

Linear Spatial Integration for Single-Trial Detection in Encephalography

Lucas Parra,* Chris Alvino,* Akaysha Tang,† Barak Pearlmutter,† Nick Yeung,‡
Allen Osman,§ and Paul Sajda||¹

*Vision Technologies Laboratory, Sarnoff Corporation, Princeton, New Jersey 08540; †Department of Psychology and Department of Computer Science, University of New Mexico, Albuquerque, New Mexico 87131; ‡Department of Psychology, Princeton University, Princeton, New Jersey 08544; §Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104; and ||Department of Biomedical Engineering, Columbia University, New York, New York 10027

Received December 31, 2001

Conventional analysis of electroencephalography (EEG) and magnetoencephalography (MEG) often relies on averaging over multiple trials to extract statistically relevant differences between two or more experimental conditions. In this article we demonstrate single-trial detection by linearly integrating information over multiple spatially distributed sensors within a predefined time window. We report an average, single-trial discrimination performance of $A_z \approx 0.80$ and fraction correct between 0.70 and 0.80, across three distinct encephalographic data sets. We restrict our approach to linear integration, as it allows the computation of a spatial distribution of the discriminating component activity. In the present set of experiments the resulting component activity distributions are shown to correspond to the functional neuroanatomy consistent with the task (e.g., contralateral sensory-motor cortex and anterior cingulate). Our work demonstrates how a purely data-driven method for learning an optimal spatial weighting of encephalographic activity can be validated against the functional neuroanatomy. © 2002 Elsevier Science (USA)

INTRODUCTION

Trial averaging is often used in brain imaging to mitigate low signal-to-interference (SIR) ratios. For example, it is the basis for analysis of event-related potentials (ERPs) (Coles and Rugg, 1995). However, for some encephalographic applications, such as seizure prediction, trial averaging is problematic. One application where the problem of single-trial averaging is immediately apparent is the brain-computer interface (BCI), i.e., interpreting brain activity for real-time communication. In the simplest case, where one wishes

to communicate a binary decision, averaging corresponds to asking the same question over multiple trials and averaging the subject's binary responses. In order to obtain high-bandwidth communication, it is desirable to do as little averaging over time or across trials as possible.

More generally, single-trial analysis of brain activity is important in order to uncover the origin of response variability, for instance, in analysis of error-related negativity (ERN). The ERN is a negative deflection in the electroencephalogram (EEG) following perceived incorrect responses (Gehring *et al.*, 1993; Falkenstein *et al.*, 2000) or expected losses (Gehring and Willoughby, 2002) in a forced-choice task. Single-trial detection of the ERN has been proposed as a means of correcting communication errors in a BCI system (Schalk *et al.*, 2000). With the ability to analyze the precise timing and amplitude of the ERN, on individual trials, one can begin to study parameters that cannot be controlled across trial, such as reaction time or error perception. Such an approach opens up new possibilities for studying the behavioral relevance and neurological origin of the ERN.

With the large number of sensors in high-density EEG and magnetoencephalography (MEG), an alternative approach to trial averaging is to integrate information over space rather than across trials. A number of methods along these lines have been proposed. Blind source separation analyzes the multivariate statistics of the sensor data to identify spatial linear combinations that are statistically independent over time (Makeig *et al.*, 1996; Vigario *et al.*, 2000; Tang *et al.*, 2002). Separating independent signals and removing noise sources and artifacts increases SIR. However, blind source separation does not exploit the timing information of external events that is often available. In most current experimental paradigms subjects are prompted with external stimuli to which they are asked to respond. The timing of the stimuli, as well as

¹ To whom correspondence and reprint requests should be addressed at Department of Biomedical Engineering, Columbia University, 351 Engineering Terrace: MC8904, New York, NY 10027. Fax: (212) 854-8725. E-mail: ps629@columbia.edu.

the timing of overt responses, is therefore available and should be exploited by the analysis method.

In the context of a BCI system, many methods have applied linear and nonlinear classification to a set of features extracted from the EEG. For example, adaptive autoregressive models have been used to extract features across a limited number of electrodes, with features combined using either linear or nonlinear classifiers to identify the activity from the time course of individual sensors (Pfurtscheller and Neuper, 2001). Others have proposed to combine sensors in space by computing maximum and minimum eigenvalues of the sensor covariance matrices. The eigenvalues, which capture the power variations of synchronization and desynchronization, are then combined nonlinearly to obtain binary classification (Ramoser *et al.*, 2000). Spatial filtering has also been used to improve the signal-to-noise ratio (SNR) of oscillatory activity. However, there has been no systematic effort to choose optimal spatial filters. In the context of the ERN, Gehring *et al.* (1993) use linear discrimination to identify characteristic time courses in individual electrodes, but do not exploit spatial information. Although many of these aforementioned methods obtain promising performance in terms of classifying covert (purely mental) processes, their neurological interpretation remains obscured.

In this article we report on using conventional linear discrimination to compute the optimal spatial integration of a large array of sensors. We exploit timing information by discriminating and averaging within a short time window relative to a given external event. We restrict ourselves to a linear integration, as it permits the computation of spatial distributions of the discriminating component² activity, which in turn can be compared to functional neuroanatomy.

We demonstrate the utility of the proposed method for three distinct data sets acquired via high-spatial density encephalography:

- Predicting explicit (overt) motor response using MEG (122 sensors).
- Classifying imagined (covert) motor activity using EEG (59 sensors).
- Detecting decision errors for a binary discrimination task using EEG (64 sensors).

MATERIALS AND METHODS

Linear Discrimination

Denoting $\mathbf{x}(t)$ as the M sensor values sampled at time instance t , we compute the spatial weighting coefficients \mathbf{v} such that

$$y(t) = \mathbf{v}^T \mathbf{x}(t) \quad (1)$$

is maximally discriminating between the times t , corresponding to two different experimental conditions. For example, in the prediction of explicit motor response experiments (described below) the times correspond to a number of samples prior to an explicit button push. The samples corresponding to a left button push are to be discriminated from samples of a right button push. For each of N trials there are T samples, totaling NT training examples. We use conventional logistic regression (Duda *et al.*, 2001) to find \mathbf{v} . After finding the optimal \mathbf{v} we average over the T -dependent samples of the k th trial to obtain a more robust result, $\bar{y}_k = (1/T) \sum_{t \in T_k} y(t)$, where T_k denotes the set of sample times corresponding to trial k . We evaluate performance using receiver operating characteristic (ROC) analysis (Swets, 1979) on the single-trial short-time averaged discrimination activities (\bar{y}_k). For visualization purposes, it is also useful to compute the trial averaged discrimination activities,

$$\bar{y}_e(t) = \frac{1}{N_{ke \in N_e}} \sum_{k \in N_e} y_k(t), \quad (2)$$

where N_e denotes the set of samples for event e (e.g., left or right button push) with time measured relative to some common reference across trials. The separation of the means together with their corresponding variances provides an indication of whether single-trial discrimination is plausible within the analysis window.

Localization of Discriminating Components

In order to provide a functional neuroanatomical interpretation of the resultant spatial weighting, we treat $y(t)$ as a component which is maximally discriminating given the linear model and task. A simple way of visualizing the origin of a component's activity is to display the coupling coefficients of the component with the sensors. The strength of the coupling roughly indicates the proximity and orientation of the component relative to the sensor. The coupling \mathbf{a} is defined as the coefficients that multiply the component $y(t)$ to give its additive contribution $\mathbf{x}_y(t)$ to the sensor readings, $\mathbf{x}_y(t) = \mathbf{a}y(t)$. However, $\mathbf{x}_y(t)$ is not observable in isolation and instead we observe $\mathbf{x}(t) = \mathbf{x}_y(t) + \mathbf{x}_{y'}(t)$, where $\mathbf{x}_{y'}(t)$ represents the activity that is not due to the discriminating component. If the contributions of the other components, $\mathbf{x}_{y'}(t)$, are uncorrelated with $y(t)$ we obtain the coupling coefficients by the least-squares solution (Haykin, 1996). Arranging the samples $\mathbf{x}(t)$ for different t as columns in the matrix \mathbf{X} and $y(t)$ as a column vector \mathbf{y} , the solution is given by

² Note that we use the term "component" instead of "source" to avoid confusion with an implied physiological source.

$$\mathbf{a} = \frac{\mathbf{X}\mathbf{y}}{\mathbf{y}^T\mathbf{y}}. \quad (3)$$

In the Appendix we present a derivation of \mathbf{a} and describe some of the underlying assumptions in its interpretation. In general, other components are not guaranteed to be uncorrelated with the discriminating component. Therefore \mathbf{a} represents the coupling of all component activities that are correlated with the discriminating component $y(t)$. We refer to \mathbf{a} as a “sensor projection,” as it measures the activity in the sensors that correlate with a given component. Our approach relies on the linearity of $y(t)$ and the fact that different components in EEG and MEG add linearly (Baillet *et al.*, 2001).

Datasets for Analysis

Predicting explicit (overt) motor response using MEG. This data set was provided by A.T. and B.P. Four subjects performed a visual-motor integration task. A “trump” experiment was defined whereby subjects were simultaneously presented with two visual stimuli on a CRT, one of which is the target and “trumps” (beats-out) the other. Subjects were instructed to push a left-hand or right-hand button, depending on which side the target (trump stimulus) was present. The subject was to discover the target by trial-and-error using auditory feedback. Each trial began with visual stimulus onset, followed by button push, followed by auditory feedback, indicating if the subject responded correctly. The interval between the motor-response and the next stimulus presentation was 3.0 ± 0.5 s. Each subject performed 90 trials, which took approximately 10 min. MEG data were recorded using 122 sensors at a sampling rate of 300 Hz and high-pass filtered to remove DC drifts.

Classifying imagined (covert) motor activity using EEG. This data set was provided by A.O. Nine subjects performed a visual stimulus-driven finger (L/R)

tapping task. Subjects were asked to synchronize an explicit or imagined tap by the left, right, or both index fingers to the presentation of a brief temporally predictable signal. Subjects were trained until their explicit taps occurred consistently within 100 ms of the synchronization signal. Subjects were presented visual stimuli indicating with which index finger to tap and if it should be an explicit or imagined tap; 1.25 s after the last instruction symbol a fixation point was replaced for 50 ms by the letter “X.” This letter served as a signal to which the instructed tap (whether overt or imagined) was to be synchronized. Each trial lasted for 6 s. After training, each subject received 10 blocks of trials. Each 72-trial block consisted of nine replications of the 8 trial types (Explicit vs Imagined \times Left vs Right vs Both vs No Tap) presented in a random order. Trials with noise due to eye blinks were not considered in the EEG analysis. The electromyogram (EMG) was recorded to detect muscle activity during imagined movements. The 59 EEG channels were sampled at 100 Hz and high-pass filtered to remove DC components.

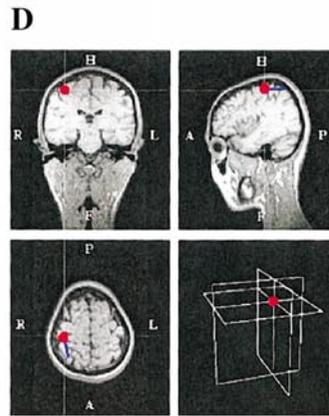
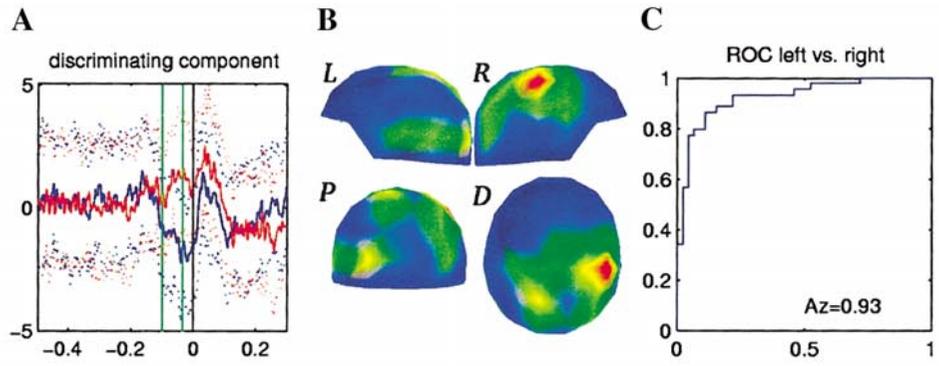
Detecting decision errors for a binary discrimination task using EEG. This data set was provided by N.Y. Seven subjects performed a visual target detection among distractors task. On each trial, subjects were presented with a stimulus for 100 ms. There were four possible stimuli, each consisting of a row of five arrows. Subjects were told to respond by pressing a key on the side indicated by the center arrow. They were to ignore the four flanking arrows. On half of the trials, the flanking arrows pointed in the same direction as the target (e.g., <<<<<); on the other half the flankers pointed in the opposite direction (e.g., <<>><<). Subjects were slower and made many more errors in the latter case. Following their response, there was an intertrial interval of 1.5 s, after which a new stimulus was presented. Subjects performed 12 blocks of 68 trials each. The 100-ms interval prior to the response was used as the baseline period (separately for each trial

TABLE 1

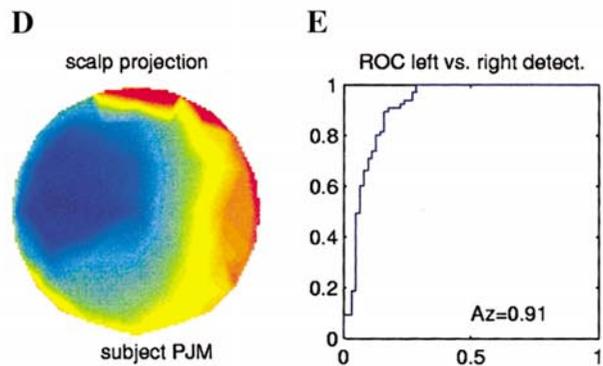
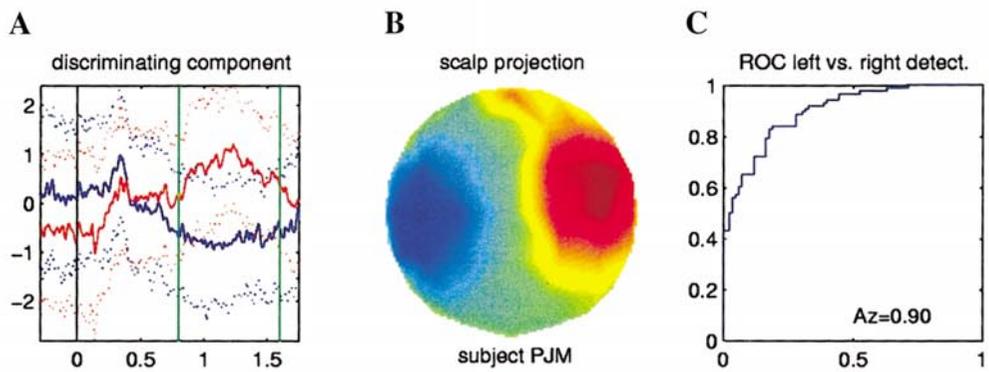
Performance of the Linear Spatial Integration Method for the Three Data Sets

	ROC area (A_z) Mean \pm SD	Fraction correct Mean \pm SD	N	N_e	Sensors	Detection time window
Explicit L/R button push prediction	0.82 \pm 0.06	0.79 \pm 0.09	4	45/45	122 MEG	100 to 33 ms prior to button push
Imagined L/R finger tap discrimination	0.77 \pm 0.10	0.71 \pm 0.08	9	90/90	59 EEG	400 ms before to 400 ms after synchronization
Response error/correct discrimination	0.79 \pm 0.05	0.73 \pm 0.05	7	40–80/300	64 EEG	0 to 100 ms after response

Note. Mean and standard deviation (SD) are reported across N subjects. N_e is the number of trials used to determine the best linear classifier (No. positive/No. negative trials). The time window used for detection is also specified.



1



2

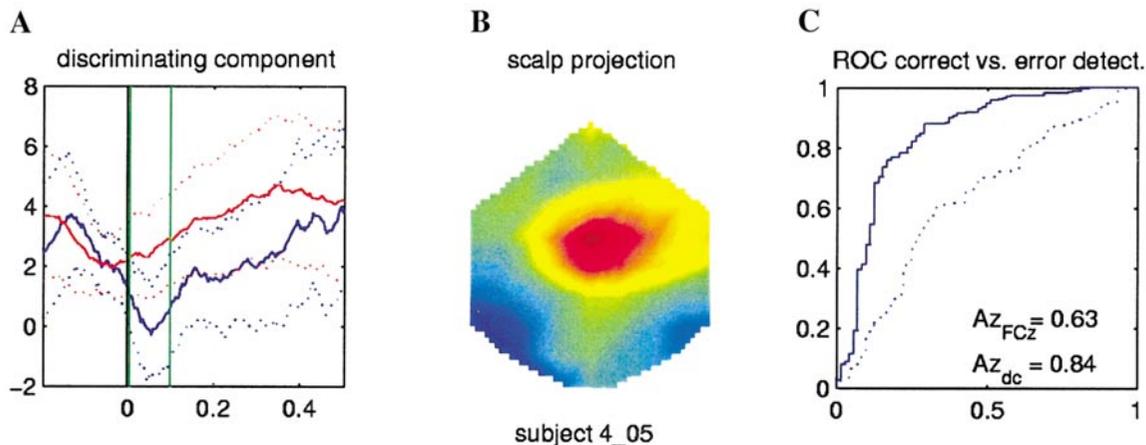


FIG. 3. Detection of decision errors with EEG. (A) Trial averages $\bar{y}_c(t)$ (solid curves) and standard deviations (dotted curves) of discriminating component for correct and error trials. The negative deflection after a button push response at $t = 0$ s is the ERN. Green vertical lines indicate the discrimination window. (B) Dorsal view of sensor projections \mathbf{a} . (C) ROC curve for error vs correct trials. Solid curve corresponds to discrimination using Eq. (1), and dotted line to discrimination with center electrode (FCz).

and electrode). The sampling rate was 250 Hz. Following the baseline period, trials were manually edited to remove those with blinks, large eye movements, instrument artifacts, and amplifier saturation.

RESULTS AND DISCUSSION

Single-trial discrimination results are shown for the three different data sets and include trial averaged discriminating component activity $\bar{y}_c(t)$, sensor projections \mathbf{a} , and detection/prediction performance using single-trial, short-time averaged \bar{y}_k . Performance is reported using ROC analysis computed with a leave-one-out training and testing procedure (Duda *et al.*, 2001). ROC analysis is a reasonable method for quantifying performance for these three data sets, since it enables one to incorporate an independent cost for false positives and false negatives. For example, in an error correction application using the ERN, it is important to detect error events with high confidence. The desired operating point of such a detector is therefore at a low false-positive rate (high specificity). In contrast, an application which looks to exploit motor imagery for communicating a binary decision is best assessed at the operating point where sensitivity equals specificity; i.e., the error rates for the two possible outcomes are

equal. A metric that quantifies the overall performance of a detector for arbitrary operating points is the area under the ROC curve (A_z). In the following sections we report A_z , as well as the fraction of correct classifications, for all three tasks. A summary of the results for the three data sets is given in Table 1.

As seen in Table 1, for all three data sets the number of trials for training is comparable to the number coefficients to be trained. This typically leads to serious problems in overtraining. We mitigate these by including multiple training samples for each trial.³ We would expect that increasing the number of independent training samples (e.g., trials) would similarly increase performance of the results presented below.

Figure 1 shows results for the data set used to predict whether a subject will press a button with their left or right hand by analyzing the MEG signals in a window prior to the button push (left hand = 1, right hand = 0 in the logistic regression model). We use an analysis window 100 ms wide centered at 83 ms prior to the button event, which at 300 Hz corresponds to

³ These samples are obviously not independent; however, they provide evidence for the natural variation of the data and thus make the estimates much more robust. They were shown, through cross-validation, to improve estimated generalization performance.

FIG. 1. MEG left/right button push prediction. (A) Trial averages $\bar{y}_c(t)$ (solid curves) and standard deviations (dotted curves) of discriminating component for left (red) and right (blue) button pushes. Time is given in seconds. Black vertical line at $t = 0$ s indicates timing of button push. Green vertical lines indicate the discrimination window. (B) Sensor projections for discrimination vector. (C) ROC curve for left vs right discrimination. Area under the curve (A_z) = 0.93. (D) Dipole-fit of \mathbf{a} overlaid on MRI image.

FIG. 2. Discrimination of imagined left/right finger taps. (A) Trial averages $\bar{y}_c(t)$ (solid curves) and standard deviations (dotted curves) of discriminating component for left (red) and right (blue) imagined. Time is given in seconds. Black vertical solid line at $t = 0$ s indicates timing of visual stimulus that defines the action to be performed (left or right imagine). The subjects are trained to execute the task at around $t = 1:25$ s. Green vertical lines indicate the discrimination window. (B) Dorsal view of sensor projections \mathbf{a} . (C) ROC curve for left vs right discrimination. For this subject the fraction of correct classification is $P = 0.79$ which corresponds to an information transfer of 0.26 bits/trial. (D) Sensor projection for same subject for explicit finger tap. (E) ROC curve for same subject for explicit finger tap.

$T = 30$. Figure 1 shows the results for one subject (AT). In the trial-averaged results one can see significant separation of the means for left vs right button push within the analysis window. Given that this separation is approximately equal to one standard deviation, this suggests that single-trial discrimination is possible. Single-trial discrimination is shown in the ROC curve, which for this subject exhibits good discriminability ($A_z = 0.93$). Figure 1 also shows the sensor projection \mathbf{a} and the location of a dipole-fit for this projection. A single equivalent current dipole fits the data with an accuracy of 64% using the least-squares “xfit” routine from Neuromag⁴ and assuming a spherical head model. This compares favorably with the 50% goodness-of-fit typically obtained for somatosensory responses when using all 122 sensors (Tang *et al.*, 2002). When considered with respect to the motor-sensory homunculus, these results indicate that the discrimination component activity originates in the sensory-motor cortex corresponding to the left hand.

Figure 2 shows results for the second data set, where the goal is to detect activity associated with a purely imagined motor response, a situation perhaps more relevant for a BCI system. Subjects are trained to imagine a tap with the left or right index finger synchronized to a brief, temporally predictable signal. Therefore there exists a known time window in which one can explicitly look for activity that discriminates between the left and the right imagined conditions. We selected an 0.8-s time window around the time where the task is to be performed; 90 left and 90 right trials were available to train the coefficients of the 59 EEG sensors. The result for the best performing subject is $A_z = 0.90$, shown in Fig. 2. The sensor projection of the 59 EEG sensors shows a clear left-right polarization over the motor area. In the context of BCI the metric of interest is the bit rate at which information can be transmitted with imagined motor activity. The information transmitted per trial is given by

$$I = 1 + p \log_2(p) + (1 - p) \log_2(1 - p), \quad (4)$$

where p is the fraction correct. For the subject shown in Fig. 2 this corresponds to $I = 0.26$ bits/trial. When averaged over the nine subjects the information transmitted is $I = 0.16$ bits/trial. Note that with a repetition rate of 6 s this experiment is not designed for an optimal transmission rate. Assuming equivalent performance with a repetition rate of 0.8 s (corresponding to the time window used here for discrimination) we obtain an average bit rate of 12 bits/min.

For comparison, we compare our approach with an alternative method, first described by Wolpaw *et al.* (1991), that is based on differences in the power spec-

trum in electrodes over the left and right motor cortex. Andersen *et al.* (1998) modify the approach by using six autoregressive (AR) coefficients to model the power spectrum of each electrode within the analysis window and classify the imagined conditions using a linear discrimination on these AR coefficients. Following Penny *et al.* (2000), we used electrodes C3 and C4 (international 10/20 electrode placement system—see Towle *et al.*, 1993) and obtain $A_z = 0.65 \pm 0.09$, and fraction correct of $p = 0.62 \pm 0.07$, which corresponds to $I = 0.054$ bits/trial or a bit rate of 4 bit/min. This is about one-third of the bit rate produced by our method of linear spatial integration.

The results, across the nine subjects, for predicting explicit finger taps from a window 300 to 100 ms prior to the taps is $A_z = 0.87 \pm 0.08$ with a fraction correct of 0.80 ± 0.08 . As shown for subject PJM in Fig. 2, sensor projections of the discrimination vector for explicit motor response are similar to the projections of the imagined motor response. This is consistent with previous findings in EEG and fMRI (Cunnington *et al.*, 1996; Porro *et al.*, 1996) and supports the approach of many current BCI systems—signals arising from the cortical areas that encode an explicit movement are also in some sense optimal for detecting the imagined movement.

Figure 3 shows the results for the target detection experiments where the goal is to detect the ERN on a single-trial basis. The ERN has a medial-frontal distribution that is symmetric to the midline, suggesting a source in the anterior cingulate (Dehaene *et al.*, 1994). It begins around the time of the perceived incorrect response and lasts roughly 100 ms thereafter. We use this time window for detection; 40 to 80 error trials and 300 correct trials were used for training and testing 64 coefficients. The sensor projection, shown in Fig. 3 for one subject, is representative of the results obtained for other subjects and is consistent with the scalp topography and time course of the ERN. The detection performance for this subject is $A_z = 0.84$ and is to be compared to $A_z = 0.63$ when detecting ERN from the front-center electrode where maximal activity is expected (FCz in the 10/20 system).

Our results demonstrate the utility of linear analysis methods for discriminating between different events in single-trial, stimulus-driven experimental paradigms using EEG and MEG. A particularly important aspect of our approach is that linearity enables the computation of sensor projections for the optimally discriminating weighting. This localization can be compared to the functional neuroanatomy, serving as a validation of the data-driven linear methods. In all three cases presented, we find that indeed the activity distribution correlated with the component that optimizes single-trial discrimination localizes to a region that is consistent with the functional neuroanatomy. This is important, for instance, to determine whether the discrimination model is capturing information directly related to the underlying task-depen-

⁴ See www.neuromag.com.

dent cortical activity, or is instead exploiting an indirect cortical response or other physiological signals correlated with the task (correlations with the stimulus, eye movements, etc.). Localization of the discriminating component activity, and its correlates, also enables one to determine the neuroanatomical correlations between different discrimination tasks, as was demonstrated for explicit and imagined motor responses in EEG.

In this work we have focused on spatial linear filtering of the magnitude of electrical or magnetic activity. This is consistent with the conventional concept of evoked responses, which capture magnitude differences under different experimental conditions. Note that in this approach the sign of the response is relevant for discrimination. Our approach does not consider temporal patterns of activity. An alternative and common approach is to consider signal powers after linear filtering in time (Wolpaw *et al.*, 1991; Pfurtscheller and Neuper, 2001; Anderson *et al.*, 1998). When considering powers (i.e., squared magnitudes), the sign of the activity is irrelevant and instead one typically captures oscillatory temporal patterns of activity. Both coherent evoked responses and oscillatory activity have been important analysis tools for encephalography, and work to combine optimal linear spatial and temporal filtering is beginning to emerge (Ramoser *et al.*, 2001).

A potential disadvantage of the method we have described is that its utility is derived from integration across a large sensor array, which may be impractical for a realistic BCI system (e.g., may be too cumbersome and/or costly). In such cases, the method still offers utility as a means for identifying a reduced sensor configuration, since the computed sensor projection can be used to prune the array in a way that minimizes loss in discrimination performance.

We close by noting that the proposed method is applicable to other encephalographic modalities with linear superposition of activity, such as functional near infrared imaging (Boas *et al.*, 2001).

APPENDIX

Derivation of Sensor Projection \mathbf{a}

Suppose our observation vector is \mathbf{x} and we have built a linear classifier, $y_1 = \mathbf{v}^T \mathbf{x}$, where y_1 is the binary number indicating some cognitive event that we are trying to detect. We assume that there are a number of such cognitive events occurring simultaneously. We represent these as a vector of binary indicators \mathbf{y} , with y_1 as its first element, and a matrix \mathbf{A} that maps these to the observation vectors; i.e., $\mathbf{x} = \mathbf{A}\mathbf{y}$. Without restriction we normalize \mathbf{y} to be zero mean. We wish to identify this mapping, namely to find the first column of \mathbf{A} , which we call \mathbf{a} and which is defined as the observation vector that we would obtain if only y_1 occurred. The most likely \mathbf{a} can be found as follows. Let

\mathbf{X} be the zero mean observation matrix for many samples, i.e., the t^{th} column is the observation for the t^{th} sample. Let \mathbf{y}_1^T be the corresponding binary column vectors across these samples given by $\mathbf{y}_1 = \mathbf{v}^T \mathbf{X}$. The definition for \mathbf{a} implies $\mathbf{X} = \mathbf{a}\mathbf{y}_1$. The maximum likelihood estimate for \mathbf{a} , given \mathbf{v} and \mathbf{X} , is given by the least-squares solution, $\mathbf{a} = \mathbf{X}\mathbf{y}_1^T(\mathbf{y}_1\mathbf{y}_1^T)^{-1}$.

We would like to determine the conditions under which the least-squares estimate of \mathbf{a} is actually proportional to the first column of \mathbf{A} . Let the matrix \mathbf{Y} be the binary matrix of the simultaneous cognitive events across trials, i.e., the t^{th} column is the cognitive events vector \mathbf{y} for the t^{th} trial. Since $\mathbf{X} = \mathbf{A}\mathbf{Y}$, we find that $\mathbf{a} = \mathbf{A}\mathbf{Y}\mathbf{y}_1^T(\mathbf{y}_1\mathbf{y}_1^T)^{-1}$. Note that \mathbf{Y} has dimensions of number of cognitive events (N) by number of samples (T), and that the quantity $\mathbf{Y}\mathbf{y}_1^T$ is the column vector of unnormalized correlations between the event indicators y_1 and the set of all cognitive events. If this is proportional to the Kronecker delta, $\delta_{i,1}$ (i.e., y_1 is uncorrelated with the indicators of the other events), then $a_i \propto \sum_j A_{i,j} \delta_{j,1} = A_{i,1}$, and therefore \mathbf{a} is proportional to the first column of \mathbf{A} .

ACKNOWLEDGMENTS

This work was supported in part by a grant from the Defense Advanced Research Project Agency, the National Foundation for Functional Brain Imaging, and the National Institutes of Health (P50 MH62196 and R01 NS37528). We thank Clay Spence, Adam Gerson, and Zuohua Zhang for fruitful discussions and assistance in the data analysis.

REFERENCES

- Anderson, C., Stolz, E., and Shamsunder, S. 1998. Multivariate autoregressive models for classification of spontaneous electroencephalogram during mental tasks. *IEEE Trans. Biomed. Eng.* **45**(3): 277–286.
- Baillet, S., Mosher, J. C., and Leahy, R. M. 2001. Electromagnetic brain mapping. *IEEE Signal Process. Mag.* **18**(6): 14–30.
- Boas, D. A., *et al.* 2001. Imaging the body with diffuse optical tomography. *IEEE Signal Process. Mag.* **18**(6): 57–75.
- Coles, M. G. H., and Rugg, M. D. 1995. Event-related brain potentials: An introduction. In *Electrophysiology of Mind*. (M. D. Rugg and M. G. H. Coles, Eds.). Oxford Univ. Press, Oxford.
- Cunnington, R., Iansak, R., Bradshaw, J. L., and Phillips, J. G. 1996. Movement-related potentials associated with movement preparation and motor imagery. *Exp. Brain Res.* **111**(3):429–36.
- Dehaene, S., Posner, M., and Tucker, D. 1994. Localization of a neural system for error detection and compensation. *Psychologic. Sci.* **5**: 303–305.
- Duda, R., Hart, P., and Stork, D. 2001. *Pattern Classification*, 2nd ed. J Wiley, New York.
- Falkenstein, M., Hoorman, J., Christ, S., and Hohnsbein, J. 2000. ERP components on reaction errors and their functional significance: A tutorial. *Biologic. Psychol.* **51**: 87–107.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., and Donchin, E. 1993. A neural system for error detection and compensation. *Psychologic. Sci.* **4**(6): 385–390.
- Gehring, W. J., and Willoughby, A. R. 2002. The medical frontal cortex and the rapid processing of monetary gains and loss. *Science* **295**: 2279–2282.

- Haykin, S. 1996. *Adaptive Filter Theory*. Prentice-Hall, Englewood Cliffs, NJ.
- Makeig, S., Bell, A., Jung, T., and Sejnowski T. 1996. Independent component analysis of electroencephalographic data. In *Advances in Neural Information Processing Systems*, Vol. 8; pp. 145–151. MIT Press.
- Penny, W. D., Roberts, S. J., Curran, E., and Stokes, M. 2000. EEG-based communication: A pattern recognition approach. *IEEE Trans. Rehab. Engin.* **8**(2): 214–215.
- Pfurtscheller, G., and Neuper, C. 2001. Motor imagery and direct brain-computer communication. *Proc. IEEE* **89**(7): 1123–1134.
- Porro, C. A., Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuiani, C., Bazzocchi, M., and di Prampero, P. E. 1996. Primary motor and sensory cortex activation during motor performance and motor imagery: A functional magnetic resonance imaging study. *J. Neurosci.* **16**(23): 7688–7698.
- Ramoser, H., Mueller-Gerking, J., and Pfurtscheller, G. 2000. Optimal spatial filtering of single trial EEG during imagined hand movement. *IEEE Trans. Rehab. Engin.* **8**(4): 441–446.
- Schalk, G., Wolpaw, J., McFarland, D., and Pfurtscheller, G. 2000. EEG-based communication: Presence of an error potential. *Clin. Neurophysiol.* **111**: 2138–2144.
- Swets, J. A. 1979. ROC analysis applied to the evaluation of medical imaging techniques. *Investig. Radiol.* **14**: 109–121.
- Tang, A., Pearlmutter, B., Malaszenko, N., Phung, D., and Reeb, B. 2002. Localization of independent components of magnetoencephalography in cognitive tasks. *Neural Comput.* **14**(8): 1827–1858.
- Towle, V., Bolanos, J., *et al.* 1993. The spatial location of EEG electrodes: Locating the best-fitting sphere relative to cortical anatomy. *Electroenceph. Clin. Neurophysiol.* **86**(1): 1–6.
- Vigario, R., Sarela, J., Jousmaki, V., Hamalainen, M., and Oja, E. 2000. Independent component approach to the analysis of EEG and MEG recordings. *IEEE Trans. Biomed. Engin.* **47**(5): 589–593.
- Wolpaw, J. R., McFarland, D. J., Neat, D. J., and Fomeris, C. A. 1991. An EEG-based brain–computer interface for cursor control. *Electroenceph. Clin. Neurophysiol.* **78**: 252–259.