Claremont Colleges
Scholarship @ Claremont

All HMC Faculty Publications and Research

HMC Faculty Scholarship

12-1-2003

Enabling Computer Decisions Based on EEG Input

Benjamin J. Culpepper

Robert M. Keller
Harvey Mudd College

Recommended Citation

Enabling Computer Decisions Based on EEG Input

Benjamin J. Culpepper and Robert M. Keller, *Member, IEEE*

Abstract—**Multilayer neural networks were successfully trained to classify segments of 12-channel electroencephalogram (EEG) data into one of five classes corresponding to five cognitive tasks performed by a subject. Independent component analysis (ICA) was used to segregate obvious artifact EEG components from other sources, and a frequency-band representation was used to represent the sources computed by ICA. Examples of results include an 85% accuracy rate on differentiation between two tasks, using a segment of EEG only 0.05 s long and a 95% accuracy rate using a 0.5-s-long segment .**

Index Terms—**Brain-computer interface (BCI), electroencephalogram (EEG), independent component analysis (ICA), neural networks.**

I. INTRODUCTION

FOR QUITE some time now, we have been using our physical actions to control computers. The most primitive interfaces (e.g., punch cards) were designed to accommodate the requirements of the computer, rather than to provide the fastest and most intuitive link between the human and the machine. However, as interface technology progressed, efforts to improve this link became more concentrated on making things easier for the human: The keyboard was invented, and then the mouse. Although there are many other, less widely used techniques for controlling computers, these tools have become the modern standard. The others include eye-tracking devices that perform the same basic function as a mouse, and speech recognition, which can be used to perform the functions of both the mouse and keyboard. Different types of input devices appeal to people with different needs (for example, the physically handicapped), and some are more appropriate for certain applications than others. Yet, they all share one similarity that is particularly relevant to this study—the requirement that the user encode his or her intentions into a stream of physical movements, which is then decoded by the computer into a set of signals to be used for control.

It seems natural to wonder whether the physical encoding process is necessary. Can we not convey our intentions to the computer in a more direct fashion by skipping the physical layer entirely? Assuming we had some method of sensing and differentiating between two or more of the mental processes, we could use these few states as a simple alphabet, from which a more complicated language of control could be built. The process of encoding intentions would take place entirely inside a person's head. For example, instead of thinking "I want that window to

move to the foreground, so I will move the mouse to it and click," a user might think, "I want that window to move to the foreground, so I will think A, B, C, A," where the latter is a sequence of mental states recognized as signals by the computer. This particular sequence would, in turn, trigger the corresponding decision in the computer. A further question asks whether this mental encoding process is even necessary. Can we use the power of the computer to decode our intentions as we think them? Unfortunately, the latter question is beyond the scope of this paper; at present, we do not know enough about how the underlying brain processes encode thoughts to perform this task well. This paper concentrates on the first approach: discrimination between mental states based on thought patterns for prescribed mental stimuli.

Electroencephalogram (EEG) signals are an important source of information for studying the underlying brain processes that make up our thoughts and actions. EEG recorders with up to 256 electrodes are currently in use, and experiments using them produce large amounts of raw data. Thus, the question naturally arises: How much can we find out about the brain's activities from all that data? Can we read a person's mind by properly deciphering their EEG? It has been known for some time that this is possible, to a modest extent. Recently, more than a few researchers have investigated using EEG signals as a new way of conveying intentions to a computer. EEGs produced during a very limited set of mental tasks can be classified according to tasks. Offline analysis of multichannel EEG has been performed by the groups of Wolpaw and Pfurtscheller since 1992 [7], [8], and more recently by Anderson *et al.* [1], [2].

In this paper, we describe the methods and results of our experiments with EEG signals recorded from one subject while the subject performed five mental tasks. The EEG signals were processed by independent component analysis (ICA) to remove artifacts, and presented to three-layer artificial neural network classifiers using a frequency-band representation.

II. RELATED WORK

Prior to acquiring our own EEG signals, we experimented with the data used in [1] and [3], both of which investigated the classification of five different mental tasks: a baseline resting task, mental multiplication, geometric figure rotation, mental letter composition, and visual counting. These tasks were chosen by Keirn and Aunon to invoke hemispheric brainwave asymmetry [3].

In [3], Keirn and Aunon recorded data from seven subjects using six channels. They transformed the data into features based on spectral estimates calculated from both the Fourier transform of the windowed autocorrelation function and a scalar autoregressive (AR) model. Their features included asymmetry ratios and power values for each channel from

four standard frequency bands: delta; theta; alpha; and beta. Asymmetry ratios $(R - L)/(R + L)$, where R and L are power values from right- and left-hemisphere electrodes, were calculated for all combinations of right and left electrodes and frequency bands. A second set of features was composed of the autoregressive coefficients from all channels concatenated together. Their most significant result was that all pairs of tasks could be discriminated 84.6% of the time using the coefficients as features.

In [1], Anderson *et al.* classified EEG from one subject who performed five mental tasks. He used the six-channel data sampled at 250 Hz from [3] and modeled signals in a half-second window as sixth-order autoregressive processes. Attempting to differentiate between five tasks resulted in a 31%–54% range of classification accuracy. He found that averaging the output of the network over consecutive inputs improved performance. By averaging the output of the best performing network over 20 consecutive overlapping time windows, which amounted to 5 s of actual EEG data, Anderson was able to obtain a range of classification accuracy from 33% to 71%.

In [2], Anderson *et al.* evaluated different signal representations. They used the data from [3], and presented 1580 patterns of quarter-second data windows to neural networks with a single hidden layer, the size of which was varied from 1 to 40 nodes. Four different representations were tried: unprocessed; low-pass filtered; Karhunen-Loeve; and frequency-band. An unprocessed feature vector was formed by concatenating one quarter second (62 samples) of data from each of the six channels (372 values per pattern). The low-pass filtered representation was similar to this, except that a finite-impulse response (FIR) low-pass filter with a cutoff frequency of 40 Hz was applied to the samples, followed by down-sampling by a factor of two (every other sample was removed).

Anderson *et al.* found that the frequency-band representation yielded the best results: 73.9% classification accuracy with 40 hidden nodes. The other representations had classification accuracies that hovered around 50%. Although the network with 40 hidden nodes did outperform those with lesser numbers of nodes, it was not by much. In general, it helped very little to increase the number of hidden nodes, although it helped the frequency-band representation the most. This seems to indicate that the network that was trained using the frequency-band representation was learning the most.

In [8], Wolpaw and McFarland let four subjects teach themselves to control electrical potentials from two bipolar electrodes located around positions C3 and C4 (see Fig. 1). The potentials were used to control a cursor on a screen and the subjects had to move the cursor into predefined areas. After 20–30 training sessions, subjects could move the cursor into given corners of the screen with 41%–70% success rate. For comparison, a random cursor movement would have a 25% success rate.

In [6], Peters *et al.* classified brain states corresponding to the intention of movement in the left and right index finger and right foot. This classification was done using a "committee" of artificial neural networks each processing an individual channel of a 56-electrode EEG. The size of the committee was chosen for optimal accuracy, and was found to be 9–12, confirming a conclusion of Wolpaw and Pfurtscheller that was obtained differently.
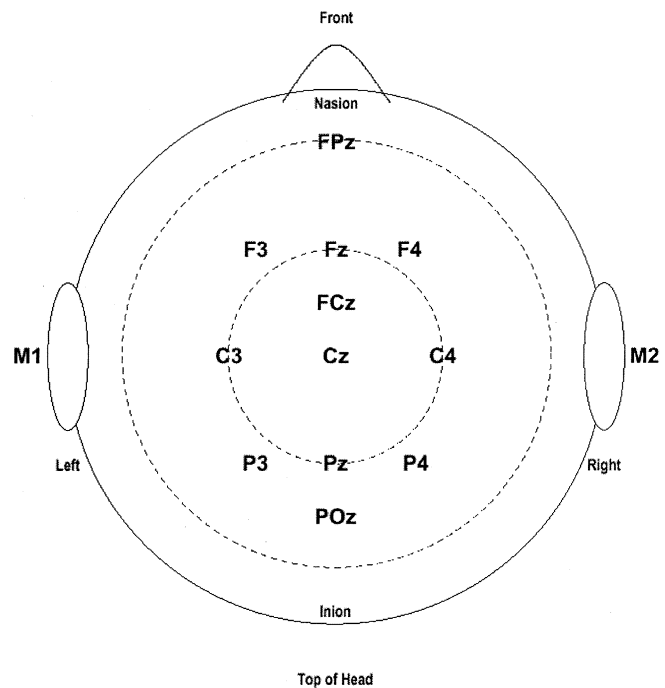

Fig. 1. Electrodes were placed at positions FPZ, F3, FZ, F4, FCZ, C3, CZ, C4, PZ, P3, POZ, and P4.

A classification rate of 83% was obtained when differentiating between two intended movements using a one second window of EEG. When the number of movements was increased to three and four, the classification accuracy dropped to around 70%.

Independent of the current experiment, which was performed in 1999, the following noteworthy related experiments were performed concurrently.

Haselsteiner and Pfurtscheller [12] compared the classification accuracy of standard multilayer perceptrons (MLPs) to time-dependent FIR MLPs. Their results demonstrated that the FIR-MLP classifier decreased the classification error rate over an ordinary MLP. The error rates were strongly tied to the subject, varying from 5% to 24% without time-dependency in the classifier, to 4%–21%. The decrease in error rate indicates that the temporal information of the input data can be used to improve classification results.

Further validation that ICA is a useful tool for segregating artifacts from EEG data and generating better feature sets for presentation to a pattern classifier is provided by Makeig *et al.* in [15], by Vigario *et al.* in [14], and by Zhukov *et al.* in [13]. In [15], Makeig *et al.* demonstrated that ICA-derived components can show much stronger spectral reactivity to motor events than activity measured for single scalp channels. In [14], Vigario *et al.* showed that the application of ICA to an averaged auditory evoked response accurately isolates the main response from subsequent components. In [13], Zhukov *et al.* effectively applied principal component analysis (PCA) in combination with ICA to the "inverse" problem: Given a subset of electrostatic potentials measure on the surface of the scalp, and the geometric and conductivity properties within the head, calculate the current sources and potential fields within the cerebrum.

In [16], Guger *et al.* investigated the possibility of using common spatial patterns to analyze EEG in real time and

provide feedback to the subjects. They demonstrated that three subjects, using feedback, could reduce their classification error within three days to 2%, 6%, and 14% in six to seven sessions. These results indicate that feedback is a promising technique for building a highly accurate brain-computer interface (BCI).

III. OUR EXPERIMENT

For the experiment presented here, we used the mental states chosen in [3], and took our own 12-channel data. We used ICA to remove the artifactual components found in our EEG, and trained three-layer neural networks of different sizes on the Fourier transform of a varied window length of EEG. Using a 0.5-s window of EEG, we obtained a classification accuracy of 94% when differentiating between the geometry and multiplication tasks. Using a 1/20th-s window of EEG, we obtained a classification accuracy of 85% when differentiating between the geometry and multiplication tasks.

A. EEG Signal Recording

The data for this experiment were taken using the neuroscience facilities at Pomona College, Claremont CA The subject was seated in a closet with dim lighting, a comfortable chair, and a computer running the program used to present visual stimulae to the subject, NeuroScan Inc.'s STIM 2.0 (http://www.neuro.com/neuroscan). The computer had internal fans that produced some slight noise, but other than that, the closet was fairly silent. A QuikCap-64 (http://www.neuro.com/neuromed/quikcap.htm) was used to record from positions FPZ, F3, FZ, F4, FCZ, C3, CZ, C4, PZ, P3, POZ, and P4, as shown in Fig. 1, and defined by the 10–20 system of electrode placement [8]. These 12 channels were referenced to electrically linked mastoids at M1 and M2. The impedance of all electrodes was kept below 20 kΩ. The data were recorded at a sampling rate of 250 Hz with a NeuroScan SynAmps Model 5038 EEG amplifier, which uses a 16-b A/D converter. A serial cable connected the computer used to present the stimulae to the SynAmps EEG amplifier, and was used to signal when a stimulus was presented. The SynAmps was programmed to do analog bandpass filtering from 0.15–30 Hz, and was calibrated with a known voltage before the recording session. Eye blinks were detected by means of a separate channel of data recorded from an electrode placed below the subject's left eye.

Data were recorded from one subject: a 21-year-old, right-handed, male, college student. Five different programs of stimulus were presented using STIM, each displaying a total of ten images. The subject was given written instructions at the beginning of each program of ten images. In general, the instructions were to view an image related to a particular mental task, and to concentrate on the task, after hearing an audible tone, until the next image was presented. The programs of stimulus took the following format, shown in Fig. 2:

- image was presented for 5 s;
- blank (dark) screen was presented for 5 s;
- 1-KHz tone (beep) sounded;
- blank screen continued for another 5 s;
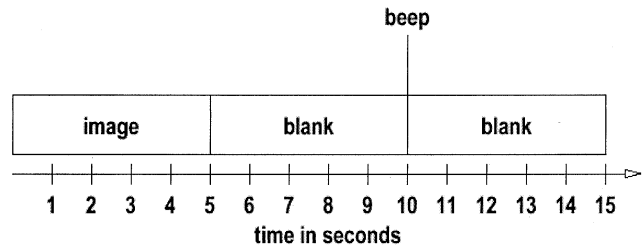- next image in the program was presented.

Fig. 2.   Format of the stimulus programs.

All tasks were performed with the subject's eyes open. The tasks used in this experiment are the same as those chosen by Keirn and Aunon in [3] to invoke hemispheric brainwave asymmetry. The five tasks were the following.

1) *Baseline Task*: The instructions given to the subject preceding the stimulus program were not to perform a specific mental task, but to relax as much as possible, make as few movements as possible, and think of nothing in particular. This task is considered a baseline task for alpha wave production and was used as a control measure of the EEG. The ten images presented were all exactly the same and consisted of a white screen.

2) *Letter Task*: The subject was shown images consisting of a black word on a white background. Each word was indicative of a friend or family member (e.g., "father," "mother," "aunt," "uncle," etc.), and the subject was asked to mentally compose a letter to that person without vocalizing or making any physical movements.

3) *Math Task*: The subject was shown images consisting of nontrivial multiplication problems, such as 89×67, and was asked to solve them without vocalizing or making any physical movements. The problems were designed so that they could not be solved in the time allowed. Although they repeated, the subject did not solve any of them to completion.

4) *Geometric Figure Rotation*: The subject was shown images of three-dimensional (3-D) figures (rendered and shaded), and asked to visualize them rotating about an axis. The figures were all 3-D extrusions of randomly drawn two-dimensional shapes.

5) *Visual Counting*: The subject was shown an image of black arabic numerals on a white background, and asked to visualize similar numerals being written onto a blackboard, one after another, sequentially in ascending order, the previous numeral being erased before the next being written.

Data were recorded for 15 s per image, or 150 s per program. Thus, one recording of all five programs resulted in 750 s of data. The five programs were each recorded twice, giving us a total of 1500 s of data containing 100 beeps, each indicating the start of a 5-s period during which the subject was to be concentrating on a particular brain state.

After having some time to examine the data, it became apparent that P_Z, channel 14 on the QuikCap, had not been recording properly. However, all the other channels were quite clear, and since P_Z is along the center-line, we surmised that
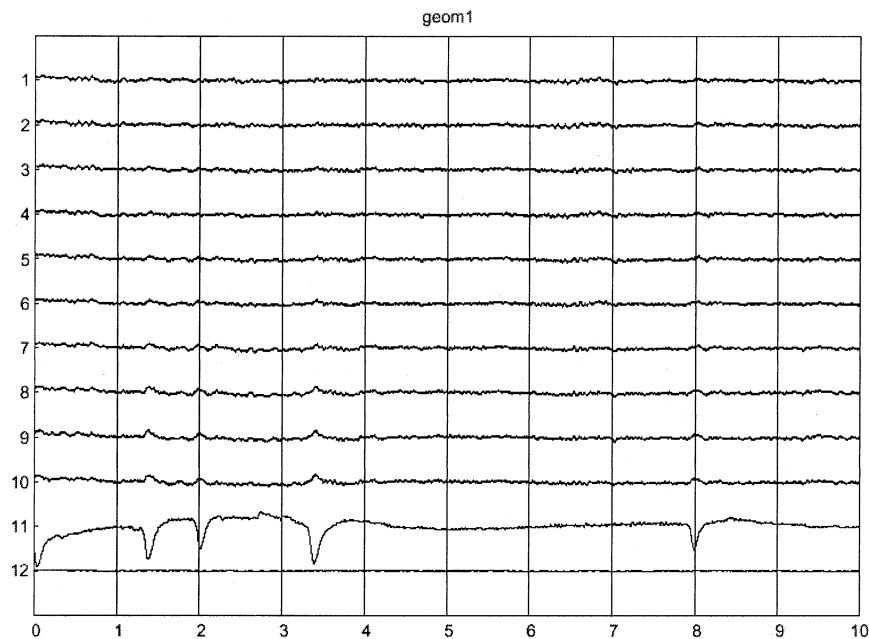
Fig. 3.   First 10 s of the geometry task, before ICA processing.

Fig. 4.   First 10 s of the geometry task, after ICA processing.

discarding it would not adversely affect the usability of the information encoded in the brainwave asymmetry ratios for each state. Discarding the malfunctioning channel left us with 11 channels of scalp data, and a single channel for detecting eye movements.

B. Artifact Removal

Contamination of EEG activity by eye movements, blinks, cardiac signals, and muscle and line noise is a serious problem for EEG interpretation and analysis. One way of dealing with this problem is to simply reject segments of EEG with unac-

ceptable amounts of noise. However, this may result in significant data loss. Fortunately, there are algorithmic alternatives to discarding data. One algorithm in particular stands out from the rest: ICA. To understand what it does and why it serves our purpose, it is of use to gain some context with regard to the type of data being analyzed with EEG.

The signals measured when we record the voltage potentials from a subject's scalp are those that result from the activity of neurons some significant distance away from the electrode that we are using to take the measurement. Each electrode is "hearing" a summation of all the neural activity in the vicinity. The signals differ from one another because they are located in

different geographical areas of the scalp. Neural activity in an area that is close in proximity to one electrode will be "louder" in the recording produced by that electrode than one further from the source. Thus, in an ideal situation, each electrode would detect a unique linear mixture of all the neural activity happening in a subject's brain.

Unfortunately, the ideal linear mixture is augmented by other electrical activity that does not pertain to neurons being fired. Typically, these noise signals are much greater in amplitude than the signals of interest, and have the effect of obliterating a good amount of useful information. Some of the noise signals, such as those resulting from eye blinks and other muscle movements, are infrequent enough such that the segments of data in which they appear can be simply discarded without losing too much. Others, such as cardiac signals and eye movements, are regular enough to make obtaining useful data a cumbersome task. The problem of removing this noise from the interesting signal can be stated as follows: from N unique linear mixtures of an undetermined number of sources, can we somehow separate out N statistically independent mixtures? In other words, can we "unmix" the statistically unrelated noise onto a separate channel from the interesting signals? In fact, it has been known for some time that this is possible.

ICA, proposed by Bell and Sejnowski in [5], is a simple neural algorithm that blindly separates mixtures of independent sources using infomax. In [5], they show that maximizing the joint entropy of the output of a neural processor minimizes the mutual information among the output components. Bell and Sejnowski offer the following two reasons why ICA is suitable for performing blind source separation on EEG data: 1) it is plausible that EEG data recorded at multiple scalp electrodes are linear sums of temporally independent components arising from spatially fixed, distinct, or overlapping brain or extra-brain networks; and 2) spatial smearing of EEG data by volume conduction does not involve significant time delays [10]. Figs. 3 and 4 exemplify the effect of ICA processing on EEG data.

C. Representation of EEG Signals

The key to training a neural network to do a reliable discrimination is finding a suitable representation of the EEG signals. Since the early days of automatic EEG processing in the medical community, representations based on a Fourier transform have been most commonly applied to the problem of discriminating and classifying EEG patterns. This approach builds upon earlier observations that there are some characteristic waveforms that fall primarily within four frequency bands: delta (1–3 Hz), theta (4–7 Hz), alpha (8–13 Hz), and beta (14–20 Hz).

In related work, Anderson *et al.* [2] found that a frequency-band representation yielded the best result of four methods that they tried. Others [13] have had success with similar representations as well. Thus, we decided to use a representation based on the power spectral densities of the sources computed by ICA. With a sample rate of 250 Hz and 12 channels of data, each 5-s window of time during which the subject was to be concentrating on a particular brain state contained 15 000 data points. After computing the ICA sources and discarding the one that was representative of eye and muscle movements, we were left with 11 channels of data. Fig. 5 shows
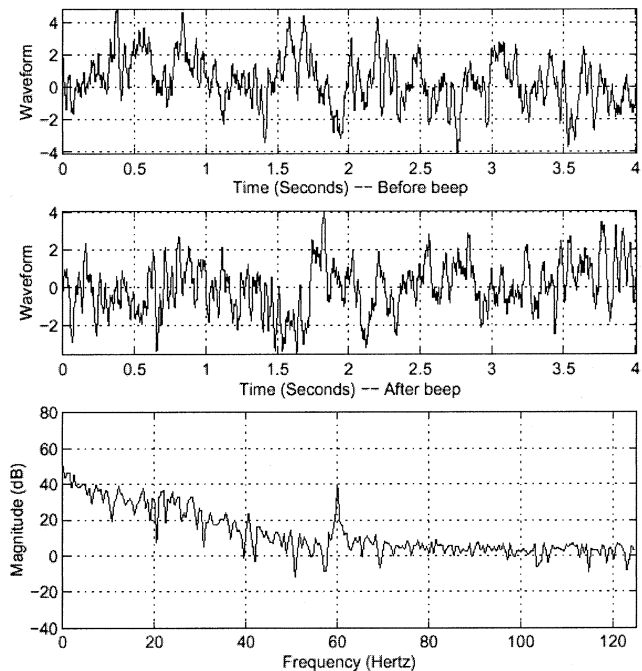

Fig. 5. Third ICA source and its Fourier transform around beep 3 of the counting task.

a portium of the third ICA souce of the counting task before and after the beep, and its Fourier transform.

Inside the period of concentration, we recorded ten windows of 11-channel EEG data, each offset by 50 samples from the one before it. For example, the first of the ten started at the beep, the second started 50 samples after the beep, etc. Since there were 100 windows per session, and two sessions of each mental state, each mental state was represented by 200 feature vectors. Of these features, half were used for training and half for validation.

The length of the window was varied from one half second (125 samples) to one-nineteenth of a second (12 samples). To point out what may not be obvious, windows longer than 50 samples overlapped other windows, while windows shorter did not. For each window, we computed the discrete Fourier transform of each channel, which left us with 11 vectors containing a number of power values equal to the number of samples in the time domain. The power spectral density was computed by taking the dot product of the Fourier transform with its conjugate, and dividing the resulting vector by the number of power values in the Fourier transform. These 11 power spectral density vectors were concatenated together to form a presentation vector.

D. Pattern Classification

Three-layer feedforward artificial neural networks were trained using Matlab with the accelerated back-propagation algorithm (tbpx). The learning rate was dynamically adjusted as the network trained, such that as the mean-squared error (MSE) of the network decreased, the learning rate increased, and when the MSE increased, the learning rate fell back down to a preset minimum. With regard to the number of hidden nodes in each hidden layer, a variety of thoughtfully selected

TABLE I
RESULTS FOR TWO-WAY DIFFERENTIATION

Window Length (seconds, samples)	Best Network (first layer, second layer)	Best Classification	Worst Classification	Training Time (minutes)
(0.5, 125)	(40, 5)	94% geom, mult	77% count, mult	7
(0.25, 62)	(40, 5)	90% base, letter	69% count, letter	6
(0.125, 31)	(100, 10)	82% base, count	69% count, letter	3
(0.05, 13)	(1000, 100)	85% geom, mult	67% count, letter	10

TABLE II
RESULTS FOR THREE-WAY DIFFERENTIATION

Window Length (seconds, samples)	Best Network (first layer, second layer)	Best Classification	Worst Classification	Training Time (minutes)
(0.5, 125)	(100, 10)	86% base, letter, mult	71% count, letter, mult	30
(0.25, 62)	(250, 50)	77% count, geom, letter	63% count, geom, letter	16
(0.125, 31)	(250, 50)	74% geom, letter, mult	56% count, letter, mult	8
(0.05, 13)	(250, 50)	66% geom, letter, mult	51% count, letter, mult	6

configurations were tried. Some of the best performing configurations were 40-5, 50-10, 100-10, 250-50, 1000-100, where for example 40-5 indicates that there were 40 hidden nodes in the first layer, and 5 in the second. All networks were trained on a six-processor Sun Ultra Enterprise 3000 with 1.5 GB of memory.

Networks were trained to differentiate between all pairs of mental states, and all triples of mental states. A one-hot encoding was used to enumerate mental states, with a "correct" classification being one in which the correct output was larger than all other outputs. Half of the total number of feature vectors were used as a validation set to prevent over-fitting, and training was stopped when the MSE of the validation set did not decrease for 200 epochs. Once this occurred, the network weights were reset to their values at the point when the network received the highest "score" upon evaluation of the validation set, where the "score" of the network was the number of validation patterns classified correctly. Each recording donated 100 patterns to the pool, and there were two recordings made of each brain state. Thus, for two-way differentiation, 200 patterns were used for training and 200 different patterns were used for validation. For three-way differentiation, 300 patterns were used for training and 300 different patterns were used for validation.

IV. RESULTS

Table I shows the results of differentiation between two mental tasks. The notation used for the "Best Network" column indicates the number of hidden nodes in each hidden layer. For example, "40-5" indicates that there were 40 nodes in the first hidden layer, and 5 in the second. In the "Best Classification" column, the percentage indicates the number of test patterns classified correctly. In Table I, the percentage is an indication of which mental tasks were used in the differentiation that resulted in that particular percentage. Clearly, some pairs of tasks are more easily differentiated between than others.

Table II shows the results of differentiation between three mental tasks. The classification accuracies indicate that this was a much harder problem for the network to learn. Another indication of the difficulty of the problem is that networks with a greater number of hidden nodes performed better. When only differentiating between two mental tasks, it was sufficient to use fewer hidden nodes: after a certain threshold, classification accuracy did not increase if the number of hidden nodes was increased. After finishing with the results for three-way differentiation, it became clear that it would not be worth the training time to compute them for four-way differentiation.

V. CONCLUSION

Accurate, two-way differentiation can be done using a short (between 0.05 and 0.5 s) window of EEG data. This is probably the most significant result of the experiment, because most applications for control systems have real-time requirements. For example, with a recognition rate of one symbol per second, it would be very difficult to steer a wheelchair or compose a letter on the computer. On the other hand, with a rate of 16 symbols per second it might be possible to accomplish this.

Increasing the number of hidden nodes in the neural network classifier increases the accuracy of classification, although the increase in accuracy is very gradual after a point. Unfortunately, since larger networks take longer to train, there is a threshold at which the return in accuracy does not justify the investment in training time.

ICA is fast and useful for removing artifacts without discarding useful data. This experiment verifies that the ICA algorithm works as a method of removing artifacts from EEG data. Furthermore, the processing time required to run the ICA algorithm was insignificant compared to the time required to train the neural networks using backprop.

ACKNOWLEDGMENT

REFERENCES

[1] C. W. Anderson, S. V. Devulapalli, and E. A. Stolz, "Determining mental state from EEG signals using neural networks," *Sci. Prog.*, vol. 4, no. 3, pp. 171–183, Fall 1995.

[2] ——, "EEG signal classification with different signal representations," in *Neural Networks for Signal Processing*. Piscataway, NJ: IEEE Press, 1995, vol. V, pp. 475–483.

[3] Z. A. Keirn and J. I. Aunon, "A new mode of communication between man and his surroundings," *IEEE Trans. Biomed. Eng.*, vol. 37, pp. 1209–1214, Dec. 1990.

[4] P. F. Harner and T. Sannit, *A Review of the International Ten-Twenty System of Electrode Placement*. Quincy, MA: Grass Instrument, 1974.

[5] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind signal separation and blind deconvolution," *Neur. Computat.*, vol. 7, pp. 1129–1159, 1995.

[6] B. O. Peters, G. Pfurtscheller, and H. Flyvbjerg, "Prompt recognition of brain states by their EEG signals," *Theor. Biosci.*, vol. 116, pp. 290–301, 1997.

[7] M. Pregenzer, G. Pfurtscheller, and D. Flotzinger, "Selection of electrode positions for an EEG-based brain-computer interface," *Biomed. Technik*, vol. 39, pp. 264–269, 1994.

[8] J. R. Wolpaw and D. J. McFarland, "Multichannel EEG-based brain-computer communication," *Electroenc. Clin. Neurophys*, vol. 90, pp. 444–449, 1994.

[9] T.-P. Jung, C. Humphries, T.-W. Lee, S. Makeig, M. J. McKeown, V. Iragui, and T. Sejnowski, "Removing electroencephalographic artifacts: Comparison between ICA and PCA," *Neural Netw. Signal Process.*, vol. VIII, pp. 63–72, 1998.

[10] S. Makeig, A. J. Bell, T.-P. Jung, and T. J. Sejnowski, "Independent component analysis of electroencephalographic data," *Adv. Neur. Inform. Process. Syst.*, vol. 8, pp. 145–151, 1996.

[11] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neur. Computat.*, vol. 7, pp. 1129–1159, 1995.

[12] E. Haselsteiner and G. Pfurtscheller, "Using time-dependent neural networks for EEG classification," *IEEE Trans. Rehab. Eng.*, vol. 8, pp. 457–463, Dec. 2000.

[13] L. Zhukov, D. Weinstein, and C. Johnson, "Independent component analysis for EEG source localization," *IEEE Trans. Eng. Med. Biol.*, pp. 87–96, May/June 2000.

[14] R. Vigario, J. Sarela, V. Jousmaki, M. Hamalainen, and E. Oja, "Independent component approach to the analysis of EEG and MEG recordings," *IEEE Trans. Biomed. Eng.*, vol. 47, pp. 589–593, May 2000.

[15] S. Makeig, S. Enghoff, T.-P. Jung, and T. J. Sejnowski, "A natural basis for efficient brain-actuated control," *IEEE Trans. Rehab. Engin.*, vol. 8, pp. 208–211, June 2000.

[16] C. Guger, H. Ramoser, and G. Pfurtscheller, "Real-time EEG analysis with subject-specific spatial patterns for a brain-computer interface," *IEEE Trans. Rehab. Eng.*, vol. 8, pp. 447–456, Dec. 2000.

Benjamin J. Culpepper received the B.S. degree in computer science from Harvey Mudd College, Claremont, CA, and the M.S. degree in computer science from San Jose State University, San Jose, CA, in 2003.

He is currently a Research Assistant with the Neuroengineering Laboratory, NASA Ames Research Center, Moffet Field, CA. He has been employed at Sun Microsystems, Palo Alto, CA, Network Appliance, Santa Clara, and Stanford University, Stanford, CA. His current interests include bioelectric human-machine interfaces, wireless sensor networks, and parallel processing.

Robert M. Keller (M'70) received the B.S. and M.S. degrees in electrical engineering from Washington University, St. Louis, MO, and the Ph.D. from the University of California, Berkeley.

He was an Assistant Professor at Princeton University, Princeton, NJ, and an Associate Professor (then Professor) at the University of Utah, Salt Lake City. He was Vice-President of R&D at Quintus Computer Systems, Inc., Mountain View, CA. He has held visiting positions at both Stanford University, Stanford, CA, and Lawrence Livermore National Laboratory, Livermore, CA. He is currently the Csilla and Walt Foley Professor of Computer Science at Harvey Mudd College, Claremont, CA. At Harvey Mudd College, he is also Director of the Computer Science Clinic, a unique program in which teams of undergraduates students conduct R&D projects for corporate and research sponsors. His current interests include parallel processing, languages, and soft computing.