EEG Signal Classification with Different Signal Representations

Charles W. Anderson
Dept. of Computer Science
Colorado State University
Fort Collins, CO 80523
anderson@cs.colostate.edu

Saikumar V. Devulapalli
Dept. of Computer Science
Colorado State University
Fort Collins, CO 80523

Erik A. Stolz
Dept. of Electrical Engineering
Colorado State University
Fort Collins, CO 80523

Abstract

If several mental states can be reliably distinguished by recognizing patterns in EEG, then a paralyzed person could communicate to a device like a wheelchair by composing sequences of these mental states. In this article, we report on a study comparing four representations of EEG signals and their classification by a two-layer neural network with sigmoid activation functions. The neural network is implemented on a CNAPS server (128 processor, SIMD architecture) by Adaptive Solutions, Inc., gaining a 100-fold decrease in training time over a Sun Sparc 10 for a large number of hidden units.

1 INTRODUCTION

Computerized analysis of EEG signals has evolved over the past three decades [5] with much of the effort directed towards a better understanding of the functioning of the brain. The work reported here has a different goal—to extract information from EEG signals with which mental states can be discriminated and thereby serve as a mode of communication for a paralyzed person. In the literature, we find two approaches towards this goal.

One approach is based on the discovery that a characteristic signal appears in the EEG approximately 300 ms following the occurrence of a relatively rare, but expected, stimulus. Such signals are referred to as event-related potentials, or ERP's. An example of how ERP's can be used to communicate with a computer is the work of [3], who used ERP's to detect which letter of the alphabet a human subject wished to select. This kind of interaction between a person and stimulus device is currently too cumbersome and slow to be useful in a real-time control application.

Much more practical would be a system for detecting patterns in normal EEG without the aid of an external stimulus device. This is the second approach, often referred to as spatial analysis, because patterns are sought in EEG signals simultaneously recorded from multiple electrodes. A number of studies have found differences in the power of the alpha band (8-13 Hz) in signals recorded from left and right hemispheres, depending on the tasks [4, 2, 9]. Asymmetries were most reliably found for motor tasks. In a recent study, [11] used a frequency-based representation of EEG from two electrodes to obtain 85% accuracy in predicting movements of the finger of a subject moving a joystick left or right.

The detection of patterns in EEG produced from normal mental states is a very difficult problem. EEG signals are recorded by surface electrodes and can contain noise as a result of electrical interference and movement of the electrodes on the scalp. Another problem is that EEG can be corrupted by eye blinks and other muscular activity that produce signals of greater magnitude than produced by cortical activity. Other problems are more cognitive in nature; the concentration of a person can vary while the person is supposedly performing a single mental task.

The work described in this article is based on previous work [7] using a Bayesian classifier trained and tested on a small subset of data recorded from subjects performing several mental tasks. We extended their study by 1) replacing the Bayesian classifier with neural networks of various sizes trained using error back-propagation, 2) using all of the data from each recording session, and 3) testing four signal representations. The objective of these experiments was to determine which of the four representations results in the best classification accuracy. If the information needed to discriminate mental state can be extracted from the unprocessed EEG signals, or at most preprocessed by projecting to a relatively small number of principal components, then we can dispense with other forms of preprocessing, such as the frequency analysis used in [7].

In one other study [10] neural networks were applied to classify EEG signals collected in a paradigm similar to that of [7]. Kohonen's SOM algorithm [8] was used to train a matrix of units to identify clusters of similar patterns and associate each cluster with a particular mental task. They trained their classifier on data for all tasks performed by one subject in one recording session and tested the resulting classifier on data from other sessions and other subjects. Most tests showed very poor classification accuracy, though the accuracy tended to be higher for some tasks, particularly the mental arithmetic task described in the next section.

In Section 2, we describe the methods used to collect, process, and classify the data. Results of experiments with data from a single subject and a pair of tasks are shown in Section 3 and conclusions are discussed in Section 4.

2 METHOD

2.1 DATA COLLECTION

All data used in this study were recorded during the previous study [7]. Selection of the set of mental tasks was guided by [4] whose results showed detectable hemispheric differences for certain tasks. Here we focus on two tasks:

Baseline—Alpha Wave Production The subject is asked to relax and try to think of nothing in particular. This will be considered to be the baseline session for alpha wave production, and other asymmetries.

Mental Arithmetic The subject is given a non-trivial multiplication problem to solve and, as in all of the tasks, is instructed not to vocalize or make overt movements while solving the problem. An example of such a task would be to multiply the numbers 49 times 78. The problems are non-repeating and designed so that an immediate answer is not apparent.

Subjects were seated in a sound-proof, dimly-lit, room. As shown in Figure 1, electrodes were placed at O_1, O_2, P_3, P_4, C_3, and C_4, standard electrode locations in the 10-20 System [6]. The electrodes were connected to Grass 7P511 amplifiers that bandpass filtered the signals at 0.1-100 Hz. The EEG signals were sampled at 250 samples per second and digitized with 12 bits of accuracy. Data was recorded from each subject for a duration of 10 seconds while the subject was performing a single task with their eyes open. Each session resulted in 250 samples/second x 10 seconds x 6 channels, or 15,000 values. For the experiments reported here, artifacts, such as eye blinks, were not removed.

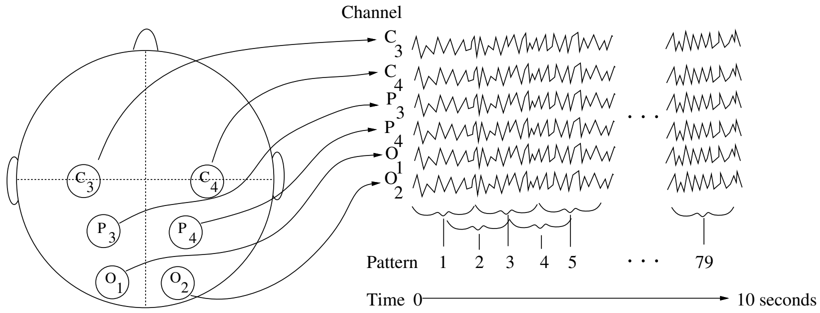

Figure 1: Location of the six surface electrodes and division of the six time series into overlapping quarter-second windows. (Data shown is fictitious.)

2.2 UNPROCESSED REPRESENTATION

In [7] it was found that quarter-second segments of the 10 second data resulted in classification accuracy approximately the same as that obtained from two-second segments. Therefore, we divided the data into 79 overlapping, quarter-second segments consisting of 62 samples from each of the 6 channels, or 372 values per segment, as shown in Figure 1. We collected 79 such patterns from one subject performing the baseline and mental arithmetic tasks 10 times each, for a total of 1,580 patterns.

It is generally believed that frequencies above 40 Hz convey little information related to mental state. To test this, a second unprocessed representation was created by applying an FIR lowpass filter with a cut-off frequency of 40 Hz. This was followed by down-sampling by a factor of 2 to reduce the number of data samples by half, resulting in an effective sampling rate of 125 Hz. Pre-filtering is required to prevent aliasing.

2.3 K-L REPRESENTATION

When classifying high-dimensional data, equivalent or better generalization accuracy is often achieved by classifying data obtained by projecting the original data onto the first n eigenvectors, where n is much smaller than the dimensionality of the original data. To perform this Karhunen-Loève decomposition (K-L), the covariance matrix of the mean-subtracted set of the 1,580, 372-dimensional, patterns was calculated.

The number of eigenvectors to project to can be chosen empirically or determined by examining the eigenvalues. One estimate of dimensionality is the global Karhunen-Loève estimate, given by the index i for which $\lambda_i/\lambda_{max} \leq 0.01$, where the λ_i are in decreasing order for $i = 1, 2, \ldots$. For the baseline and mental arithmetic data used in our experiments, i = 50. Therefore, the K-L representation was formed by projecting each 372-dimensional pattern onto the first 50 eigenvalues.

2.4 FREQUENCY-BAND REPRESENTATION

In [7] features were extracted from spectral density estimates using asymmetry ratios given by $(R-L)/(R+L)$, where R is the area under the spectral density curve of a right hemisphere channel for a specific frequency band and L is defined similarly for the corresponding left hemisphere channel. These asymmetry ratios were calculated for each possible right-to-left combination of channels and for each of four frequency bands: delta (0-3 Hz), theta (4-7 Hz), alpha (8-13 Hz), and beta (14-20 Hz). This results in 36 asymmetry ratios. In addition the 24 power values themselves (R and L) were added for a total of 60 features. In the current study, the spectral density was estimated from autoregressive (AR) parameters calculated using the Burg method [12]. An AR model of order 6 is used here since it yielded good results in [7].

2.5 CLASSIFICATION

Classification was performed by fully-connected, feed-forward neural networks with single hidden layers of varying size. All units use an asymmetric sigmoid activation function. The work reported here involved a single pair of tasks, the mental arithmetic and baseline tasks, so the networks require only one output unit. An output value greater than 0.5 represents the mental arithmetic task and a value less than 0.5 represents the baseline task. The target value for the output during training is set to 0.9 for mental arithmetic and 0.1 for baseline.

Our data was obtained from ten different 10-second recording sessions for each task, so it naturally divides into ten parts. We trained on eight parts, used a ninth part as a cross-validation, and tested on the tenth part. These three partitions were used as follows. The network is trained on the training set for a given number of epochs. After each epoch, the RMS error for the current network is calculated for the cross-validation set. After training, the weights for the epoch at which the cross-validation error was lowest are obtained and the RMS error of the network using these weights is calculated for the test set. This test error is the estimate of how well the network will classify novel data. Thus, we train, cross-validate and test on data from different recording sessions. With ten parts, we can choose 90 different pairs for cross-validation and testing. The results in the next section are averages over 90 runs constructed in this way.

Training was performed with the `bnlib` error back-propagation library on a 128-node CNAPS Server II from Adaptive Solutions, Inc.. [1] present details of the implementation. The CNAPS implementation ran for 4 minutes to train a 40 hidden unit net for 1,000 epochs on the unprocessed data representation, whereas a Sparc 10 required about 9 hours, over 100 times longer.

Networks were trained on the training data until the squared error for 200 successive epochs did not decrease. The epoch at which the squared error on the cross-validation data was at a minimum was identified, and the network was retrained to this epoch, starting with the same initial weights. At this point the network's weights were written to a file. All of this was performed on the CNAPS hardware. The network was evaluated on a Sun Sparc by calculating the number of test patterns classified correctly. A pattern is said to be correctly classified if the network's output falls on the correct side of 0.5.

3 RESULTS

Table 1 shows the results of all classification experiments as the average percent of test patterns classified correctly, out of 158 patterns for all but the unprocessed, 125 Hz, representation, which consisted of 154 patterns. This table includes 90% confidence intervals, based on 90 repetitions.

Clearly the best classification accuracy is achieved with the frequency-band representation, giving an average accuracy of about 74% for a network with 40 hidden units, though the accuracy varies little for other network sizes, including a network with a single hidden unit. This shows that the dependence between the signals in the various representations and the correct classification is primarily linear on average. The performance of the raw and K-L representations is significantly lower, ranging from 50% to 53%. These results suggest that the energy within standard frequency bands is more useful in discriminating the two

Hidden Units	Unprocessed (250 Hz)	Unprocessed (125 Hz)	K-L	Frequency-Bands
40	53.2 ± 0.6	51.1 ± 0.6	51.7 ± 0.6	73.9 ± 0.7
5	52.8 ± 0.6	52.1 ± 0.8	51.8 ± 0.7	72.9 ± 0.6
1	52.1 ± 0.5	52.1 ± 0.7	50.4 ± 0.5	73.1 ± 0.6

Table 1: Percent of test patterns classified correctly for different sized networks and different signal representations.

mental tasks than is the raw data, or dimensionally-reduced data. This hypothesis must be tested by further experimentation. It appears that the performance of the raw and K-L representations is increasing with network size, but the differences are not statistically significant. We have not yet tried networks with more than 80 hidden units.

The results reported above dealt with each quarter-second segment in isolation. One way to consider more than one interval is to average the output of the network over successive segments. We have investigated this using the frequency-band representation for one combination of training, cross-validation, testing data and obtained the following preliminary results.

Table 2 shows how the number of successive quarter-second segments over which the network output is averaged affects classification accuracy. For the mental arithmetic task, only five segments are needed to get 100% classification accuracy for the single test data set considered. However, 28 successive segments were required to get 100% accuracy for the baseline task. Thus, correct discrimination between baseline and mental arithmetic could be achieved by averaging over 28 successive segments, or about 3.6 seconds (1/4 + 27 x 1/8).

One hypothesis for this difference is that a subject performing the baseline task is less focused on a particular set of mental states than is a subject performing the mental arithmetic task. Plots of the network output, averaged over five segments, is shown for both tasks in Figure 2. Recall that the target for the network's output for the mental arithmetic task is 0.9, while for the baseline task it is 0.1. These plots suggest that while performing the baseline task, the subject's mental state varied considerably. During the middle of the 10 second recording session, the subject's mental state more resembled the mental arithmetic state than the baseline state, using the frequency-bands representation.

When this process of averaging over successive segments is repeated for different groupings of the data into training, validation, and test sets, we occasionally find cases for which the classification accuracy does not converge on 100% as more segments are averaged. We are currently performing experiments to determine the average accuracy over different groupings of the data. Preliminary results indicate that the accuracy on average converges near 90% accuracy as the averaging is extended to cover full 10 second recording period.

Number of Successive Segments	Mental Arithmetic Task	Baseline Task
1	79%	67%
2	88	76
3	92	84
4	93	78
5	100	81
.	.	.
.	.	.
.	.	.
27	100	98
28	100	100
.	.	.
.	.	.
.	.	.

Table 2: Percent correct when network output is averaged over successive segments.

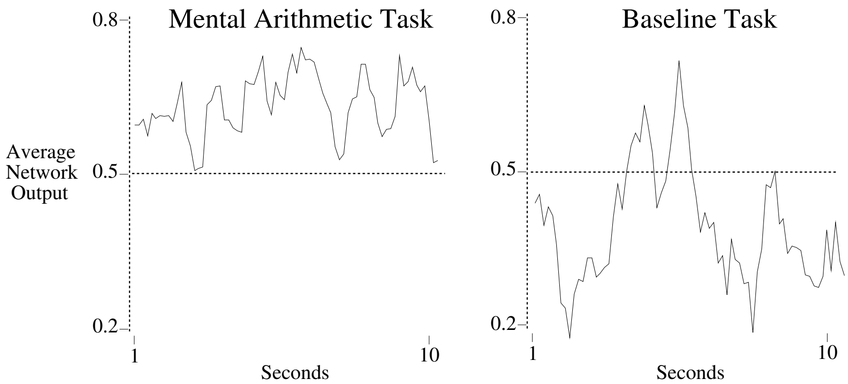

Figure 2: Network output using frequency-band representation averaged over five successive segments.

4 CONCLUSIONS

The frequency-based representation resulted in significantly more accurate classification than the unprocessed or K-L representations. A second conclusion from this study is the utility of the parallel implementation of the error back-propagation algorithm. A much greater number of network sizes and initial weight vectors could be evaluated on the CNAPS server than could be completed in a comparable amount of time on a serial machine.

The practicality of this result is limited by the long period, 3.6 seconds, that was required to get reliable classification. This is too long for use in controlling a wheelchair. However, this estimated time period is conservative, because averages were over quarter-second segments that overlapped only by an eighth of a second. Instead, the network output could be calculated for all quarter-second segments, with successive segments overlapping by one time sample, or $1/250^{th}$ of a second. This might reduce the length of the data sample needed to reliably discriminate tasks.

We are currently studying other representations, including AR model coefficients, Hidden Markov Models, LVQ, and "bottleneck" networks to perform nonlinear principal components analysis. Generalization studies across subjects with additional tasks are being performed.

Acknowledgements

The authors would like to thank Jorge Aunon and Zachary Keirn for providing the EEG data, Ed Orosz for installing and interpreting the data. This research was funded by the National Science Foundation through grant IRI-9202100.

References

[1] Charles W. Anderson, Sai Devulapalli, and Erik Stolz. Determining mental state from EEG signals using neural networks. *Scientific Programming*, to appear.

[2] J. C. Doyle, R. Ornstein, and D. Galin. Lateral specialization of cognitive mode: II EEG frequency analysis. *Psychophysiology*, 11:567–578, 1974.

[3] L. A. Farwell and E. Donchin. Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and Clinical Neurophysiology*, 70:510–523, 1988.

[4] D. Galin and R. E. Ornstein. Hemispheric specialization and the duality of consciousness. *Human Behavior and Brain Function*, 1973.

[5] Alan S. Gevins. Overview of computer analysis. In A. S. Gevins and A. Rémond, editors, *Methods of Analysis of Brain Electrical and Magnetic Signals*, volume 1 of *EEG Handbook*, chapter 3, pages 31–83. Elsevier Science Publishers Biomedical Division, 1987.

[6] H. Jasper. The ten twenty electrode system of the international federation. *Electroencephalographic Clinical Neurophysiology*, 10:371–375, 1988.

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

[8] T. Kohonen. *Self-Organization and Associative Memory*. Springer-Verlag, Berlin, 1984.

[9] R. H. Kraft, O. R. Mitchell, M. L. Languis, and G. H. Wheatley. Hemispheric asymmetry during six- to eight-year old performance of piagetian conservation and reading tasks. *Neuropsychologia*, 22:637–643, 1984.

[10] Shiao-Lin Lin, Yi-Jean Tsai, and Cheng-Yuan Liou. Conscious mental tasks and their EEG signals. *Medical & Biological Engineering & Computing*, 31:421–425, 1993.

[11] G. Pfurtscheller, D. Flotzinger, and J. Kalcher. Brain-computer interface - a new communication device for handicapped persons. *Journal of Microcomputer Applications*, 16(3):293–299, 1993.

[12] J. G. Proakis and D. G. Manolakis. *Digital Signal Processing*. MacMillan, New York, 1992.