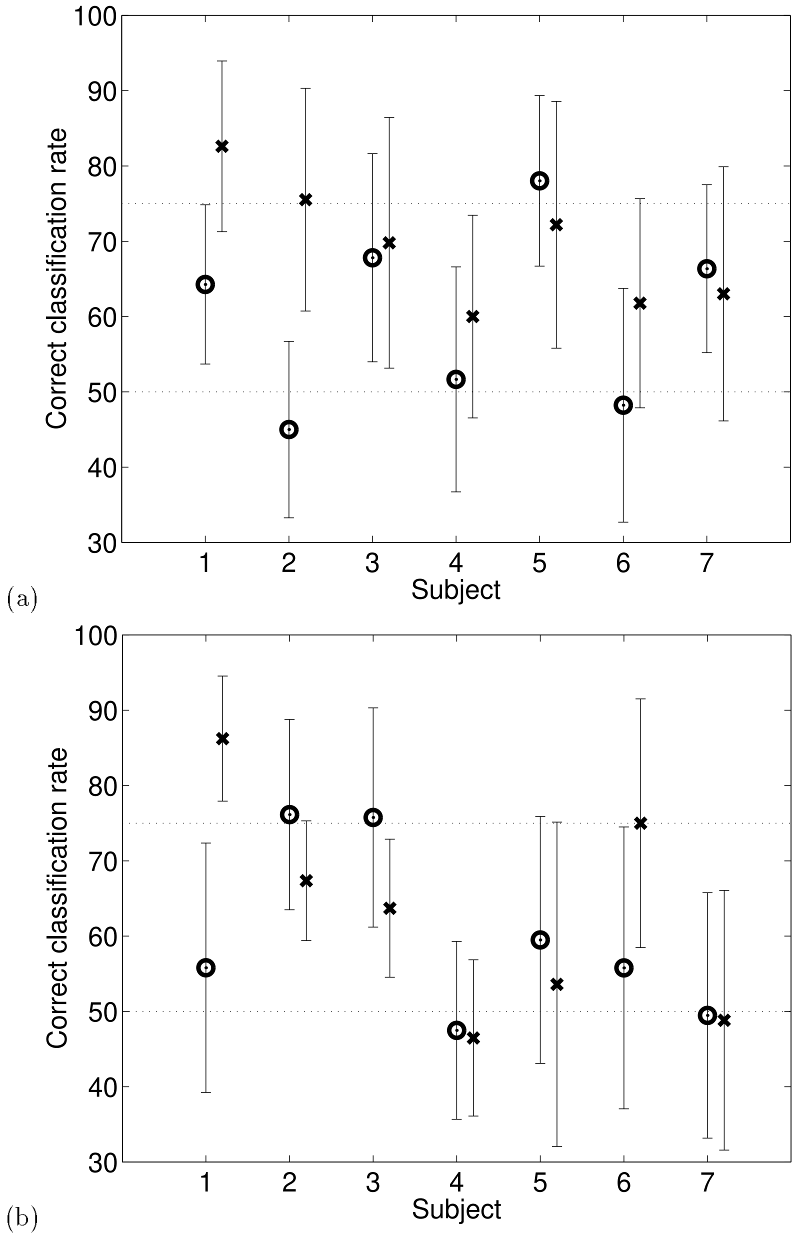

Figure 1: The figure shows classification accuracies on seven subjects for the imagery versus baseline pairing (marked as circles) and the imagery versus maths pairing (marked as crosses). The error bars are plotted at one standard deviation. Each result is from approximately 1200 data segments derived from 12 repetitions of each task where each task lasts up to ten seconds. Plot (a) is for **stationary cursor** recordings: for the imagery versus baseline pairing, discrimination is good in one subject, significant in three and not possible in the other three. For the imagery versus maths pairing, discrimination is good in four subjects and significant in three. Plot (b) is for **moving cursor** recordings: for the imagery versus baseline pairing, discrimination is good in two subjects and not possible in the other five. For the imagery versus maths pairing, discrimination is good in two subjects, significant in two subjects and not possible in the other three.

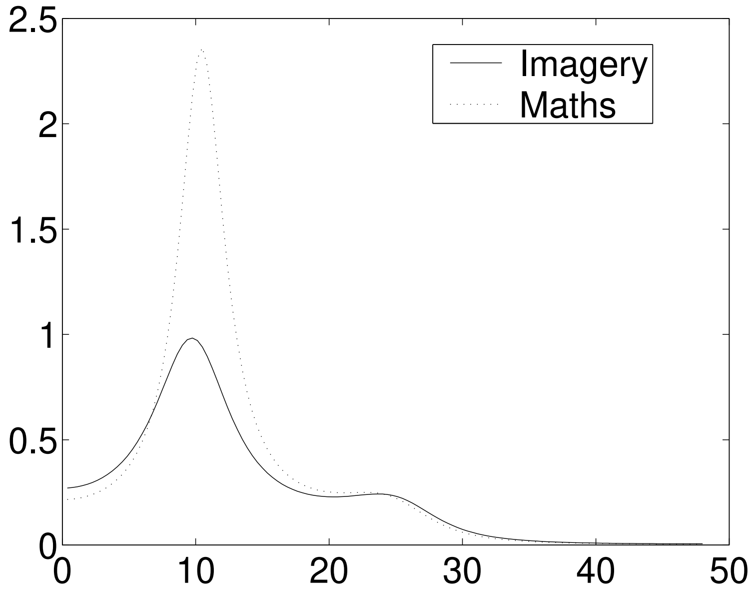

Figure 2: Enhanced spectra for subject five performing motor imagery and maths tasks.

are the typical LAR features and corresponding spectra associated with each task. This is equivalent to the enhanced averaging method described by Gevins and Morgan [5]. Figure 2 shows enhanced spectra for subject 5 performing the imagery versus maths pairing. This is quite representative of all the subjects with the majority of differential activity in the mu-band (8-13 Hz). Some subjects also showed differences in the theta (4-7 Hz) and beta (14-20 Hz) bands.

4 Discussion and future plans

We note that four out of the seven subjects achieved at least 75% accuracy in the moving cursor trials. They would would therefore be able to immediately use a wordprocessor package via our EEG-based interface using, for example, the protocol developed by Birbaumer *et al.* [4]. Two of these subjects would use an imagery versus baseline strategy and the other two imagery versus maths.

But what of the other three subjects and what of interfacing with other devices (wheelchair control, for example, may require much better than 75% accuracy) ? As the present technology stands they would not be able to use the interface but we can see the technology improving by following one of two distinct research paths.

Firstly, we could persevere with methods involving little or no biofeedback training and focus on other ways of improving classification accuracy. One method might be to also use 'readiness' potentials (RPs) [2] or 'slow cortical potentials' [4]. Recent research shows that RPs and spectral changes (mu-rhythm desynchronization) are uncorrelated [17]. The RPs would therefore provide an additional source of information. A second method is to look at using additional features such as complexity [15] or to pre-process the signals in more adventurous ways eg. using an Independent Component Analysis embedding [3]. A third method is to use HMMs. HMMs have not, as yet, been applied to the whole database but qualitative results on one subject suggest the approach is promising. The main strength of HMMs is in analysing *sequences* of data (via Viterbi decoding). We therefore envisage that HMMs can best be utilised by retrospectively analysing sequences of intended cursor movements in order to better decode the transmitted message. Use can also be made of error-correcting codes.

Secondly, we could train subjects using a biofeedback approach. This involves the interaction of two adaptive controllers; the user and the computer. One promising approach for handling this is to use nonstationary classification algorithms [11] which acknowledge that the statistics of each class (ie. how the user moves the cursor up or down) can change with time.

We plan to incorporate all of the more recent technical developments into a new version of the recognition software and to test the resulting system on a large number of normal subjects. The interface will then be tested on patients having severe neurological disabilities but known to be cognitively aware.

Acknowledgements

The authors would like thank Nahid Haque, EEG Research Assistant at the Royal Hospital for Neuro-disability, and the subjects who took part in the experiment. W. D. Penny was supported by the UK Engineering and Physical Sciences Research Council (EPSRC), grant number GR/K79062, and equipment and technical support were funded by the Neuro-disability Research Trust.

References

[1] C. Anderson and Z. Sijercic. Classification of EEG Signals from four subjects during five mental tasks. In *Proceedings of the conference on engineering applications in neural networks (EANN'96)*, pages 407–414. Systems Engineering Association, 1996.

[2] R. Cunnington. Movement-related potentials in Parkinson's disease: motor imagery and movement preparation. *Brain*, 120:1339–1353, 1997.

[3] D. Lowe. Feature space embeddings for extracting structure from single channel wake EEG using RBF networks. In *Neural networks for signal processing: Proceedings of the IEEE workshop*, pages 428–437, 1998.

[4] N. Birbaumer et. al. A spelling device for the paralysed. *Nature*, 398:297–298, 1999.

[5] A.S. Gevins and N.H. Morgan. Classifier-directed signal processing in brain research. *IEEE Transactions on Biomedical Engineering*, 33(12):1054–1068, 1986.

[6] Z.A. Keirn and J.I. Aunon. A new mode of communication between man and his surroundings. *IEEE Transactions on Biomedical Engineering*, 37(12):1209–1214, 1990.

[7] D.J.C. Mackay. The evidence framework applied to classification networks. *Neural Computation*, 4(5):720–736, 1992.

[8] D.J. McFarland, L. McCane, L.A. Miner, T.M. Vaughan, and J.R. Wolpaw. EEG Mu and Beta Rhythm Topographies with Movement Imagery and Actual Movement. In *Society for Neuroscience Abstracts*, page 1277, 1997.

[9] W.D. Penny, D. Husmeier, and S.J. Roberts. The Bayesian Paradigm: second generation neural computing. In P.J. Lisboa, editor, *Artifical Neural Networks in Biomedicine*. Springer, 1999.

[10] W.D. Penny and S.J. Roberts. Dynamic models for nonstationary signal segmentation. *Computers and Biomedical Research*, 1998. To appear.

[11] W.D. Penny and S.J. Roberts. Dynamic Logistic Regression. In *International Joint Conference on Neural Networks (IJCNN)*, 1999.

[12] W.D. Penny and S.J. Roberts. Experiments with an EEG-based computer interface. Technical report, Department of Electrical and Electronic Engineering, Imperial College, 1999.

[13] W.D. Penny, S.J. Roberts, and M.J. Stokes. Imagined hand movements identified from the EEG mu-rhythm. *Journal of Neuroscience Methods*, 1998. Submitted.

[14] G. Pfurtscheller, D. Flotzinger, and J. Kalcher. Brain-Computer Interface - a new communication device for handicapped persons. *Journal of Microcomputer applications*, 16:293–299, 1993.

[15] S.J. Roberts, W. Penny, and I. Rezek. Temporal and spatial complexity measures for EEG-based brain-computer interfacing. *Medical and Biological Engineering and Computing*, 37(1):93–99, 199.

[16] S.J. Roberts and W.D. Penny. Real-time Brain-Computer Interfacing. Technical report, Department of Electrical Engineering, Imperial College, 1999. Also submitted to Engineering in Medicine and Biology.

[17] C. Toro, G. Deuschl, R. Thatcher, S. Sato, C. Kufta, and M. Hallett. Event-related desynchronization and movement-related cortical potentials on the ECOG and EEG. *Electroencephalography and Clinical Neurophysiology*, 93(5):380–389, 1994.

[18] J.R. Wolpaw, D.J. McFarland, D.J. Neat, and C.A. Forneris. An EEG-based Brain-Computer Interface for Cursor Control. *Electroencephalography and Clinical Neurophysiology*, 78:252–259, 1991.