Translating Thoughts Into Actions by Finding Patterns in Brainwaves

Charles W. Anderson and Jeshua A. Bratman
Department of Computer Science
Colorado State University, Fort Collins, CO 80523
anderson@cs.colostate.edu bratman@cs.colostate.edu

***Abstract*— Principal Components Analysis (PCA) is often used to project high-dimensional signals to lower dimensional subspaces defined by basis vectors that maximize the variance of the projected signals. Data containing variations of relatively short duration and small magnitude, such as those seen in EEG signals, may not be captured by PCA when applied to time series of long duration. Here, PCA is applied independently to short segments of data and the basis vectors themselves are used as features for classification. In addition, time-embedding the EEG by augmenting each sample with previous samples prior to PCA results in a representation that captures EEG variations in space and time. The resulting features are classified into categories corresponding to which mental task a subject is performing in a brain-computer interface (BCI) paradigm. Approximately 80% of test samples are correctly classified as one of five mental tasks. In addition, an on-line artifact removal method is demonstrated and an inexpensive hardware and software system for BCI research is described.**

I. Introduction

A brain-computer interface (BCI) is a hardware and software system that records electroencephalogram (EEG) signals from human subjects and is used in an on-line fashion to take actions, such as computer cursor movement or the selection of a letter for a typing task, by identifying patterns in EEG corresponding to various kinds of mental imagery performed by the subject. Recent efforts in the brain-computer interface (BCI) field have resulted in exciting demonstrations of the potential for BCI applications. Wolpaw, et al., [15] showed that some subjects can learn two-dimensional control of a computer cursor. Millàn, et al.,[14] demonstrated the control of a wheelchair instrumented with range sensors that prevented collisions. Most of these advances are still based on the discrimination between two or three mental imagery, such as imagined hand or foot movements. The ability to discriminate more mental tasks would enhance the speed with which a disabled person could communicate through a BCI.

Progress in the BCI field is impeded primarily by the difficulty of finding and representing the patterns in EEG that correspond to various mental tasks. To-date, most BCI systems rely on our current understanding of changes in EEG over motor cortex areas during imagined movements. The identification of EEG from other kinds of mental tasks will require new signal representations that capture patterns in EEG over multiple electrodes and over time. Here such a representation is described that is based on singular value decomposition (SVD) of short time windows of multi-channel EEG. Results show that with this representation a simple classifier (linear discriminant analysis) can correctly identify about 80% of the time which of five particular mental tasks subjects are performing, using about four seconds of EEG data. Before describing these results, a method for removing unwanted variations, called artifacts, from the EEG is shown which is also based on SVD.

Another major factor limiting research activity in BCI is the lack of affordable and extendable BCI hardware and software. Here we describe the publicly-available software we have developed and use with an EEG amplifier by NeuroPulse-Systems, LLC. A complete BCI system based on this amplifier with our software running on a laptop computer costs a total of about $6,500.

II. Artifact Removal

A computationally efficient method for removing artifacts from EEG on-line is the maximum signal fraction analysis (SFA). SFA was initially developed as a method to reduce noise in satellite imagery [3] and it was further developed in the context of signal separation [6], [5], [10]. Here SFA is defined as it relates to the problem of separating the components in a sequence of EEG signals. Common artifacts in EEG signals are caused by eye blinks and eye movements and by other muscle movements. The artifacts due to the eyes tend to be slow waves that are present in data recorded from frontal and central electrodes. Artifacts due to muscle movement are higher frequency. Therefore, one approach to artifact removal is to separate the EEG signals into components from slow to fast variability, then eliminate the extreme components from the signal. SFA provides such a spectrum of components, as follows.

Let a time series of multi-channel EEG data be represented by the matrix $X_{n\times d}$ where n is the number of samples and d is the number of electrodes. Each column of X contains the data recorded from a single electrode and each row is the data from all electrodes at one point in time. Characterize the "noise" in the signal by the temporal difference, Y, in X, where $Y = X - X_s$, and the subscript s represents a shift forward in time by one sample. The generalized singular value decomposition (GSVD) of X and Y can be used to extract a spectrum of components ordered

This work was supported by the National Science Foundation Grant No. 0208958.

by the degree to which they represent signal versus noise. The GSVD of X and Y is given by

$$
\begin{aligned}
X &= UCW^T \\
Y &= VDW^T
\end{aligned}
$$

where the matrices U and V are orthogonal, W is an invertible matrix, and C and D are diagonal with $C^2+D^2 = I$. The GSVD of X and Y may be computed with the following Matlab code.

```
Y = X(1:end-1,:)-X(2:end,:);
[U,V,W,C,D] = gsvd(X,Y,0);
extractedSources = U;
```

The matrix U contains the components ordered by increasing amounts of variation. A visual inspection of a graph of these components can reveal which components should be eliminated from the signals. To remove a component the corresponding column of U must be set to zero. This operation can be represented as a right multiplication by a matrix Z where $Z_{ii} = 1$ if a signal is not an artifact, $Z_{ii} = 0$ if signal i is an artifact, and $Z_{ij} = 0$ for all $i \neq j$. Now the artifact filter, F, can be calculated by $F = (CW^T)^{-1}ZCW^T$. The GSVD defines W as an invertible matrix making this definition well defined. This filter is applied to the original EEG as XF and has the effect of removing the contribution of components that were considered artifacts. Application of the filter is not limited to X. It can be applied to new data sets as well. Further details on SFA for artifact removal is provided by Knight, et al., [11], [2].

The use of SFA for artifact removal is demonstrated by applying it to 10 channels of EEG recorded from a subject performing a mental multiplication task. Figure 1a shows 10 channels of EEG recorded for five seconds. Eye movements are visible in the lowest frequencies of FP1 and FP2 and 60 Hz line noise is apparent in all channels. Figure 1b shows the 10 components resulting from SFA. The first component is mostly 60 Hz power-line interference. We can also see that most of the eye movement signal has been isolated in the tenth component. The first, second, and tenth components were selected to be removed and the corresponding filter constructed. Applying the filter to the original EEG results in the filtered EEG shown in Figure 1c. 60 Hz power line noise is no longer apparent. The eye movement has been removed from FP1 and FP2 (and other channels) revealing details not seen in the original EEG.

III. Classification of EEG Represented by Short-Time PCA

Principal component analysis (PCA) is commonly used to project data samples to a lower-dimensional subspace that maximizes the variance of the projected data. For many data sets, PCA is also used to isolate the information in the data into meaningful components, such as “eigenfaces” [13] and “eigenlips” [9] in applications involving analysis of face images.

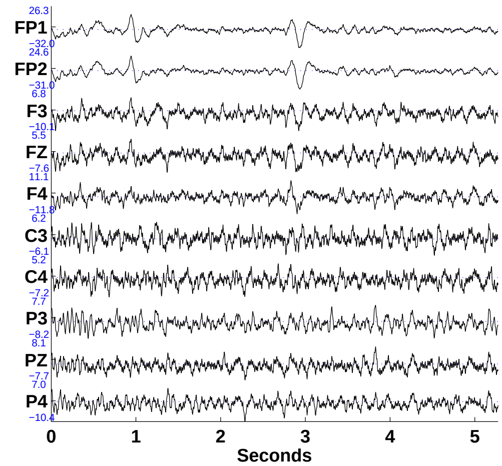

a. Original EEG

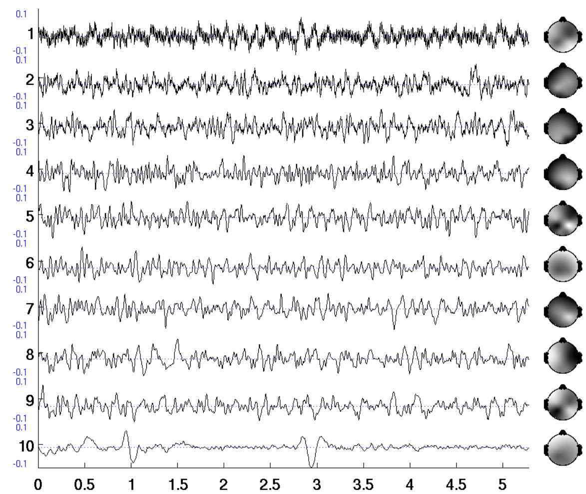

b. SFA Components

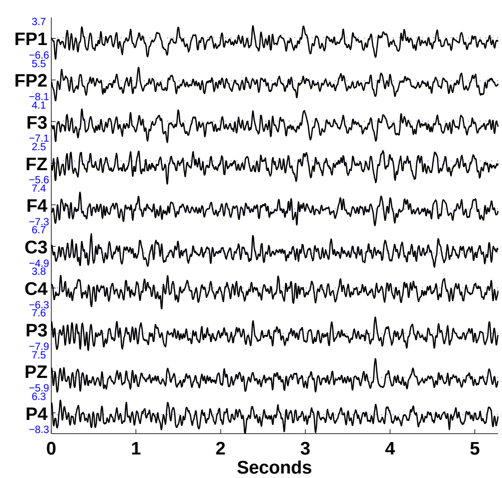

c. Filtered EEG

Fig. 1. Creation and use of SFA for removing artifacts from EEG 10 channels of EEG. a) Five seconds of 10 channel EEG sampled at 256 Hz; b) Components (columns of U) separated by SFA, ordered from smallest to largest signal-to-noise ratio; c) EEG after removing SFA Components 1, 2, and 10.

For classification problems, PCA is usually applied to a collection of samples from all classes with the hope that the projection of new samples onto the PCA basis form components whose amplitudes are related to the class. This approach may fail to capture variations that appear in the data over short time intervals. Such variations contribute little to the overall variance of the data, but may be critical in classifying samples into the correct classes.

Features of short duration can be captured by applying PCA to short windows in time of the data. This results in multiple bases, one for each window. To project data samples using these multiple bases, they must somehow be combined into a single basis. An alternative approach is used here. Rather than projecting the data to form features on which classification is performed, the bases themselves are taken as the features. Our hypothesis is that the directions of significant variation within each window will capture the information needed to correctly classify the data in the window. We refer to this method as short-time PCA, or STPCA. The basis vectors are calculated by singular value decomposition (SVD) of the matrix of samples in each window.

The EEG samples are augmented by samples delayed in time, forming a time-embedded representation described in the next section and in [2], [8]. With this modification, PCA becomes a tool for simultaneously analyzing spatial and temporal aspects of the data. A related approach using common spatial patterns was recently described in [12].

Let x_i be a column vector of EEG voltages for d electrodes at sample time i. Consecutive samples are combined to form the $l+1$ dimensional time embedding $\hat{x}_i = (x_{i+l}^T, x_{i+l-1}^T, \ldots, x_i^T)^T$, for $i = 0, \ldots, n-l$. Windows of w consecutive time-embedded samples that overlap by p samples are collected into matrices $W_j = (\hat{x}_{1+(w-p)(j-1)}, \ldots, \hat{x}_{(w-p)j+p})$, for $j = 1, \ldots, \lfloor \frac{n-p-l}{w-p} \rfloor$. Each W_j is $d(l+1) \times w$. The SVD, $W_j = U_j S_j V_j^T$, of each window, W_j, is performed to obtain the left singular vectors in U_j. To remove the variation in sign of basis vectors over multiple windows that results from SVD, columns of U_j for which the first component is negative are multiplied by -1. Concatenating all columns of U_j results in a $(d(l+1))^2$ dimensional vector representing the j^{th} window of multichannel, time-embedded EEG data.

A. Linear Discriminant Analysis

Linear Discriminant Analysis (LDA) is a simple probabilistic approach to classification in which each class is assumed to follow a normal distribution [4]. The parameters for the distribution of each class are estimated, and with Bayes Rule are combined to form linear discriminant functions. For Class k with sample mean μ_k and covariance matrix Σ averaged over all classes, the discriminant function is

$$\delta_k(x) = x^T \Sigma^{-1} \mu_k - \frac{1}{2} \mu_k^T \Sigma^{-1} \mu_k.$$

Defining weights $w_k = \Sigma^{-1}\mu_k$ and bias $b_k = -\frac{1}{2}\mu_k^T \Sigma^{-1} \mu_k$, each discriminant function simplifies to $\delta_k(x) = x^T w_k + b_k$.

B. Results

EEG data used here was provided by an earlier study [7]. EEG signals were recorded from subjects performing the following five mental tasks: 1) resting task, in which subjects were asked to relax and think of nothing in particular; 2) mental letter writing, in which subjects were instructed to mentally compose a letter to a friend without vocalizing; 3) mental multiplication of two multi-digit numbers, such as 49 times 78; 4) visual counting, in which subjects were asked to imagine a blackboard and to visualize numbers being written on the board sequentially; and 5) visual rotation of a three dimensional block figures. For each trial, EEG was recorded from six electrodes ($d = 6$) at positions $(C_3, C_4, P_3, P_4, O_1, O_2)$ for 10 seconds sampled at 250 Hz. Each task was repeated five times (for a total of five trials per task). The order in which tasks were performed was randomized, and subjects did not practice the tasks beforehand. Another five trials of each task were recorded during a second session on a following day.

Window size and overlap were $w = 50$ and $p = 40$. The time embedding involved three samples ($l = 2$). Data was partitioned into training and testing sets as follows. First, all five trials from the first day comprised the training set and all five trials from the second day comprised the testing set. The LDA classifier was computed from the training set samples and the trained classifier was tested by applying it to the testing set samples. This evaluates how well a trained classifier generalizes to data recorded on another day. Then, training one trial at a time from the second day was moved from the testing set to the training set.

It was found that the classification accuracy on testing data is increased by combining the classification result over consecutive windows of EEG by selecting the most commonly predicted class over consecutive windows. Results for two subjects are shown in Figure 2 as fraction of test samples correctly classified versus the number of consecutive windows whose predicted classes were combined. There is a clear trend in increasing classification accuracy as more trials from the second day are included in the training set.

The amount of EEG in seconds, t, sampled at f Hz, that corresponds to combining the classes of m consecutive windows of size w with overlap p is $t = \frac{m(w-p)+p}{f}$. So, in this experiment, $m = 1$ consecutive window is equivalent to approximately $t = 0.2$ seconds, $m = 40$ is $t = 1.8$ seconds and $m = 100$ is $t = 4.2$ seconds. For Subject 1, the correct classification rate for test trials 9 and 10 goes from about 62% for 0.2 seconds of data to about 78% for 1.8 seconds. The results for Subject 2 reach about 68% correct for 4.2 seconds of data. Recall that this is the rate of predicting the correct task out of five, so a random classifier would result in only 20% correct.

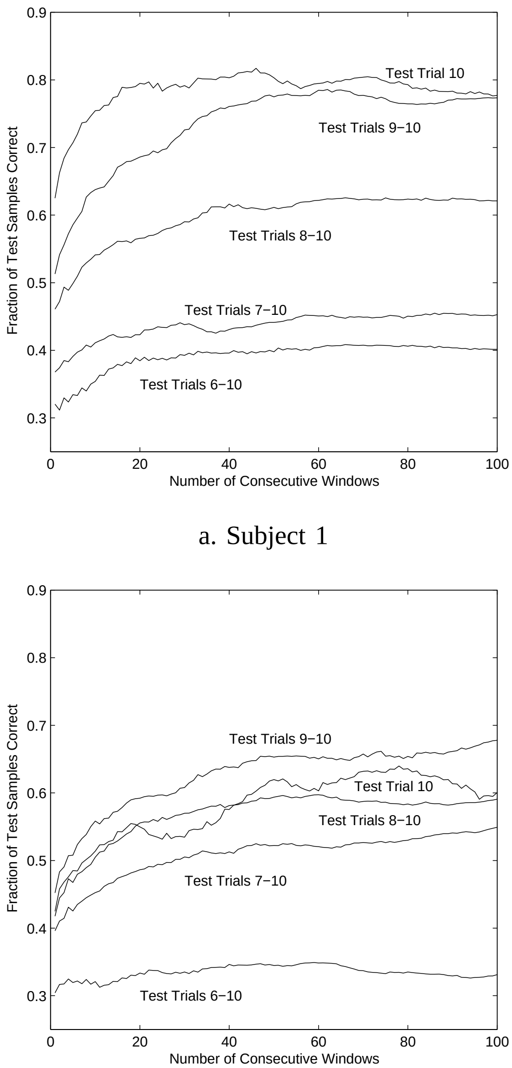

a. Subject 1

b. Subject 2

Fig. 2. Fraction of test samples correctly classified versus number of consecutive windows whose predicted classes are combined. Different curves correspond to different partitions of trials into training and testing sets, indicated by the range of trials indices in the testing set. Trials 1 through 5 were recorded on the first day and Trials 6 though 10 were recorded on a second day.

		Predicted				
		Task 1	Task 2	Task 3	Task 4	Task 5
Actual	Task 1	100.0	0.0	0.0	0.0	0.0
	Task 2	3.4	11.0	3.8	34.2	47.6
	Task 3	0.0	0.0	100.0	0.0	0.0
	Task 4	0.0	0.0	1.6	96.0	2.4
	Task 5	2.0	0.0	15.3	0.0	82.7

a. Subject 1

		Predicted				
		Task 1	Task 2	Task 3	Task 4	Task 5
Actual	Task 1	90.9	0.0	4.5	0.0	8.7
	Task 2	39.4	35.8	0.0	0.0	24.8
	Task 3	21.2	31.4	45.6	1.8	0.0
	Task 4	0.0	2.4	0.0	72.6	25.0
	Task 5	0.0	0.0	0.0	0.0	100.0

b. Subject 2

TABLE I

CONFUSION MATRICES FOR SUBJECTS 1 AND 2 SHOWING THE PERCENT OF TEST SAMPLES CORRECTLY CLASSIFIED.

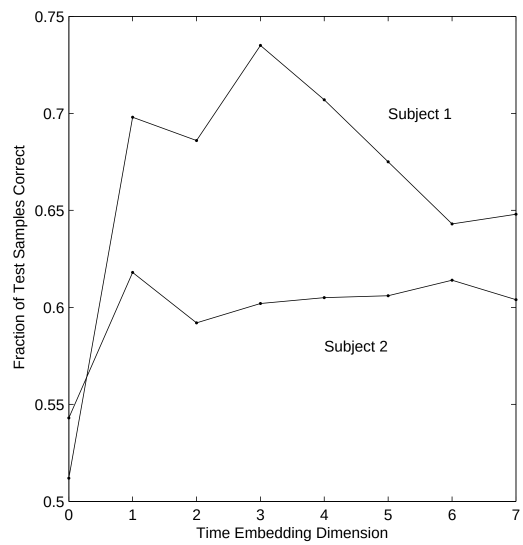

Fig. 3. Fraction of samples from test trials 9 and 10 that are correctly classified versus the time embedding dimension, l.

Details of the classification performance on test trials is revealed by confusion matrices showing the percent of test samples correctly classified for each pair of actual and predicted classes. Table I shows the confusion matrices for each Subject, using $m = 100$ consecutive windows. Clearly some tasks are better classified than others, but depend on the Subject. Data from Subject 1 when performing the resting, multiplication, and visual rotation tasks are classified with the most accuracy, while data from the letter writing task is classified as either visual rotation or mental counting. For Subject 2, resting and mental counting are well classified, but letter writing data is confused with resting and counting, and multiplication data is confused with resting and letter writing.

The above results are for the a time embedding dimension of $l = 2$. The sensitivity of the results to this dimension is determined by repeating the above classification experiments using a variety of values. Figure 3 shows the results. The number of consecutive windows is $m = 20$. A higher embedding dimension of 3 or 4 produces better results for Subject 1, but a dimension of 1 is best for Subject 2.

For six-channel EEG with time embedding dimension $l = 2$, the representation, U_i of the i^{th} window of data consists of $(6(l+1))^2 = 324$ components. The variations in the weights of the LDA discriminant functions over the five

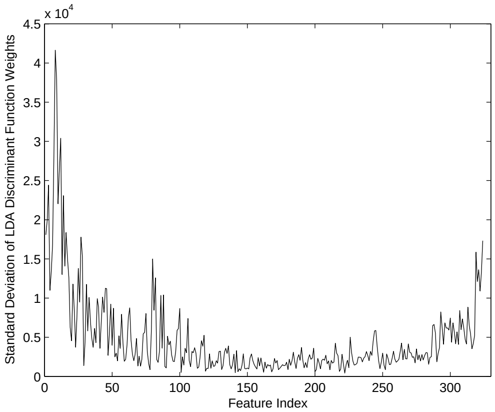

Fig. 4. Standard deviation of LDA weights over the five linear discriminant functions for the five tasks. Components of the first and the last basis vectors have the most variance and are thus the most significant.

classes indicate which of 324 features are most significant for the classification results. Figure 4 shows the standard deviation of the weights for an LDA classifier trained on data from Subject 1. Results are similar for Subject 2. Recall that U_i is composed of the basis vectors ordered from ones with largest singular values to those with smallest singular values. Figure 4 shows that the components in the first basis vector, the one that captures the most variance in the data, have LDA weights that vary the most over the LDA discriminant functions, so play a significant role in distinguishing the tasks. Somewhat surprising is that components in the last basis vectors, ones that capture the least variance in the data, are also significant. These low-variance directions may be capturing small variations in the data that relate strongly to the mental task being performed.

IV. EEG Acquisition and Custom BCI Software

The expense of clinical EEG acquisition systems is out of reach of small research groups. To facilitate the entry of new researchers into the BCI field, we have studied a combination of an affordable EEG amplifier with custom EEG acquisition and BCI software. The EEG amplifier is the 24-channel amplifier, called MS-24R, from NeuroPulse-Systems, LLC, with a cost of $5,000. Additional parts are Electro-cap electrode caps with supplies, an impedance meter, and a SCSI adaptor for a notebook computer, with a total cost of $1,500. So, without the purchase of a notebook computer, the total cost is about $6,500. This list of hardware, and the software described below, is available on the net [1].

Our software has been designed to perform on-line EEG recording, filtering, and classification. The objective of the software design is to provide a convenient method of connecting various pieces of BCI experimentation such as interfacing with an EEG amplifier, selecting channels, filtering, extracting features, classifying, and providing input interfaces. The system was written in C++ and runs on the Linux operating system.

The software is designed to be modular in terms of all the pieces that might need to be added for future experimentation. Filters, feature extractors, and classifiers are all built as external libraries which can be plugged into the main system. This makes experimentation with new algorithms convenient.

In addition to the filters, feature extractors, and classifiers, the user interfaces are designed to be modifiable and easy to implement. These user interfaces are the method of converting classified mental tasks from a user wearing the EEG cap into useful computer input. The interfaces currently designed are a simple pie menu, where each slice represents a class, and a keyboard using a similar pie menu. Figure 5 shows this pie menu in action. In the top example, there are three mental tasks, each of which is displayed on one slice of the pie. As data is sampled from a user wearing an EEG cap and classified as one of these three classes, a bar will grow from the center of the menu towards one of the slices. When a bar reaches a pie slice, all bars are reset and the process starts again.

A more complicated and useful interface is the keyboard pie menu shown in the bottom half of Figure 5. This pie menu is similar to the one previously described, but instead of just resetting when a bar reaches a pie slice, the contents of the selected slice expand to cover the rest of the pie. When a single letter or command is finally selected, it is sent to the computer as keystroke and the whole menu is reset to its initial state.

Initial tests of the entire system involved the control of a wirelessly connected robot using commands from the BCI. A subject wore an EEG cap with four bipolar electrodes used. Three tasks were trained. To simplify the classification problem in this preliminary experiment, three tasks were used that were primarily muscle movements instead of pure mental tasks. They included blinks, jaw clenches, and right hand movement. Each task was trained for three sequences of five seconds each. After training, an LDA classifier was trained to classify time-embedded data. After training, the system was attached wirelessly to a simple robot, and commands were associated with each task. To turn the robot left, the subject would blink, to turn the robot right, the subject would move their right hand, to move the robot forward, the subject would clench their jaw. Using these tasks, the subject effectively steered the robot around the room with these three controls.

V. Conclusion

Experiments showed that EEG representations based on short-time PCA can be classified by simple linear discriminant analysis (LDA) with an accuracy of about 80% correct classification of the correct mental task out of five. This data was obtained from two subjects; tests on additional subjects are warranted to investigate the generality of this result.

Fig. 5. Two user interfaces being tested in the publicly-available BCI software [1]. Top: An interface for the three mental tasks, count backwards, imagine left hand or right hand imagery. Bottom: An interface for three mental tasks to type by choosing a section of the circle in a hierarchical menu. The selected subset of letters is distributed among the three rings at each level, until a single letter is selected.

Analysis of the classifiers' weights revealed that short-time PCA basis vectors late in the sequence play significant roles, suggesting that the low-variance activity represented by these vectors is strongly related to the mental task. This hypothesis warrants further study.

Information gleaned from analyses like those summarized in Figures 4 can be used to select subsets of features to greatly reduce the dimensionality of the data and possibly improve the generalization performance of the classifiers. Extending this analysis to consider the time course of significant electrodes and basis vector directions could lead to hypotheses of the underlying cognitive activity.

VI. Acknowledgments

The authors thank Joseph Anderson for assisting with EEG recording.

References

[1] CEBL: CSU Eeg Brain-computer interface Lab. `http://www.cs.colostate.edu/eeg/eegSystem.html`.

[2] C. W. Anderson, J. N. Knight, T. O'Connor, M. J. Kirby, and A. Sokolov. Geometric subspace methods and time-delay embedding for EEG artifact removal and classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(2):142–146, 2006.

[3] A. A. Green, M. Berman, P. Switzer, and M. D. Craig. A transformation for ordering multispectral data in terms of image quality with implications for noise removal. *IEEE Transactions on Geoscience and Remote Sensing*, 26(1):65–74, January 1988.

[4] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, 2001.

[5] D. R. Hundley, M. J. Kirby, and M. Anderle. Blind source separation using the maximum signal fraction approach. *Signal Processing*, 82(10):1505–1508, 2002.

[6] Douglas R. Hundley, Michael J. Kirby, and Markus G. Anderle. A solution procedure for blind signal separation using the maximum noise fraction approach: Algorithms and examples. In *Proceedings of the 3rd International Conference on Independent Component Analysis and Blind Signal Separation*, 2001.

[7] 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.

[8] M. Kirby and C. Anderson. Geometric analysis for the characterization of nonstationary time-series. In E. Kaplan, J. Marsden, and K. R. Katepalli Sreenivasan, editors, *Springer Applied Mathematical Sciences Series Celebratory Volume for the Occasion of the 70th Birthday of Larry Sirovich*, chapter 8, pages 263–292. Springer-Verlag, 2003.

[9] M. Kirby, F. Weisser, and G. Dangelmayr. A model problem in the representation of digital image sequences. *Pattern Recognition*, 26(1):63–73, 1993.

[10] Michael Kirby. *Geometric Data Analysis: An Empirical Approach to Dimensionality Reduction and the Study of Patterns*. John Wiley & Sons, Inc., 2001.

[11] J. N. Knight. Signal fraction analysis and artifact removal in EEG. Master's thesis, Department of Computer Science, Colorado State University, 2003.

[12] S. Lemm, B. Blankertz, G. Curio, and K.-R. Müller. Spatio-spectral filters for improving the classification of single trial EEG. *IEEE Transactions on Biomedical Engineering*, 52(9):1541–1548, Sep 2005.

[13] M. Turk and A. Pentland. Face recognition using eigenfaces. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.

[14] Gerolf Vanacker, Jos del R. Millán, Eileen Lew, Pierre W. Ferrez, Ferran Galan Moles, Johan Philips, Hendrik Van Brussel, and Marnix Nuttin. Context-based filtering for assisted brain-actuated wheelchair driving. *Computational Intelligence and Neuroscience*, 2007.

[15] J. Wolpaw and D. McFarland. Control of a two-dimensional movement signal by a non-invasive brain-computer interface in humans. In *Proceedings of the National Academy of Sciences*, volume 101 (51), pages 17849–17854, December 21 2004.