NeuroImage 51 (2010) 1425–1437

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Predicting variations of perceptual performance across individuals from neural activity using pattern classifiers

Koel Das, Barry Giesbrecht, Miguel P. Eckstein *

Department of Psychology and Institute for Collaborative Biotechnologies, University of California Santa Barbara, Santa Barbara, CA 93106-9660, USA

ARTICLE INFO

Article history:
Received 10 December 2009
Revised 2 March 2010
Accepted 9 March 2010
Available online 17 March 2010

Keywords:
Neural correlates of human behavior
ERP analysis
Pattern classifiers

ABSTRACT

Within the past decade computational approaches adopted from the field of machine learning have provided neuroscientists with powerful new tools for analyzing neural data. For instance, previous studies have applied pattern classification algorithms to electroencephalography data to predict the category of presented visual stimuli, human observer decision choices and task difficulty. Here, we quantitatively compare the ability of pattern classifiers and three ERP metrics (peak amplitude, mean amplitude, and onset latency of the face-selective N170) to predict variations across individuals' behavioral performance in a difficult perceptual task identifying images of faces and cars embedded in noise. We investigate three different pattern classifiers (Classwise Principal Component Analysis, CPCA; Linear Discriminant Analysis, LDA; and Support Vector Machine, SVM), five training methods differing in the selection of training data sets and three analyses procedures for the ERP measures. We show that all three pattern classifier algorithms surpass traditional ERP measurements in their ability to predict individual differences in performance. Although the differences across pattern classifiers were not large, the CPCA method with training data sets restricted to EEG activity for trials in which observers expressed high confidence about their decisions performed the highest at predicting perceptual performance of observers. We also show that the neural activity predicting the performance across individuals was distributed through time starting at 120 ms, and unlike the face-selective ERP response, sustained for more than 400 ms after stimulus presentation, indicating that both early and late components contain information correlated with observers' behavioral performance. Together, our results further demonstrate the potential of pattern classifiers compared to more traditional ERP techniques as an analysis tool for modeling spatiotemporal dynamics of the human brain and relating neural activity to behavior.

Introduction

A fundamental goal of cognitive neuroscience is to understand how the human brain processes and represents the environment and how these representations are used to guide adaptive behavior. Traditionally, the link between brain activity and information processing has been investigated by revealing correlations between measures of neural activity and measures of behavioral performance. The utility of electroencephalography (EEG) in making this link was recognized by Woodworth (1938) soon after Hans Berger (1929) reported the first published observations of the human electroencephalogram. EEG is particularly well suited to investigating brain–behavior relationships because of its temporal resolution, but also because the EEG signal can be characterized by various features (e.g., latency, amplitude, power, phase), each of which can be used to relate neuronal responses with behavioral measures of information processing efficiency (response time, accuracy, sensitivity). For instance, the earliest published reports related EEG amplitudes and latencies with overt conditioned responses (Walter et al., 1964), trial-by-trial variations in response time (RT) (Kutas et al., 1977), and individual differences in intelligence (Ertl and Schafer, 1969). More recently, correlations have been established between ERP component amplitudes and changes in perceptual sensitivity during spatial attention tasks (Mangun and Hillyard, 1988,1990) and individual differences in working memory capacity (Vogel and Machizawa, 2004; Vogel et al., 2005). Complementing the analyses of latency and amplitude, power modulations in specific frequency bands have also been linked with behavioral performance (Ergenoglu et al., 2004; Hanslmayr et al., 2007; Klimesch et al., 1993).

While the traditional approach of using EEG to investigate the relationship between neural activity and information processing measured by behavioral performance on specific cognitive tasks has provided key insights into brain–behavior relationships, this approach does not take full advantage of the multivariate nature of the EEG data: particularly the fact that task-related neural activity is distributed across

* Corresponding author. Office: Psychology East (Building 251), Room 3806, USA. Fax: +1 805 893 4303.
E-mail addresses: das@psych.ucsb.edu (K. Das), giesbrecht@psych.ucsb.edu (B. Giesbrecht), eckstein@psych.ucsb.edu (M.P. Eckstein).

1053-8119/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2010.03.030

time. Multivariate pattern classifiers provide a method of integrating neural activity into a decision variable that can be used to make categorical decisions on a trial-by-trial basis for which performance metrics (i.e., accuracy of decisions) directly comparable to behavioral performance (i.e. single-trial EEG; Gerson et al., 2005; Philiastides et al., 2006; Philiastides and Sajda, 2006) can be computed. The technique has been applied both to EEG (Philiastides et al., 2006; Philiastides and Sajda, 2006) and functional magnetic resonance imaging, fMRI, (Haynes and Rees, 2005; Kamitani and Tong, 2005; Norman et al., 2006). For the case of EEG, studies have used pattern classifiers to successfully predict the visual stimulus presented to the observer, observer choices and task difficulty (Philiastides et al., 2006; Philiastides and Sajda, 2006).

The aim of the present study is to quantitatively compare pattern classifiers and more traditional ERP metrics in their ability to predict the variability in perceptual performance across different individuals from EEG activity. First, we compared the recently proposed single-trial EEG analysis using pattern classifiers to metrics related to more traditional methods based on EEG-trial averaging (Event Related Potentials, ERP). In particular, we compare the pattern classifiers to traditional ERP metrics (peak amplitude, mean amplitude and peak latency) in the ability to predict behavioral performance in an object identification task (face vs. car) across twenty human participants. The face/car paradigm evokes well-known EEG components including N1, N170 (Gauthier et al., 2003; Taylor et al., 1999) and in the past have been used to explore links between brain and behavior using pattern classifiers (Philiastides et al., 2006; Philiastides and Sajda, 2006). Second, we compared the predictive ability of different classifier algorithms (LDA, SVM and CPCA) and various training methods including using neural data restricted to correct decision trials and high decision confidence trials. Finally, but most importantly, we investigated the temporal distribution of classification performance to elucidate the time-epochs coding neural activity that is predictive of observers' perceptual performance.

Materials and method

Experimental setup

Screen displays of face and car were presented as stimuli to the observers who performed the behavioral task of identifying the correct label of the image (face/car) and their neural signals were recorded via electroencephalography. The details of the experiment are given in the following section.

Stimulation and display

The stimuli set consists of 290×290 pixel 8-bit gray scale images of faces and cars taken from the Max Planck Institute for Biological Cybernetics face database (Troje and Bülthoff, 1996). Images were displayed on an 19-inch ViewSonic Color E90F monitor (resolution 1024×768) with a refresh rate of 75 Hz in a darkened room. All images were filtered to achieve a common frequency power spectrum (the average of all images). Twelve images of each class, face and car, (six frontal view, six 45° rotated) were used as stimuli. Gaussian noise was then added to these 24 base images to build a stimuli set of 1000 images (500 face, 500 car). Noise was generated by filtering independent white Gaussian noise fields (standard deviation of 3.53 cd/m^2) by the average power spectrum of the car/face stimuli. The noise fields were added to the original car/face images. Observers were placed at a distance of 125 cm from a display set at a mean luminance of 25 cd/m^2 with a maximum luminance of 50 cd/m^2 and images subtended 4.57° of visual angle. Contrast energy (CE) of all face and car stimuli were matched to be 0.3367°^{2}, where CE is defined as the sum of the squared contrast values of the stimuli multiplied by the spatial extent of a pixel: $CE = \sum\sum \frac{s^2(x,y)}{L_o^2}\Delta x \Delta y$, where s = luminance value of signal at pixel(x, y) and L_o is the mean luminance.

Observers and procedure

Twenty naive observers (ages: 18–26) participated in the study. Observers were initially presented with 1000 stimuli-familiarization trials on the first day and 100 more practice trials immediately preceding the current experiment on the second day. No EEG activity was measured during the practice sessions. The actual study consisted of 1000 trials split into 5 successive sessions, each having 200 trials. After each session observers had a break. The observers fixated on a central cross and pressed a mouse button to indicate the beginning of the trial. The stimulus appeared for 40 ms after a variable delay of 0.5 − 1.5 s. The stimulus was followed by a blank screen presented for 0.5 − 1.5 s after which the response window was presented (Fig. 1). Observers were asked to rate how confident they were that they saw either a face or a car, with a rating of 1 indicating complete confidence that a face was presented and a rating of 10 indicating complete confidence that a car was presented. Confidence responses were recorded by mouse clicks on the rating buttons of the response window. Participants were instructed only to move the mouse when the response window appeared and premature mouse clicks were given feedback and recorded.

Electroencephalogram data acquisition and pre-processing

Each subject's electroencephalogram was recorded from 64 Ag/AgCl sintered electrodes mounted in an elastic cap and placed according to the International 10/20 System. The horizontal and vertical electro-oculograms (EOG) were recorded from electrodes placed 1 cm lateral to the external canthi (left and right) and above and below each eye, respectively. The data were sampled at 512 Hz, re-referenced offline to the signal recorded from the central midline electrode (Cz), and then band-pass filtered (0.01−100 Hz). Trials containing ocular artifacts (blinks and eye movements) detected by EOG amplitudes exceeding ±100 mV or by visual inspection were excluded from the analysis. The average ERP waveforms in all conditions were computed time-locked to stimulus onset and included a 200 ms pre-stimulus baseline and 500 ms post-stimulus interval.

Analysis

The analysis of the EEG data used both classical ERP measures and pattern classifiers.

ERP metrics

ERP activity was quantified using three measures, peak N170 amplitude, mean N170 amplitude and peak latency for N170 for face and car trials and difference wave. Fig. 2 (top row) shows the mean face trial and mean car trial averaged over all 20 observers for electrode PO8. The peak amplitude and latency was measured using an automatic peak detection routine within a window of 140−210 ms after onset of the stimulus on the ERP. For each subject and condition, mean N170 was calculated using a ±40 ms window centered on the peak N170 latency. The difference wave was constructed by

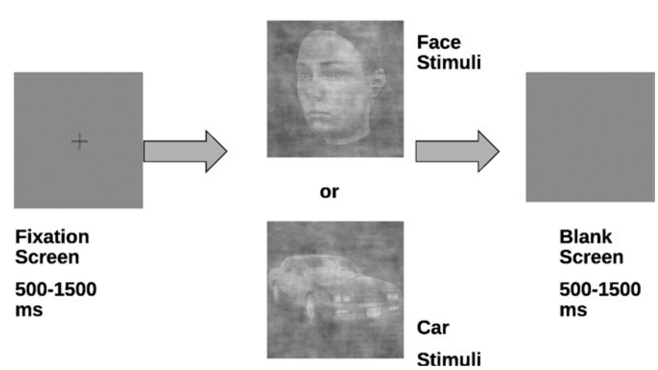

Fig. 1. Illustration of psychophysical procedure for study.

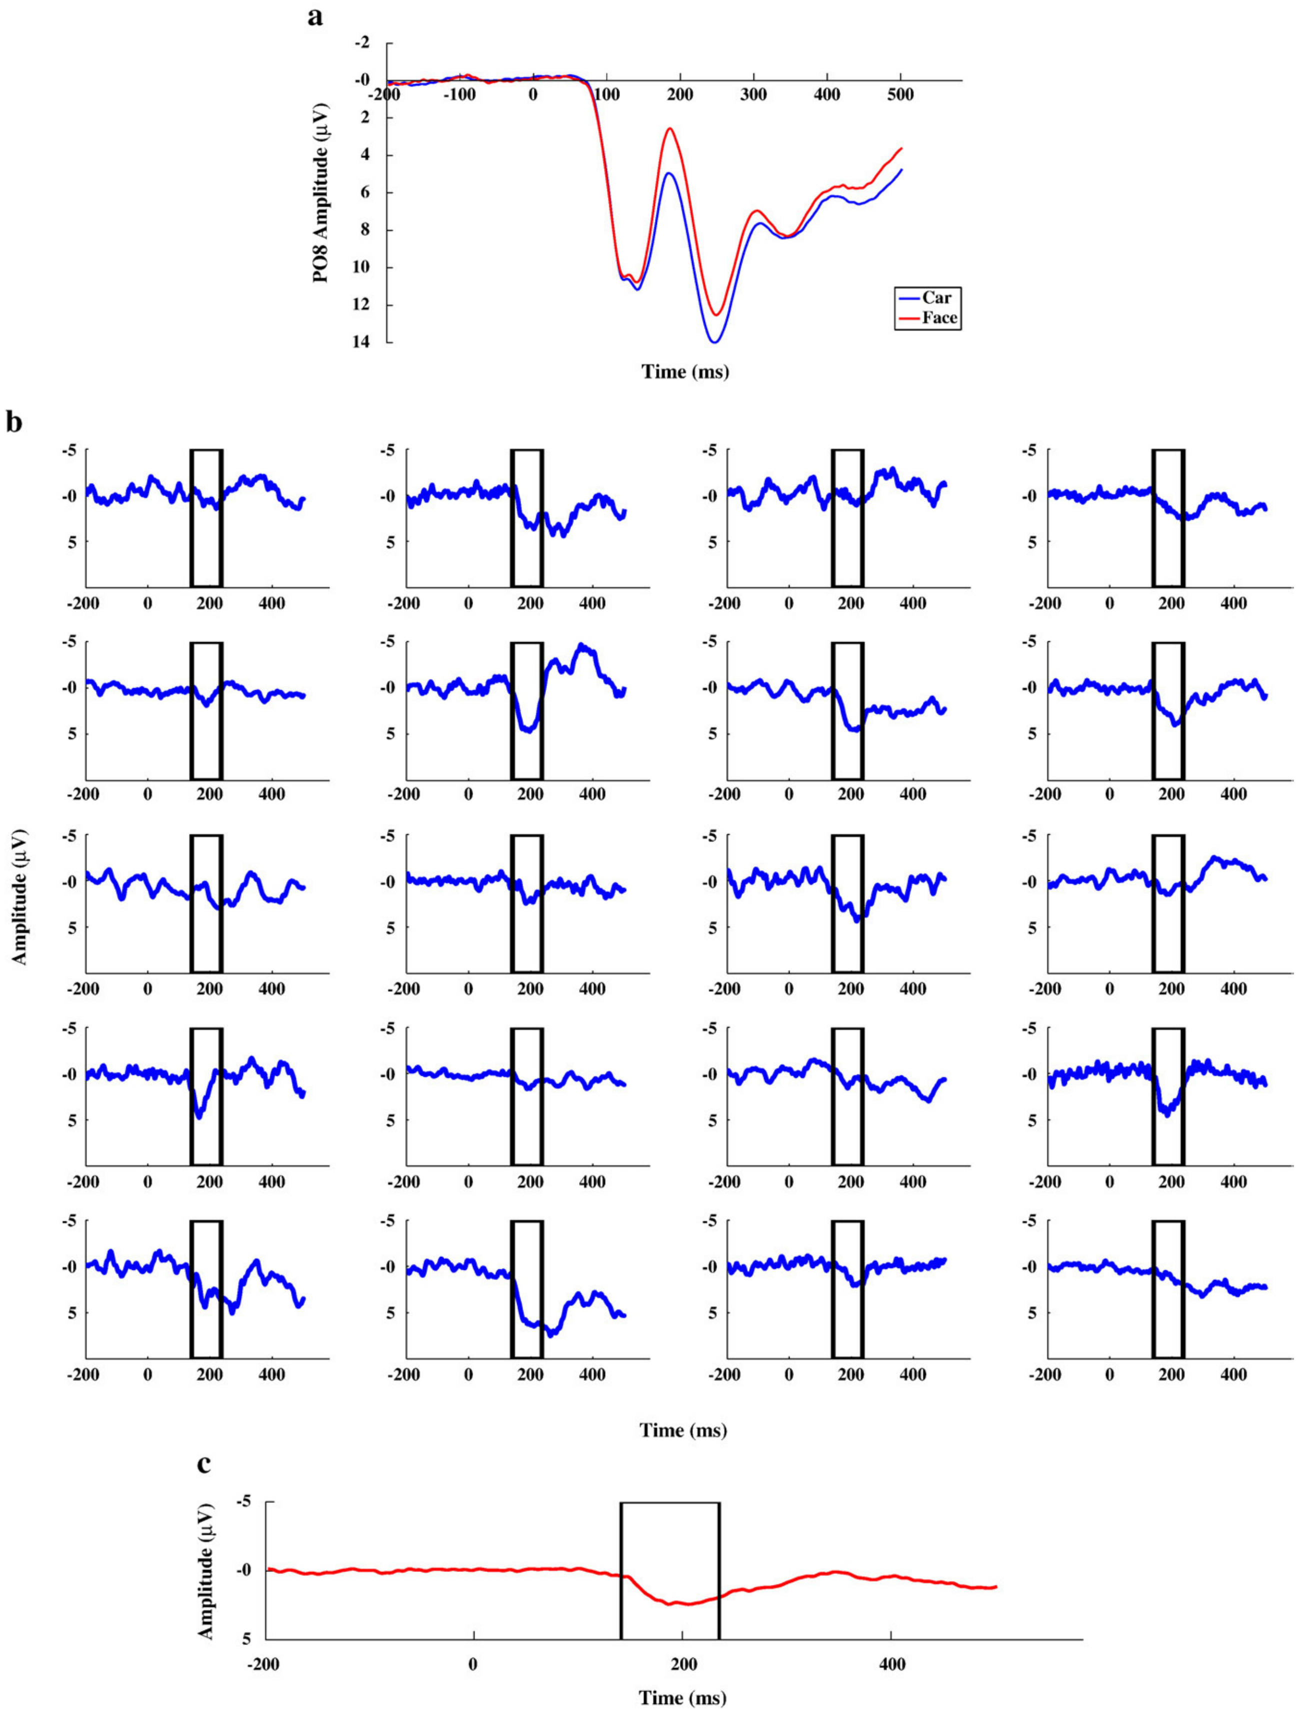

Fig. 2. ERP response: ERP for face and car trials for PO8 electrode averaged across all observers (top row). Time interval (140–210 ms) used for N170 analysis: difference wave for PO8 for 20 observers shown in blue which includes all the peaks (middle row). The mean difference wave for 20 observers is shown in red (bottom row).

subtracting the mean face trial from the mean car trial. Fig. 2 illustrates using difference wave, the time window taken for each observer for PO8 for peak detection and verifies that all peaks are included during N170 analysis.

Pattern classifier-based metrics

Recently, multivariate pattern classifiers has been successfully used in research related to functional magnetic resonance (Haynes and Rees, 2005; Kamitani and Tong, 2005; Norman et al., 2006) and/or EEG (Parra et al., 2005; Philiastides et al., 2006; Philiastides and Sajda, 2006). In this study, we investigated the efficacy of three types of pattern classifiers in predicting the variability of perceptual performance across individuals. Single-trial EEG classification is a challenging task since the data is typically of high dimensionality and suffers from the *small sample size problem*, which arises when the dimensionality of the data exceeds the size of the training database due to limitations in the training time. Under the small sample size conditions, a large portion of the data space is sparse and carries very little or no useful information. The irrelevant subspace is discarded through extraction of a small set of useful features in order to obtain meaningful data statistics. Here we have used classical discriminant analysis tools like Linear Discriminant Analysis, non-parametric pattern classifiers (Support VectorMachines) and a recently proposed non-linear classification technique (Classwise Principal Component Analysis) to analyze the EEG signals.

Definition of variables. Let $X \in \mathbb{R}^{N_e \times t}$ be a 500 ms long segment of EEG signal, where N_e is the number of electrodes ($N_e = 63$, electrode Cz being the reference electrode is not considered) and t is the number of time points sampled at 512 Hz ($t = 0.500 \times 512 = 256$). Let $x \in \mathbb{R}^n$ be a vectorized version of $\mathbf{X}$, where $n = N_e \times t$. The vectorization scheme can be chosen arbitrarily, but must be applied consistently. Note that the n coordinates of x contain both spatial and temporal information, which can be recovered by applying the inverse of the vectorization scheme. For temporal window analysis, $X \in \mathbb{R}^{N_e \times 20}$ is a 40 ms long EEG signal, ($t = 0.040 \times 512$ and $N_e = 63$). Let the number of trials be N_t ($=1000$) and $\{x_1, \ldots, x_{N_t}\}$ be the input data fed to each of the three pattern classifiers.

Linear Discriminant Analysis (LDA). LDA is perhaps the most widely used feature extraction technique. The objective of LDA is to perform dimensionality reduction while enhancing class separability, normally by maximizing an objective function. The most popular form of LDA relies on the maximization of the Fisher critereon (Fisher, 1936):

$$J(\mathbf{T}) = \arg \max_{T \in \mathbb{R}^{m \times p}} \frac{|\mathrm{T} S_b \mathrm{T}^{\mathrm{T}}|}{|T S_w T^{\mathrm{T}}|}, \quad (1)$$

where $\mathbf{S}_b$ and $\mathbf{S}_w$ represent between-class and within-class scatter matrices (Duda et al., 2001; Fukunaga, 1990). $\mathbf{T}$ is a set of eigenvector associated eigenvalues based on the generalized eigen-decomposition. However, for large-scale data, the traditional LDA approach faces a couple of challenges. Firstly, large-scale data is associated with large covariance matrices which are difficult to store and manipulate. Secondly, large-scale data inevitably leads to the small sample size conditions ($n > N_t$) and singular scatter matrix Sw, and so the classical solution to Eq. (1) based on the eigenvalue/eigenvector decomposition of $\mathbf{S}_w^{-1}\mathbf{S}_b$ (Duda et al., 2001; Fukunaga, 1990), is not directly applicable. We have used Principal Component Analysis (PCA) to reduce dimension of the EEG signals thereby avoiding the small sample size problem before applying LDA.

Support Vector Machines (SVM). In recent years, Support Vector Machines (SVM) (Vapnik, 1998) have been extensively used in machine learning and pattern recognition. The goal of SVM is to maximize the margin between two classes. This is achieved by picking the hyperplane so that the distance from the hyperplane to the nearest data point is maximized. Given data $\{x_1, \ldots, x_{N_t}\}$ with their labels $\{L_i = 1, -1\}$, the hyperplane takes the form $\mathbf{w}^{\top}\mathbf{x} - b = 0$,where $\mathbf{w}$ is a normal vector perpendicular to the hyperplane. Moreover, all the data satisfy the following constraint:

$$w^{\top} x_i - b \geq +1 \quad for \quad L_i = +1, \quad (2)$$

$$w^{\top} x_i - b \leq -1 \quad for \quad L_i = -1, \quad (3)$$

This constraint can be written as:

$$L_i \left(w^{\top} x_i - b \right) \geq 1 \quad (4)$$

Now, let us consider the hyperplanes when the two equalities in Eqs. (2) and (3) hold. These points lie on the hyperplane $H_1{:}\mathbf{w}^{\top}\mathbf{x}_i - b = 1$; and the points lie on the hyperplane $H_2{:}\mathbf{w}^{\top}\mathbf{x}_i - b = -1$. By using some geometry knowledge, the distance between H_1 and H_2 is $\frac{2}{||\mathbf{w}||}$. The problem now is to maximize $\frac{2}{||\mathbf{w}||}$, or equivalently, minimize $||\mathrm{w}||$ subject to the constraint (Eq. 4). More precisely,

$$\min \left(||\mathrm{w}||^2 \right) \quad s.t. \quad L_i \left(w^T x_i - b \right) \geq 1, \quad i = 1, \cdots, N_t. \quad (5)$$

This is a standard optimization problem, and can be solved by using the Lagrange method. For more details, please refer to (Vapnik, 1998). In this paper, we use linear SVM for the purpose of simplicity. Note here that SVM has no distribution assumption. This suggests that SVM might perform well under non-Gaussian distribution.

Classwise Principal Component Analysis (CPCA). Lastly, we used a recently proposed (Das and Nenadic, 2009) computationally efficient, locally adaptable, pattern classification technique for analyzing the EEG data. The method is based on classwise PCA (CPCA) and results in a simple piecewise linear dimensionality reduction technique. In this technique, the strength of PCA as a dimensionality reduction technique is exploited, while preserving the class-specific information to facilitate subsequent classification. The main idea behind the technique is to identify and discard a useless (non-informative) subspace in data by applying PCA to each class. The classification is then carried out in the residual space, where the small sample size conditions and the curse of dimensionality are no longer concerns. Fig. 11 (A) illustrates the major difference between the classical PCA method and CPCA, applied to a binary class case. The algorithm is a two step procedure. In the first step, CPCA extracts a piecewise linear subspace from the training trials and can be referred to as feature extraction step. In the second step, unknown test trial is classified in one of the subspaces estimated during the feature extraction step. The two steps of CPCA are explained in detail in the Appendix **A** section.

Training conditions for pattern classifier

Pattern classification performance was evaluated for 6 different training conditions as shown in Table 2. In the first case in Table 2, the pattern classifier was trained with EEG trials using actual labels of stimulus identity (face/car), regardless of the observer's decision (denoted as “Tr Stim (all trials)”). In the second case, the training database used the observer's choice of whether they thought the image was a face/car as the label rather than actual stimulus identity (“Tr Choice (alltrials)”). The third training condition rejected all trials where the observer's choice of face/car differed from the actual stimuli (“Tr Stim/Choice(correct trials)”). In this case, the training phase was based on only EEG response for correct trials, where the choice of the observer matched that of the actual stimuli identity. The last two training conditions considered ERPs where the observers made their choices with higher confidence (rating of 1:3 for face, and

8:10 for car). "Tr Stim (hi-conf trials)" in Table 2 assigns the actual stimuli as the label while "Tr Choice (hi-conf trials)" in Table 2 marks the choice of the observers as the label.

Evaluation using Area Under the ROC

Performance evaluation for individual observers of both behavioral ratings and neural decision variables was evaluated using Area under the Receiver Operating Curve (AUC). Performance as measured by AUC was calculated using a non-parametric method that quantifies for each scalar value (behavioral rating or neural decision variable using ERP measures/pattern classifiers) corresponding to the face stimuli (across F face trials), the probability that it will exceed the responses to all the variables corresponding to the car stimuli across all C car trials in the test data set.

$$AUC = \frac{1}{FC}\sum_{f=1}^{F}\sum_{c=1}^{C} step\left(\lambda_f - \lambda_c\right) + \frac{1}{2}\delta\left(\lambda_f - \lambda_c\right) \tag{6}$$

where λ_f is the scalar value corresponding to the *f*th face stimuli, λ_c is the scalar value to the *c*th car, *F* and *C* are the total number of face and car trials and *step* is the heavyside step function defined as:

$$step(x) = \begin{cases} 1, & if\ x > 0 \\ 0, & if\ x < 0 \end{cases}$$

The function δ is the impulse function defined as

$$\delta(x) = \begin{cases} 1, & if\ x = 0 \\ 0, & if\ x \neq 0 \end{cases}$$

The first term (step function) inside the summation measures the frequency in which a given response to a face stimuli exceeds the response to the car stimuli. The second term with the impulse function handles the instances in which the values for the face and car stimuli are a tie (i.e., the frequency of correct decisions is $\frac{1}{2}$ the frequency of the ties). Use of a parametric binormal model to calculate the Area under ROC Curve resulted in similar results to the non-parametric method.

Evaluation of ERP metric using Area Under the ROC

Area under the curve using ERP measures was calculated using two separate analyses. In the first analysis, mean ERP waveforms were computed for each of the 10 sessions (100 trials/session) separately for car and face trials. Peak amplitude, mean amplitude and latency were extracted for each mean ERP waveform. The resulting 10 scalar values for each ERP measure (e.g. peak amplitudes) for car trials and 10 scalar values for face trials were then used to compute AUC to evaluate the ability to discriminate car from faces. The calculated AUC was then correlated with subject behavioral performance (AUC).

In the second AUC analysis, the peak amplitude, mean amplitude and peak latency was extracted for each individual trial, and these EEG measures for car and face trials were then used to calculate the AUC. The AUC was subsequently correlated with AUC calculated from behavioral data for the 20 observers.

The first analysis extracts the ERP measures after averaging a subset of trials (within one session) while the second analysis extracts the EEG measures from individual trials.

Evaluation of pattern classifier metrics using Area Under the ROC

The analysis using different pattern classifiers was performed using 10-fold stratified cross validation (Kohavi, 1995). The dataset was randomly divided into 10 non overlapping folds of equal size, each having 100 trials. One of the folds was designated as test data while the remaining 9-folds constituted the training set. The pattern classifier performance was evaluated by using the area under the Receiver Operating Curve (ROC), AUC. To reduce the variance of estimated area under the curve (AUC), the overall 10-fold CV procedure was repeated 10 times, and the overall performance is given by the mean AUC.

Figures of merit to evaluate neural metrics

The ability of various metrics to predict behavioral perceptual performance was evaluated by using a Pearson linear correlation and rank ordering of individuals' performance based on neural metrics. The rank ordering measure is more similar to a Spearman Rank Correlation which does not penalize departures from linearity and is complementary to the Pearson correlation.

Pearson correlation

Once the ERP metrics and pattern classifiers were calculated, the sample correlation between each ERP metric (Table 1) and behavioral performance was calculated using Pearson's correlation (r), where

$$r = \frac{cov(x,y)}{\sigma_x \sigma_y} \tag{7}$$

cov is the covariance matrix and σ is the standard deviation between the behavioral data and the ERP metric (x and y). The standard error for each metric was calculated by performing a jack knife resampling technique (Duda et al., 2001). The jack knifing approach is a computational technique based on removing samples from the available dataset and recalculating the generic estimator. The ERP metrics included peak N170 amplitude, mean N170 amplitude and peak latency of the N170 for both face and car trials and the difference wave. The mean AUC was used for the pattern classifiers using different training conditions explained in the next section. Statistically significant correlation coefficients using a 95% confidence interval are emphasized in the results.

Neural metric-based rank ordering of individual's behavioral performance

We performed an additional 2 AFC simulation to further quantify the neural correlation between the ERP measures and the behavioral performance. In the simulation we chose any 2 random observers without replacement out of 20. We found the corresponding ERP measures (peak N170 amplitude and latency and mean N170 for P08 using the difference wave and AUC using pattern classifier) for each of the 2 chosen observers. For each of the 4 measures, we found the observer having the better ERP metric. We also found which of the chosen observers had better behavioral performance. If the observer having better ERP metric matched with the observer having better behavioral performance, it was considered a correct classification. We repeated this procedure for all possible combinations (190 simulation trials) and reported the average correct classification rate for all 4 measures. The standard error for rank ordering was computed using jack knife resampling technique.

Table 1
Summary of different EEG measures used for finding neural correlates.

Measures	Metric
ERP measures	N170 peak amplitude
	N170 mean amplitude
	N170 peak latency
	AUC using N170 peak amplitude
	AUC using N170 mean amplitude
	AUC using N170 peak latency
Pattern classifiers using all electrodes	AUC

Results

Correlations between neural measures and behavioral performance

ERP peak, mean, and latency

Based on classic studies showing that the N170 is larger over right lateral occipital electrodes than the corresponding left electrodes (Bentin and Allison, 1996) and more recent pattern classification studies (Philiastides et al., 2006; Philiastides and Sajda, 2006), we computed the correlation between the various ERP metrics and behavioral performance for a single electrode over right lateral occipital cortex (PO8). Fig. 3 plots each observer's peak N170, mean N170 and N170 latency ERP against the corresponding observer's behavioral performance (Area under the ROC curve, AUC). The measures were calculated using the ERP evoked on car trials (black), face trials (blue), and the difference wave (red). The peak N170 and the mean N170 calculated from the difference wave showed moderate positive correlations with behavior ($r=0.546$; $r=0.52$; $p<0.05$). Otherwise, when the metrics were calculated from the face and car trials separately, the correlations between ERP metrics and behavioral performance are low (~$0.08-0.23$, all $p>0.05$) and not statistically different from $r=0$. In addition, the correlation between the N170 latency and behavioral performance was low and not statistically significant ($p>0.05$), irrespective of whether it was based on EEG activity from car trials ($r=0.086$), face trials ($r=0.099$) or the difference wave ($r=0.16$).

To provide a more complete characterization of the spatial distribution of the correlation between each ERP metric and behavioral performance, we computed the correlation at all electrodes. The resulting topographical maps of the r-values are shown in Fig. 4. Spatial locations with dark red colors denote higher positive correlation values, dark blue colors denote higher negative correlations and gray corresponds to statistically insignificant correlations ($p>0.05$). Consistent with the analysis of the single electrode, the strongest correlations are observed when the metrics were computed based on the difference wave. Specifically, both peak and mean amplitude measures resulted in significant positive correlations over right lateral occipital, left temporal, and frontocentral electrodes. The N170 latency was negatively correlated with behavioral performance in frontal electrodes (both left and right). Relative to the correlation with the difference wave, the correlations between performance and the face and car trials shown are not robust and largely confined to midline occipital electrodes.

We computed the AUC using each ERP metric to provide a meaningful comparison between ERP metrics and pattern classifier metric. However as seen from the correlation coefficient in the scatter plots in Fig. 5, the correlation between behavioral performance (AUC) and ERP performance was not statistically significant ($r<0.44$; $p>0.05$) using all three ERP metrics. The low correlation is not surprising as ERP metrics show large trial-to-trial variations and the AUC computed from ERP measures do not provide robust neural correlates for predicting behavior. Note here that neural correlates

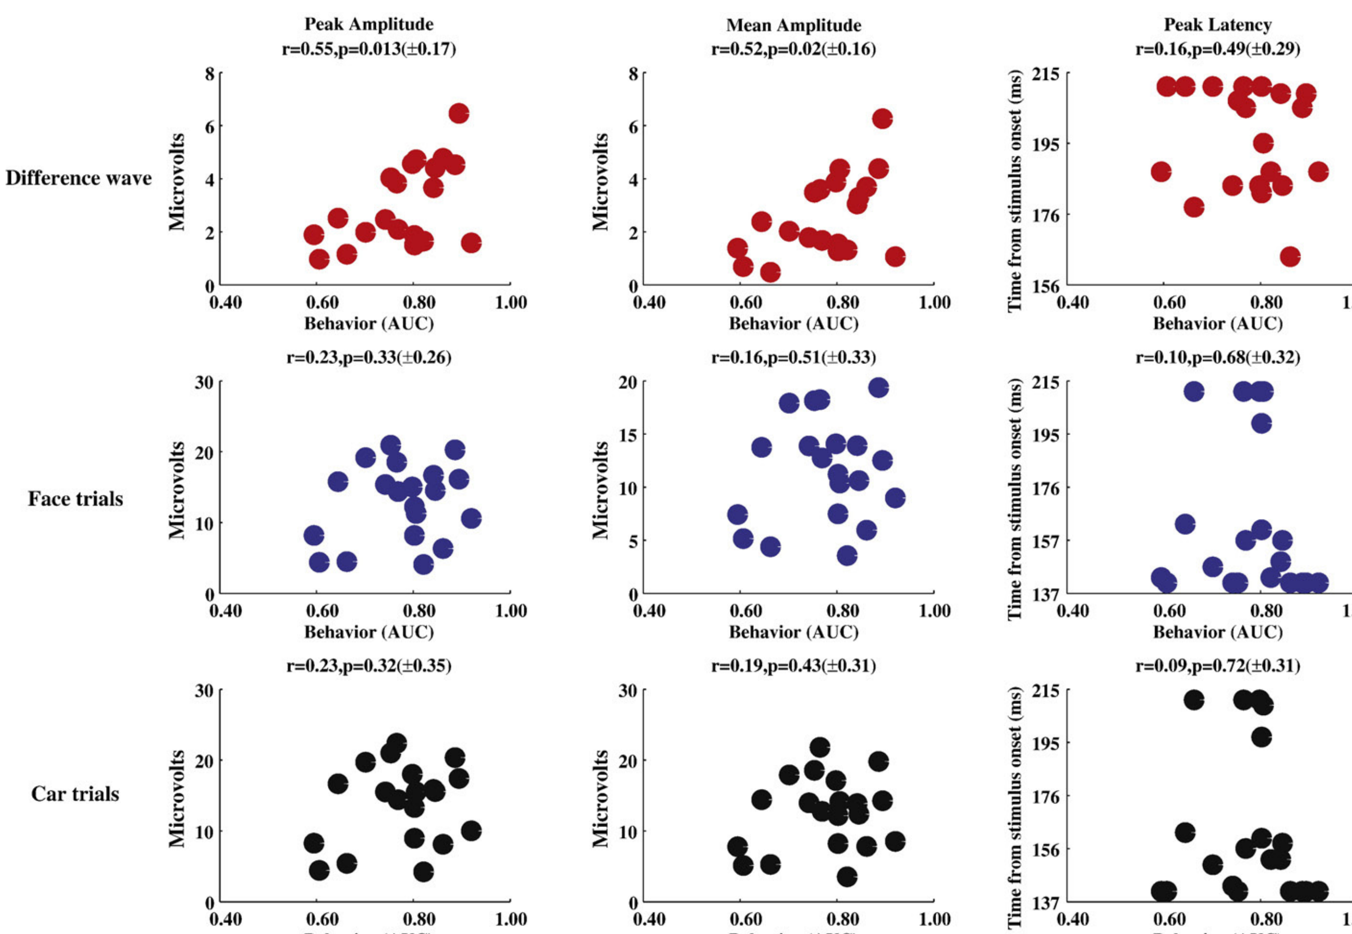

Fig. 3. Scatter plots showing correlation between behavioral performance (AUC) in face/car perceptual task and different ERP measures for P08 for 20 observers using difference wave shown in red (top row), face trials in blue (center row) and car trials in black (bottom row). ERP measure include peak N170 amplitude, mean N170 amplitude and peak N170 latency.

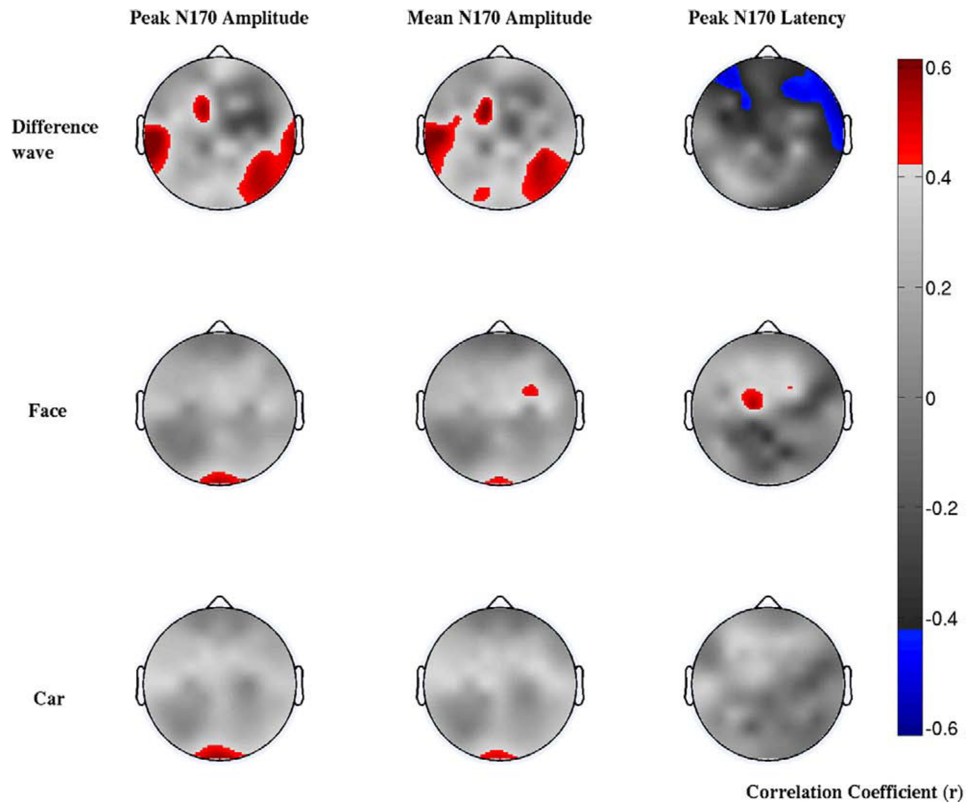

Fig. 4. Topographical map of correlation coefficients between ERP measures using all electrodes and behavioral performance (AUC) in face/car perceptual task for 20 observers using difference wave (top row), face trials (center row) and car trials (bottom row). Different ERP measures include peak N170 amplitude, mean N170 amplitude and peak N170 latency. Statistically insignificant correlation coefficients (r) using 95% confidence interval ($0.44>r>-0.44$) are shown in gray. All topographical maps are generated by EEGLAB (Delorme and Makeig, 2004) using all electrodes.

calculated using grand average of all trials are more stable and ERP metrics (peak N170 and mean N170) using difference wave correlates positively, albeit moderately, with the behavioral performance ($r>0.44$).

Pattern classifier performance (Area under the ROC, AUC)

We used three pattern classifiers (LDA, SVM, CPCA) to predict the perceptual performance of observers performing the visual task. The pattern classifiers were evaluated using AUC and Table 2 shows the correlation between pattern classifier performance (AUC) and behavioral performance (AUC) for different training conditions. We initially evaluated the pattern classifiers when it was trained using all trials and using high-confidence trials only (ratings of 1:3 for face and 8:10 for car). The correlation between one of the pattern classifiers (CPCA) and individual observer's performance for each training condition is shown in Fig. 6 (Left = all trials; right = high-confidence trials). Both modes of classifier training resulted in a statistically significant positive correlation ($r=0.67$; $p=0.001$ and $r=0.704$; $p=0.0005$) between classifier performance identifying car/face stimulus and subjects' behavioral performance. Fig. 7 shows the correlations between each of the EEG metrics investigated (pattern classifier and N170 peak, mean and latency based on ERP difference wave) and observers' behavioral performance. The correlation coefficients for the pattern classifier based metrics are higher than those from single electrode ERP metrics ($p<0.05$). In addition, the standard error of the correlations obtained from the pattern classifiers were lower than the correlations for the N170 peak and mean (e.g., stdterr = 0.17 for the N170 peak correlation vs. stderr = 0.07 for the pattern classifier correlation) suggesting that pattern classifiers are statistically more stable which is possibly related to their integration of EEG activity across electrodes.

In addition to the two training conditions described above, the classifiers were trained using three addition modes: a) training on all trials but with the EEG activity for each trial assigned to the category (car or face) based on the observers' choice rather than the actual category of the image (Tr Choice, all trials); b) training on EEG activity for those trials in which the observer made a correct decision (Tr Stim/Choice, correct trials); c) training on EEG activity for trials in which the observer expressed high confidence of their decision and with the EEG activity assigned to the category based on the observers' choice.

All three pattern classifiers resulted in statistically significant correlations with behavioral performance and all surpassed ($p<0.05$; Wilcoxon signed-ranks test) the ERP metrics. Across all training conditions the CPCA method resulted in higher correlations with behavioral performance ($p<0.005$; sign-rank test with Bonferroni correction for ten tests) but the differences in correlations across algorithms were modest (~0.013−0.058).

Table 2 summarizes the correlations between performance of different pattern classifiers and behavioral performance for five different training conditions. This analysis revealed some differences across the various training methods in the resulting correlation between the pattern classifiers' performance and the observers' behavioral performance (~0.03−0.08). In particular, all three pattern classifiers trained on neural data restricted to trials having high-confidence ratings correlated best with the behavioral performance and were shown to be statistically different from other training methods using pairwise comparisons of Wilcoxon signed-ranks test with Bonferroni corrections applied.

Correlation coefficients for all five training methods were statistically different from $r=0$ ($p<0.05$) for all the three pattern classifiers used. The remaining analyses using pattern classifiers and comparison with ERP metrics was performed using CPCA due to its superior performance but all pattern classifiers produced similar results.

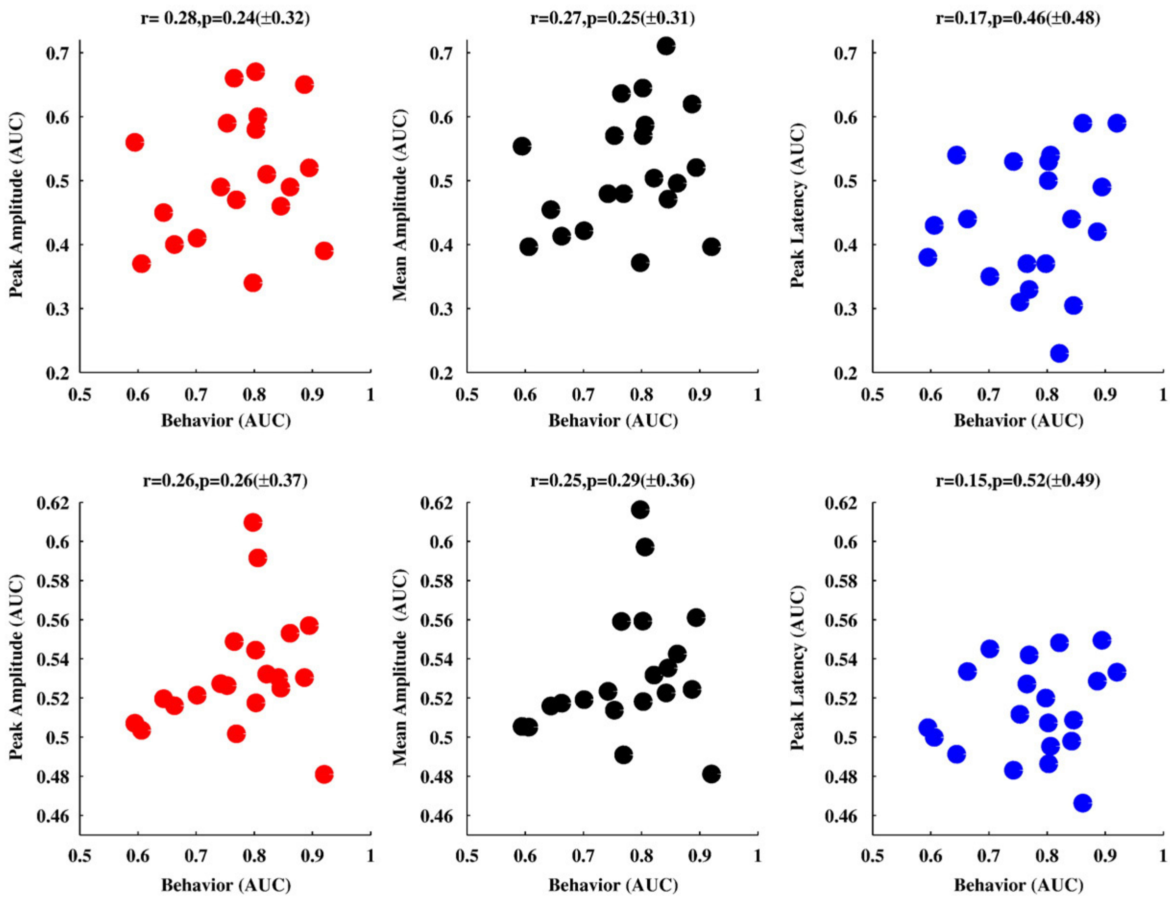

Fig. 5. Scatter plots showing correlation between behavioral performance (AUC) and AUC computed from ERP measures (peak N170 amplitude, mean N170 amplitude and peak N170 latency) using P08 electrode: top row shows AUC calculated using mean ERP waveforms for 10 sessions (100 trials/session). Bottom row shows AUC calculated using ERP measures for each individual trials.

Accuracy at rank ordering subjects behavioral performance from neural measures

To further evaluate the ability of the different EEG metrics to predict the variation in perceptual performance across individuals we performed a simple simulation calculating the accuracy of the different metrics to rank order the behavioral performance of two randomly sampled observers based on the neural activity (190 simulation trials, see Materials and methods, "Neural metric-based rank ordering of individual's behavioral performance" section, for details). Fig. 8 shows the mean percent correct rank ordering rank of two random observers for the four metrics: pattern classifier performance[1] (AUC), N170 mean ERP, N170 peak ERP and N170 latency. All ERP metrics were based on observers' difference ERP wave between face and car trials. Consistent with the correlation results, accuracy rank ordering subjects based on their behavioral performance was higher for the pattern classifier (82%) when compared to the ERP metrics (peak N170 amplitude 67.37%, mean N170 amplitude 65.79% and peak N170 latency 54.74%; $p<0.05$).

Table 2
The correlation coefficient between behavioral data and EEG measures using pattern classifiers under different training conditions for 20 observers.

Pattern classifiers using all electrodes	Metric	Training condition	r
CPCA	AUC	Tr Stim (all trials)	**0.6720**(±0.07)
		Tr Choice (all trials)	**0.6766**(±0.10)
		Tr Stim/Choice (correct trials)	**0.6794**(±0.09)
		Tr Stim (hi-conf trials)	**0.7049**(±0.09)
		Tr Choice (hi-conf trials)	**0.6552**(±0.08)
SVM	AUC	Tr Stim (all trials)	0.6371(±0.10)
		Tr Choice (all trials)	0.6245(±0.12)
		Tr Stim/Choice (correct trials)	0.6561(±0.13)
		Tr Stim (hi-conf trials)	0.6873(±0.11)
		Tr Choice (hi-conf trials)	0.64(±0.11)
LDA	AUC	Tr Stim (all trials)	0.6339(±0.10)
		Tr Choice (all trials)	0.6177(±0.12)
		Tr Stim/Choice (correct trials)	0.6660(±0.12)
		Tr Stim (hi-conf trials)	0.69(±0.15)
		Tr Choice (hi-conf trials)	0.6432(±0.12)

Temporal window analysis

We divided the EEG signals into time intervals of 40 ms starting at 200 ms preceding the stimulus presentation and extending to 480 ms post-stimulus. The pattern classifier was separately trained and tested on EEG data for each of the 40 ms temporal windows. Fig. 9A shows the correlation coefficient between the pattern classifier performance and the observers' behavioral performance for each 40 ms temporal window starting at 200 ms pre-stimulus onset. The correlation

[1] The pattern classifier investigated was trained on EEG activity for those trials in which observers expressed high confidence of their decision and assigned to the object category based on the identity of the image presented.

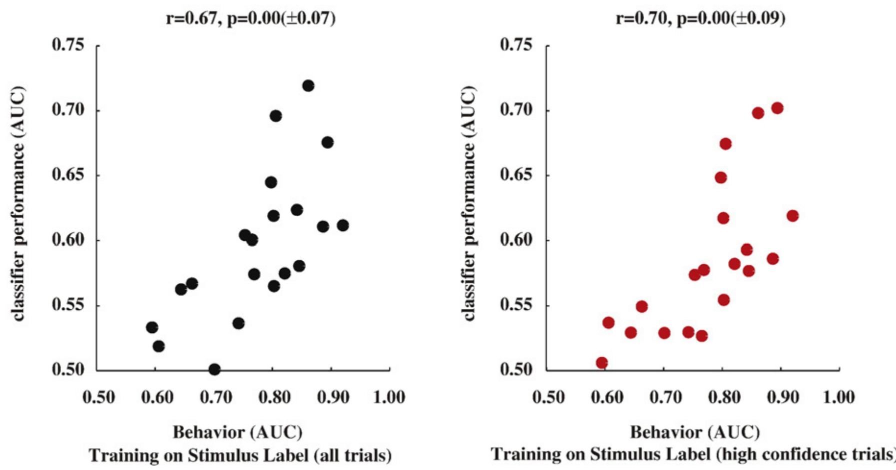

Fig. 6. EEG measures using pattern classifiers: a) AUC of pattern classifiers trained on all EEG trials using true labels of stimulus identity vs. the behavioral performance (AUC) of all 20 observers. b) AUC of pattern classifiers trained on only EEG trials with high-confidence user ratings using true labels of stimulus identity vs. the behavioral performance (AUC) of all 20 observers.

increases after stimulus onset and starts to asymptote at the temporal window starting at 120 ms. Fig. 9B presents the accuracy at rank ordering observers' behavioral performance based on the pattern classifier applied separately to each temporal window. The results are consistent with the correlation data, accuracy is at chance (50%) before stimulus onset and then increases steeply before asymptoting (or increasing at a slower rate) 120 ms after stimulus onset. For comparison, Fig. 9C shows the accuracy (AUC) of the classifier (averaged across observers and separately run for each temporal window) predicting whether a car or face stimulus was present. Although the pattern of results is similar to the correlation and rank ordering of observers' behavioral performance, there are some important differences. In particular, for the temporal window centered around 120 ms, the accuracy predicting the stimulus is close to chance but the correlation and rank ordering accuracy are higher.

Temporal analysis using both ERP difference wave and pattern classifiers are shown in Fig. 10. The top 2 rows of Fig. 10 shows the difference wave in each 40 ms time windows and the correlation between observers' behavioral performance and mean amplitude of the difference wave in the respective windows. The bottom row shows the discriminative filter weights (see Appendix) obtained from a pattern classifier for the same windows. Highly positive and highly negative filter weights (red and blue regions in Fig. 10, bottom row) denote spatial regions carrying discriminant information between faces and cars, as determined by the pattern classifier. From the figure, it follows that there is significant information differentiating faces and cars around 170 ms post-stimulus onset which shows up in the difference wave as well as in the pattern classifier filter weights. Statistically significant correlation exists between observers' behavioral performance and the mean amplitude of the difference wave from 170 ms onwards. Furthermore, the discriminative filter weights suggest that the right hemisphere contain greater discriminant information compared to the left hemisphere which is consistent with previous research (Bentin and Allison, 1996; Kanwisher et al., 1997).

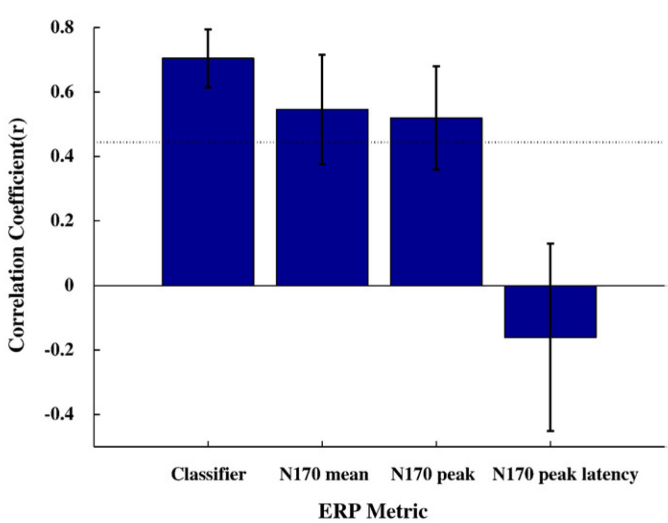

Fig. 7. Correlation between behavioral performance(AUC) and different EEG measures. The pattern classifier used was trained on EEG trials with high-confidence user ratings using true labels of stimulus identity, the three ERP measures used are for difference wave and P08 electrode. ($|r|>=0.444$ for statistically significant r using 95% confidence interval is marked with dotted line).

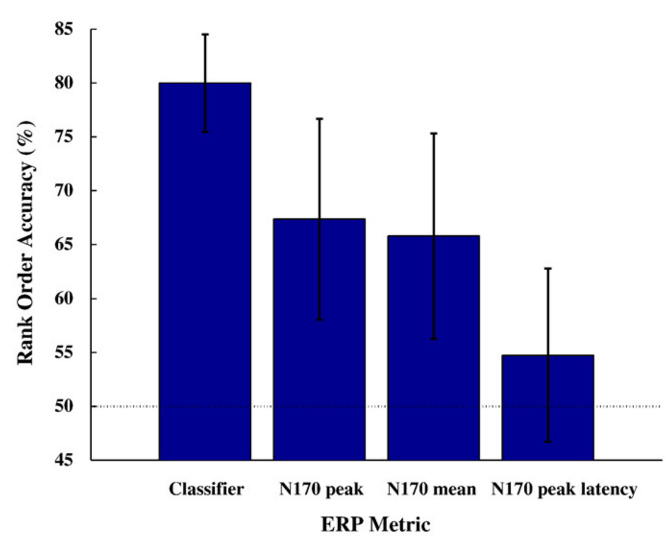

Fig. 8. Accuracy of rank ordering observers using different EEG measures, taking 2 observers at a time: The pattern classifier used was trained on EEG trials with high-confidence user ratings using true labels of stimulus identity, the three ERP measures used are for difference wave and P08 electrode (chance performance is marked with dotted lines).

Discussion

EEG-based metrics to predict variation in perceptual performance across individuals

A number of studies have demonstrated that pattern classifiers applied to EEG (single-trial EEG) can be used to predict the category of the image presented to the observers, the observers' choices and also task difficulty (Philiastides et al., 2006; Philiastides and Sajda, 2006). However, ability to reliably predict variation of perceptual performance across individuals will necessarily depend on how well the techniques predict observer choices and the variability in perceptual performance across individuals. If there is relatively small variability across individuals' performance then small inaccuracies in predicting observer choices will have strong detrimental effects in the ability to use neural activity to rank observers based on their behavioral performance. Here, we demonstrate that pattern classifiers applied to EEG data can be used to predict individual variations in perceptual performance in a categorization task (faces vs. cars). We compared three pattern classifiers to other more traditional ERP metrics including peak amplitude of the N170, mean amplitude of the N170 and peak latency of the N170. We used two methods to assess the ability of the metrics to predict human perceptual performance: a linear Pearson's correlation and accuracy in rank ordering behavioral performance based on EEG activity. Both evaluation methods showed that the pattern classifier was superior in their ability to predict observers' behavioral performance than the three ERP metrics evaluated (Figs. 7 and 8). However, the N170 peak and mean amplitudes calculated from the difference wave of the PO8 electrode correlated moderately with perceptual performance across individuals. This result is consistent with previous studies showing that the difference in ERP can capture aspects of display properties and human behavioral performance (Philiastides et al., 2006; Philiastides and Sajda, 2006; Vogel and Machizawa, 2004). We also calculated area under the curve using the three ERP metrics for a more direct comparison with the pattern classifiers and demonstrated that area under the curve using ERP metrics does not correlate significantly with the behavioral performance.

One lingering question is whether the correlation between the pattern classifier performance and behavior depends on the specific algorithm and training method used. To answer this question we investigated three pattern classifiers including linear pattern classifier (LDA), Support Vector Machines (SVM) and a non-linear classifier (CPCA) and five different training methods (see Table 2). The enhanced performance of CPCA over the remaining classifiers can probably be attributed to the non-linear nature of feature extraction process of the classifier. Although the CPCA algorithm trained on high-confidence trials resulted in the highest correlation with perceptual performance, the difference across algorithms and training methods were small and suggest a robustness of results across pattern classifier algorithms and training methods.[2]

Spatial distribution of electrodes predicting variation in behavioral performance across individuals

Although the interpretation of the spatial distribution of the correlations between behavior and ERP measures is constrained by the limited spatial resolution of EEG and the inverse problem, there are two aspects of the present results that converge with the existing literature on the face-selective N170 ERP component. First, the finding that electrodes over lateral occipital and temporal cortex of both hemispheres show significant positive correlations between performance and amplitude (both peak and mean, Figs. 4 and 10) is consistent with functional neuroimaging (Kanwisher et al., 1997) and intracranial recording (Allison et al., 1999; McCarthy et al., 1999; Puce et al., 1999) studies demonstrating that regions of the anterior fusiform gyrus in both hemispheres exhibit more robust responses to faces compared to other categories of stimuli. Similarly, the present finding that these correlations may be stronger in the right hemisphere than in the left (Fig. 4) also parallels classic fMRI (Kanwisher et al., 1997) and ERP studies (Bentin and Allison, 1996) demonstrating that while the face-selective response can be measured in both hemispheres, the amplitude of the response is larger in the right occipital and temporal cortex. Together these converging lines of evidence suggest that the face-selective response carries key information that aids discrimination across individuals.

Second, the robust positive correlations between amplitude and behavior at frontal electrode sites (Fig. 10) both during and after the typical N170 time window are likely indicative of the involvement of dorsolateral prefrontal cortex in perceptual decision making and categorization (Freedman et al., 2001; Kim and Shadlen, 1999; Li et al., 2009,2007; for reviews see Ashby and Maddox, 2005; Heekeren et al., 2008). This interpretation predicts that the faster these systems are engaged should be predictive of better performance. Consistent with this interpretation, we found that during the N170 time window, onset latency is negatively correlated with performance at frontal electrodes (Fig. 4) such that those individuals who had shorter onset latencies had higher performance. The scalp topography of the behavioral correlations and the classifier weights are consistent with the known regional specificity of face processing and decision making, as well as with recent EEG-informed fMRI results using a similar paradigm (Philiastides and Sajda, 2007).

Neural time-epochs predicting variation in behavioral performance across individuals

Previous studies have identified EEG waveforms at specific time-epochs (N170, N200) associated with the presentation of faces when compared to other non-face objects (Halgren et al., 2000; Jeffreys, 1996; Liu et al., 2000; Rossion et al., 2003). In addition recent applications of pattern classifiers to EEG data, (Philiastides et al., 2006; Philiastides and Sajda, 2006) have characterized the N170 early component as being related to perceptual processes and distinct from a later component (330–400 ms) linked to decision processes and a component at 220 ms (D220) related to task difficulty. Finally, a study (Liu et al., 2002) using magnetoencephalography (MEG) has found an early trial-averaged activity (M100) associated with the categorization of face. Unlike previous studies, we did not find that the neural activity predicting the performance across individuals was narrowly concentrated in particular temporal epochs. Instead, the neural activity predicting observers' behavioral performance was distributed through time (see Fig. 9) and increased monotonically starting at 120 ms after stimulus onset. Arguably, this result is reasonable given that we might expect that observers that perform at higher performance level might show N170 perceptual components that more strongly discriminate face from car stimuli, different D220 components reflecting the varying task difficulty experienced by each observer and also a different amount of information in the late decision component. Note that if the classifier could perfectly predict the observer trial-to-trial choices from the EEG activity in the late component then the classifiers' performance would perfectly correlate with observers' decisions. In this context, it is not surprising that the highest accuracy predicting observers' performance (and correlation) rests in the later times associated with EEG decision component.

[2] The pattern classifiers chosen in the study jointly extract neural information in spatiotemporal domain and do not require expensive parameter optimization. Use of data-driven spatial processing approaches like common spatial pattern (CSP) (Muller-Gerking et al., 1999; Ramoser et al., 2000), resulted in near chance performance; however optimization of parameters such as frequency bands, time window was not explored.

Conclusion

Within the past several years pattern classification techniques applied to measures of brain activity have provided key insights into the neuronal signatures of information processing on a trial-by-trial basis (Haynes and Rees, 2005; Kamitani and Tong, 2005; Norman et al., 2006; Philiastides and Sajda, 2006). The present results demonstrate that pattern classification techniques are a powerful tool for using the spatiotemporal dynamics of patterns in neural data to predict individual differences in behavioral performance. We expand on previous work by systemically comparing different pattern classifier algorithms and training methods to various standard ERP techniques to demonstrate that irrespective of algorithm and training method, pattern classifiers outperform traditional ERP measures in their ability to predict human perceptual performance. Although the precise timing and spatial distribution of the predictive efficacy is likely to change depending on the type of stimulus and task demands (Philiastides and Sajda, 2006), these findings converge with numerous EEG and fMRI studies that have used pattern classification techniques to provide insights into the neuronal signatures of human information processing.

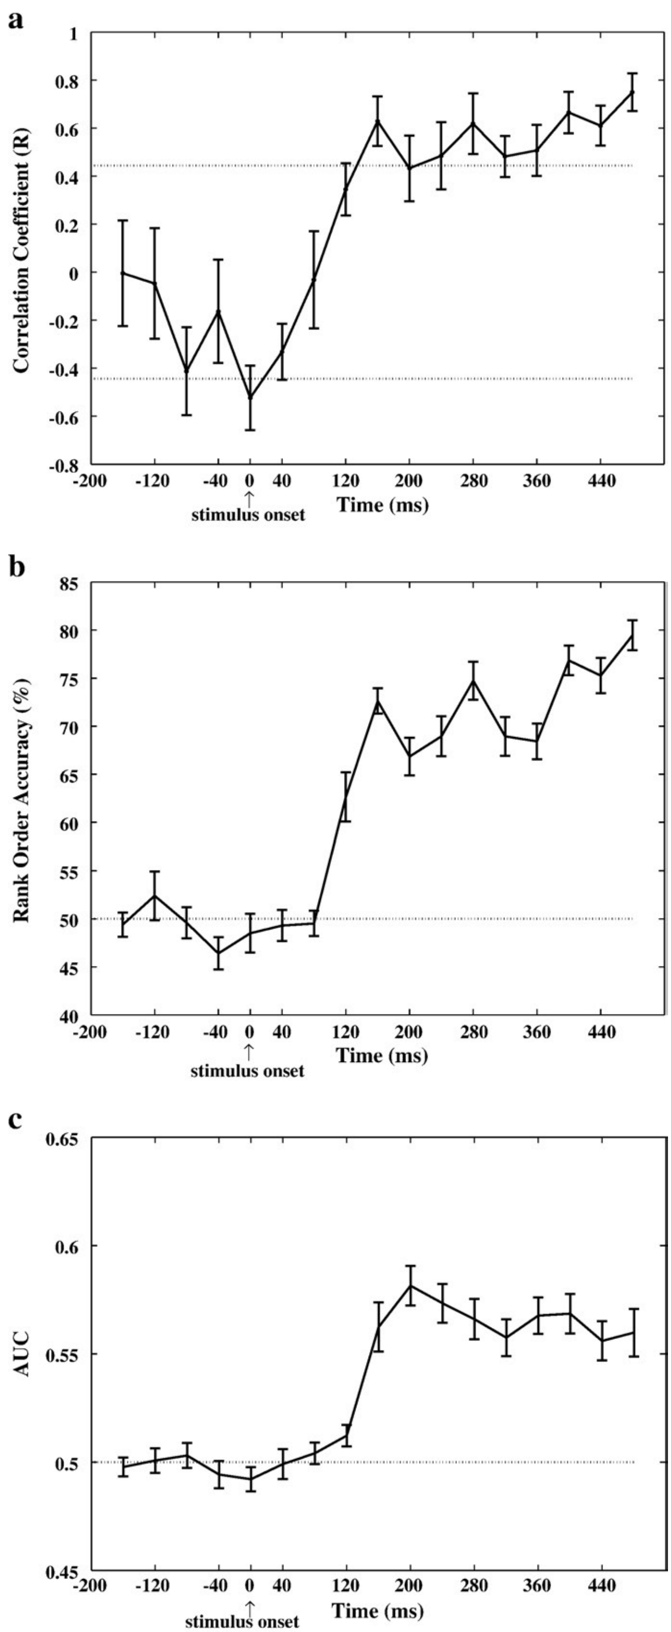

Fig. 9. EEG analysis with pattern classifiers using time intervals starting 200 ms before the stimulus appears to 480 ms post-stimulus onset. a) Correlation between AUC of pattern classifiers of 20 observers taken over time intervals and their behavioral performance (AUC) ($|r| >= 0.444$ for statistically significant R using 95% confidence interval is marked with dotted line). b) Accuracy of rank ordering observers using pattern classifiers taken over time intervals (chance performance is marked with dotted lines). EEG analysis with pattern classifiers using time intervals starting 200 ms before the stimulus appears to 480 ms post-stimulus onset. c) AUC using pattern classifiers averaged across 20 observers predicting the true stimuli, cars/faces, as a function of temporal window (chance performance is marked with dotted lines).

Acknowledgments

Funding for this project was graciously provided by Army grant W911NF-09-D-0001.

Appendix A. Classwise Principal Component Analysis (CPCA)

The details of the algorithm will be given for $c = 2$, and the extension to an arbitrary number of classes is straightforward.

Feature extraction

Let $\omega_i (i = 1,2)$ denote two classes with means μ_i, and covariances Σ_i, and let $x^* \in \mathbb{R}^n$ be unknown (test) data to be classified. In the first step, x^* is represented in 2 subspaces, S_1 and S_2 (see Fig. 11), by means of the following transformation

$$x_i^* = F_i^T (x^* - \mu_i) \quad i = 1, 2 \tag{A.1}$$

where the columns of $F_i \in \mathbb{R}^{n \times m_i'}$, are taken as the basis vectors of S_i. The two classes are transformed in a similar fashion (Figs. 11B,C). In the simplest scenario, $F_i = V_i$, where $V_i \in \mathbb{R}^{n \times m_i'}$ consists of the m_i (m_i to be chosen) principal components of the class ω_i. To account for classes whose principal directions are nearly parallel, and hence the projections of the two classes to S_i are highly overlapped, F_i is augmented with $V_b \in \mathbb{R}^{n \times 1}$, where $V_b \propto \mu_1 - \mu_2$. This step ensures that class differences arising from the two means are accounted for. For c-class cases, V_b readily generalizes to a basis spanned by the columns of the between-class-scatter matrix, commonly used in LDA applications (Duda et al., 2001). To keep all projections orthogonal, the columns of $F_i = [V_i | V_b]$ are orthonormalized through the Gram–Schmidt procedure.

While the above procedure typically yields S_i of sufficiently low dimension ($m_i' \ll n$), where the size of data is no longer an obstacle, further improvements in terms of classification accuracy are possible with simple feature extraction techniques applied directly to the subspace S_i. If linear feature extraction techniques are used (e.g. LDA), the mathematical formalism (Eq. A.1) remains the same, with mere modifications in the definition of F_i. More specifically, $F_i = [V_i | V_b] T_i$, where $T_i \in \mathbb{R}^{m_i' \times m}$ is the feature extraction matrix of the chosen

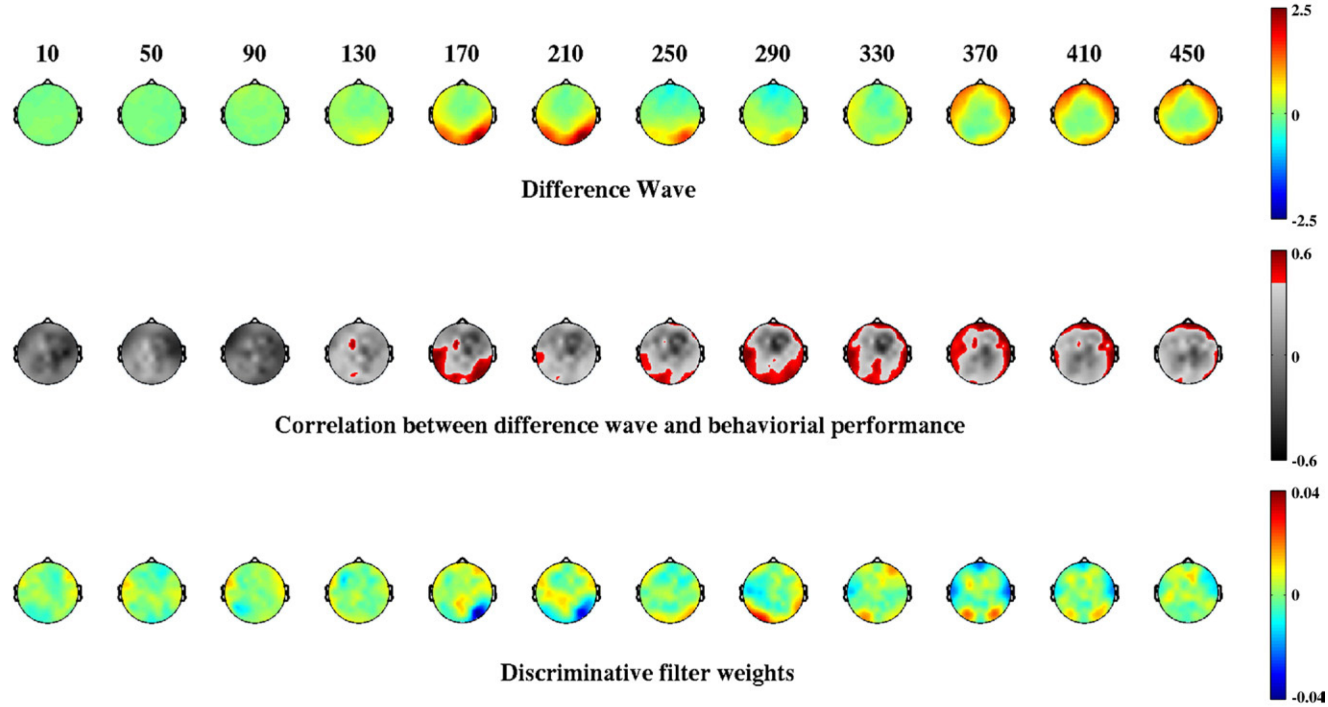

Fig. 10. Temporal analysis using ERP difference wave and filter weights from pattern classifiers: top row shows the difference wave across 40 ms time windows, followed by correlation between observers' behavioral performances (AUC) and difference wave using mean amplitude across time windows (center row). Only areas showing significant correlation (correlation coefficient $|r>0.444|$) are marked in red. The bottom row shows discriminative filter weights (see Appendix) obtained from pattern classifiers for the corresponding temporal windows. Filter coefficients with large absolute value (red and blue) denote spatial regions carrying information discriminating between face and car.

method. An information-theoretic technique called Approximate Information Discriminant Analysis (AIDA) is used, whose advantages over LDA and similar techniques have been discussed at length in Das and Nenadic (2008). Unlike LDA, AIDA has no constraints regarding the final dimension, m, of the feature space. However, EEG data is generally so sparse (small n_t) that the choice of m is severely limited.

Classification

Due to piecewise linear nature of the feature extraction method, the test data x^* is represented in 2 feature subspaces (Figs. 11B,C). To complete the feature extraction process, one of the subspaces must be eliminated. It turns out that this question can be solved within a classification framework, which is the ultimate goal of the technique. Therefore, the formal completion of the feature extraction process can be viewed as a bi-product of the classification process.

For simplicity, we will assume that the classes are Gaussian with prior probabilities, $P(\omega_i)$. A straightforward application of the Bayes classifier at the first subspace yields

$$P(\omega_{i1}|x_1^*) = \frac{p(x_1^*|\omega_{i1})P(\omega_i)}{p(x_1^*)} \quad i = 1,2 \tag{A.2}$$

where $p(x_1^*) = \sum_{i=1}^{2} p(x_1^*|\omega_{i1})P(\omega_i)$, and $P(\omega_{i1}|x_1^*)$ are the posterior probabilities of the two classes in the first subspace. Since (Eq. A.1) is an affine transformation, both classes remain Gaussian, i.e. $p(.|\omega_{i1}) \sim N(F_1^T(\mu_i - \mu_1), F_1^T \Sigma_i F_1)$.

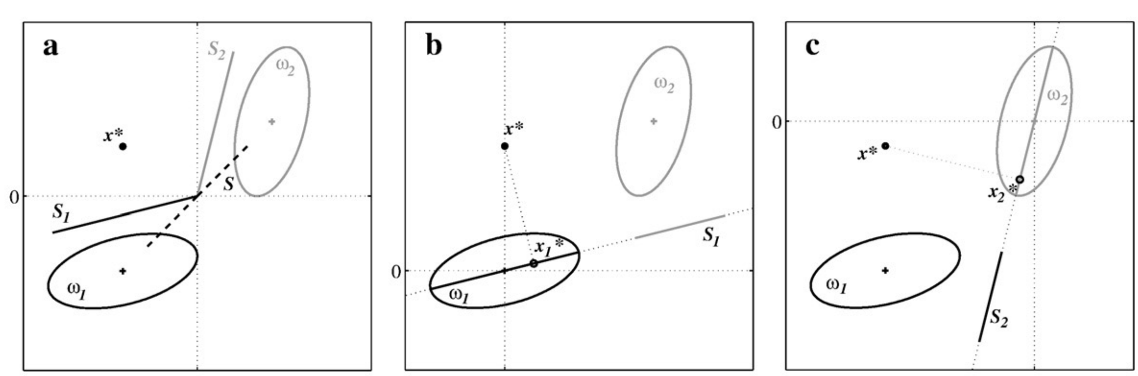

Fig. 11. (a) PCA (dashed) vs. CPCA subspace for 2-class case, where the classes ω_1 and ω_2 are represented as Gaussian contours. (b) Test data x projected on reduced subspace S_1 to produce test feature x_1 (C) x^* projected on S_2 to produce x_2^*.

Similarly, x^* is transformed to the second subspace (see Figs. 11C), and posterior probabilities of the two classes in the second subspace is calculated as

$$P(\omega_{i2}|x_2^*) = \frac{p(x_2^*|\omega_i 2)P(\omega_i)}{p(x_2^*)} \quad i = 1, 2 \tag{A.3}$$

and $p(x_2^*)$ is defined analogous to $p(x_1^*)$. Also note that $p(.|\omega_{i2}) \sim \mathcal{N}(F_2^T(\mu_i - \mu_2), F_2^T \Sigma_i F_2)$. Once the posterior probabilities for each subspace were determined, the scalar value (λ) to be used in Area under the Receiver Operating Curve (ROC), AUC, calculation was computed for each trial by taking the difference between the maximum posterior of each class.

$$\lambda = \arg \max_{i=1,2} P(\omega_{1i}|x_1^*) - \arg \max_{i=1,2} P(\omega_{2i}|x_i^*) \tag{A.4}$$

Discriminative Filter Weights

While features arising in CPCA are abstract, the corresponding filters (the columns of $\mathbf{F}_i$) have a clear physical interpretation. If the filter coefficients are organized into a spatiotemporal array, they will explicitly point to brain areas and time scales involved in encoding the differences between the two classes (face/car). Examples of such filters for specific time points are provided in Fig. 10.

References

Allison, T., Puce, A., McCarthy, G., 1999. Electrophysiological studies of human face perception. i: Potentials generated in occipitotemporal cortex by face and non-face stimuli. Cereb. Cortex 9, 415–430.

Ashby, F., Maddox, W., 2005. Human category learning. Ann Rev. Psych. 56, 149–178.

Bentin, S., Allison, T., 1996. Electrophysiological studies of face perception in humans. J. Cogn. Neurosci. 8 (6), 551–565.

Berger, H., 1929. Ãœber das elektroencephalogram des menschen. Arch Psychiatr Nervenkr 87, 527–570.

Das, K., Nenadic, Z., 2008. Approximate information discriminant analysis: a computationally simple heteroscedastic feature extraction technique. Pattern Recogn. 41, 1548–1557.

Das, K., Nenadic, Z., 2009. An efficient discriminant-based solution for small sample size problem. Pattern Recogn. 42, 857–866.

Delorme, A., Makeig, S., 2004. Eeglab: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Methods 134, 9–21.

Duda, R.O., Hart, P.E., Stork, D.G., 2001. Pattern Classification. Wiley-Interscience. .

Ergenoglu, T., Demiralp, T., Bayaraktaroglu, Z., Ergen, M., Beydagi, H., Uresin, Y., 2004. Alpha rhythm of the EEG modulates visual detection performance in humans. Cogn. Brain Res. 20, 376–383.

Ertl, J.P., Schafer, E.W.P., 1969. Brain response correlates of psychometric intelligence. Nature 223, 421–422.

Fisher, R.A., 1936. The use of multiple measurements in taxonomic problems. Ann. Eugen. 7, 179–188.

Freedman, D., Riesenhuber, M., Poggio, T., Miller, E.K., 2001. Categorical representation of visual stimuli in the primate prefrontal cortex. Science 291 (5502), 312–316.

Fukunaga, K., 1990. Inroduction to Statistical Pattern Recognition. Academic Press.

Gauthier, I., Curran, T., Curby, K.M., Collins, D., 2003. Perceptual interference supports a non-modular account of face processing. Nature 6, 428–432.

Gerson, A.D., Parra, L.C., Sajda, P., 2005. Cortical origins of response time variability during rapid discrimination of visual objects. Neuroimage 28 (2), 342–353.

Halgren, E., Raij, T., Marinkovic, K., Jousmaki, V., Hari, R., 2000. Cognitive response profile of the human fusiform face area as determined by MEG. Cereb. Cortex 10, 69–81.

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C.S., Bäuml, K.-H., 2007. Prestimulus oscillations predict visual perception performance between and within subjects. NeuroImage 37, 1465–1473.

Haynes, J.D., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat. Neurosci. 8 (2), 686–691.

Heekeren, H.R., Marrett, S., Ungerleider, L., 2008. The neural systems that mediate human perceptual decision making. Nat Rev. Neurosci. 9, 467–479.

Jeffreys, D., 1996. Evoked studies of face and object processing. Vis Cogn. 3, 1–38.

Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nat. Neurosci. 8, 679–685.

Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J. Neurosci. 17 (11), 4302–4311.

Kim, J.-N., Shadlen, M.N., 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. 2, 176–185.

Klimesch, W., Schimke, H., Pfurtscheller, G., 1993. Alpha frequency, cognitive load, and memory performance. Brain Topogr. 5, 241–251.

Kohavi, R., 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection: Int. Joint C. Art. Int. , pp. 1137–1145.

Kutas, M., McCarthy, G., Donchin, E., 1977. Augmenting mental chronometry: the p300 as a measure of stimulus evaluation time. Science 197, 792–795.

Li, S., Mathew, S., Kourtzi, Z., 2009. Learning shapes the representation of behavioral choice in the human brain. Neuron 62, 441–452.

Li, S., Ostwald, D., Giese, M., Kourtzi, Z., 2007. Flexible coding for categorical decisions in the human brain. J. Neurosci. 27 (45), 12321–12330.

Liu, J., Harris, A., Kanwisher, N., 2002. Stages of processing in face perception: an MEG study. Nat. Neurosci. 5 (9), 910–916.

Liu, J., Higuchi, M., Marantz, A., Kanwisher, N., 2000. The selectivity of the occipitotemporal m170 for faces. NeuroReport 11, 337–341.

Mangun, G.R., Hillyard, S.A., 1988. Spatial gradients of visual attention: behavioral and electrophysiological evidence. Hum. Factors 29, 195–212.

Mangun, G.R., Hillyard, S.A., 1990. Allocation of visual attention to spatial locations: tradeoff functions for event-related brain potentials and detection performance. Percept. Psychophys. 47, 532–550.

McCarthy, G., Puce, A., Allison, T., 1999. Electrophysiological studies of human face perception. ii: Response properties of face-specific potentials generated in occipitotemporal cortex. Cereb. Cortex 9, 431–444.

Muller-Gerking, J., Pfurtscheller, G., Flyvbjerg, H., 1999. Designing optimal spatial filters for single-trial EEG classification in a movement task. Clin. Neurophysiol. 110 (5), 787–798.

Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cognit. Sci. 10 (9), 424–430.

Parra, L.C., Spence, C.D., Gerson, A.D., Sajda, P., 2005. Recipes for the linear analysis of EEG. NeuroImage 28 (2), 326–341.

Philiastides, M.G., Ratcliff, R., Sajda, P., 2006. Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. J. Neurosci. 26 (35), 8965–8975.

Philiastides, M.G., Sajda, P., 2006. Temporal characterization of the neural correlates of perceptual decision making in the human brain. Cereb. Cortex 16 (4), 509–518.

Philiastides, M.G., Sajda, P., 2007. EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. J. Neurosci. 27 (48), 13082–13091.

Puce, A., Allison, T., McCarthy, G., 1999. Electrophysiological studies of human face perception. iii: Effects of top-down processing on face-specific potentials. Cereb. Cortex 9, 445–458.

Ramoser, H., Muller-Gerking, J., Pfurtscheller, G., 2000. Optimal spatial filtering of single trial EEG during imagined hand movement. IEEE T. Neur. Sys. Reh. 8 (4), 441–446 dec.

Rossion, B., Joyce, C.A., Cottrell, G.W., Tarr, M.J., 2003. Early lateralization and orientation tuning for face, word, object processing in the visual cortex. Neuroimage 20, 1609–1624.

Taylor, M.J., McCarthy, G., Saliba, E., Degiovanni, E., 1999. ERP evidence of developmental changes in processing of faces. Clin. Neurophysiol. 110 (5), 910–915.

Troje, N., Bülthoff, H.H., 1996. Face recognition under varying poses: the role of texture and shape. Vis. Res. 36, 1761–1771.

Vapnik, V., 1998. Statistical Learning Theory. . Wiley.

Vogel, E., Machizawa, M.G., 2004. Neural activity predicts individual differences in visual working memory capacity. Nature 428, 748–751.

Vogel, E.K., McCollough, A.W., Machizawa, M.G., 2005. Neural measures reveal individual differences in controlling access to working memory. Nature 438, 500–503.

Walter, W.G., Cooper, R., Aldridge, V.J., McCallum, W.C., Winter, A.L., 1964. Contingent negative variation: an electric sign of sensori-motor association and expectancy in the human brain. Nature 203, 380–384.

Woodworth, R., 1938. Experimental Psychology. Holt, New York.