ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 32 (2006) 465 – 476

Single-trial discrimination of truthful from deceptive responses during a game of financial risk using alpha-band MEG signals

Anil K. Seth,* John R. Iversen, and Gerald M. Edelman

The Neurosciences Institute, San Diego, CA 92121, USA

Received 22 December 2005; revised 15 February 2006; accepted 28 February 2006
Available online 4 May 2006

We studied whether magnetoencephalography (MEG) could detect deceptive responses on a single-subject, trial-by-trial basis. To elicit spontaneous, ecologically valid deception, we developed a paradigm in which subjects in a simulated customs setting were presented with a series of pictures of items which might be in their baggage, and for each item, they decided whether to 'declare' (tell the truth) or 'smuggle' (lie). Telling the truth involved a small but certain monetary penalty, whereas lying involved both greater monetary risk and greater potential reward. Most subjects showed decreased signal power in the 8–12 Hz (alpha) range during deceptive responses as compared to truthful responses. In a cross-validation analysis, we were able to use alpha power to classify truthful and deceptive responses on a trial-by-trial basis, with significantly greater predictive accuracy than that achieved using simultaneously recorded skin conductance signals. Average predictive accuracy for spontaneous deception was greater than 78%, and for some subjects, predictive accuracy exceeded 90%. Our results raise the possibility that alpha power modulation during deception may reflect risk management and/or cognitive control.

Introduction

Deception involves the intentional concealment, distortion or fabrication of information for the purpose of gaining an advantage or leading another into error. The most widely used method for deception detection remains the polygraph (Office of Technology Assessment, 1990) which measures peripheral physiological signals including skin conductance responses (SCRs), heart rate and respiration. Because these signals do not directly reflect cognitive or neural processes, there has been growing interest in the use of brain signals for deception detection. Scalp-recorded event-related potentials (ERPs) have been used to detect concealed information in the 'guilty knowledge test' [GKT (Lykken, 1959, 1991)], in which subjects are instructed to conceal a knowledge item from the investigator (Allen et al., 1992; Farwell and Donchin, 1991; Rosenfeld et al., 1988). The utility of this approach may depend on the well-known P300 'oddball' response of ERPs [enhanced amplitude ~300 ms following rarely presented and meaningful stimuli (Donchin and Coles, 1988), see Rosenfeld, 2001 for a review]. More recently, researchers have turned to functional magnetic resonance imaging (fMRI) in order to gain insight into the spatial locations of neural mechanisms involved in deception (Davatzikos et al., 2005; Kozel et al., 2004a,b; Langleben et al., 2002, 2005; Lee et al., 2002; Nunez et al., 2005; Spence et al., 2001). These studies have implicated a variety of frontal, temporal and parietal areas in the production of deceptive responses [see Spence et al., 2004 for a review]. In the present study, we looked for modulation of magnetoencephalographic (MEG) signals by deception. In comparison to fMRI, MEG offers high temporal resolution. In contrast to ERP studies, we focused on ongoing (i.e., oscillatory) components of brain signals. Although the results of our study have several implications for deception detection 'in the field', we should emphasize that it was not our intention here to develop a practical 'lie detector'.

An important challenge for experimental studies of deception is to arrange for subjects to lie spontaneously (i.e., without being explicitly instructed to lie), in a controlled experimental situation. Previous ERP and fMRI studies have included elements of both spontaneous deception and instructed deception. Among ERP studies, both Rosenfeld et al. (1999) and Johnson et al. (2005) instructed their subjects to lie on about 50% of trials; subjects decided spontaneously, for each trial, whether to lie or not. Among fMRI studies, Ganis et al. (2003) allowed subjects to decide spontaneously the content of a lie (they were still instructed when to lie). In the present study, we have developed a paradigm in which subjects decide for themselves whether, and when, to lie. In our paradigm, subjects were presented with a series of pictures of items varying in monetary value. For each item, they decided whether to 'declare' (tell the truth) or 'smuggle' (lie), in a simulated customs setting. Telling the truth involved a small but certain financial penalty, whereas lying involved both greater financial risk and greater potential reward. Although subjects were

* Corresponding author. Fax: +1 858 626 2099.
E-mail addresses: seth@nsi.edu (A.K. Seth), iversen@nsi.edu (J.R. Iversen), edelman@nsi.edu (G.M. Edelman).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2006.02.050

instructed to lie or tell the truth for certain items, they were given no such instructions for other items. An advantage of this design is that spontaneous deception can be contrasted with instructed deception within the same experimental session.

We measured several aspects of the MEG signal, while subjects engaged in the above paradigm. In addition to power in the standard frequency ranges ('theta', 4–8 Hz; 'alpha', 8–12 Hz; 'beta', 15–30 Hz; and 'gamma', 30–50 Hz) (Nunez, 1995), we measured the brain response to an amplitude-modulated auditory tone (a 'frequency tag', see Materials and methods) as a 'probe' into cortical dynamics (Patel and Balaban, 2000; Silberstein, 1995). Of these various signals, we found that power in the alpha range was the most useful for detection of deceptive responses. For most subjects, mean alpha power across the MEG sensor array was higher during truthful responses than during deceptive responses. We did not observe consistent task-dependent modulation of any other frequency range nor of the tag response.

Modulation of alpha oscillations has been associated with cortical 'idling'(Pfurtscheller et al., 1996), with working memory load (Jensen et al., 2002), with attentional shifts (Yamagishi et al., 2003) and, recently, with risk (Oya et al., 2005). In our study, deceptive responding incurred increased financial risk as compared to truthful responding but did not differentially task working memory or attention in any obvious way. As with most deception studies, deceptive responding in our paradigm may also involve increased 'cognitive control', for example, conflict monitoring (Ruff et al., 2001) and inhibition of competing responses (Braver et al., 2001). Our results therefore raise the possibility that alpha power modulation during deception may reflect the operation of neural mechanisms mediating risk management and/or cognitive control.

A second important challenge for experimental studies of deception is the discrimination of truthful responses from deceptive responses at the level of single subjects and single trials, rather than after averaging data across multiple subjects and/or multiple trials. This challenge is significant for assessing rigorously the robustness of any empirically observed difference between truthful responding and deceptive responding, as well as for potential practical applications. Whereas a skin conductance-based GKT procedure produces data that can be assessed on a trial-by-trial basis, ERP data are less easily analyzed this way and typically require averaging across multiple trials [see Allen, 2002; Rosenfeld, 2001 for reviews]. Recently, multivariate fMRI data have been successfully analyzed on a trial-by-trial basis (Davatzikos et al., 2005; Langleben et al., 2005). Apart from an exploratory study by Thornton (1995), to our knowledge, there are no robust reports of trial-by-trial classification based on ongoing MEG or EEG signals. Moreover, no previous studies have compared the classification accuracy of brain signals with that achievable by physiological measures such as the SCR, within the same experimental session.

In the present study, we used logistic regression to classify, on a single-subject, trial-by-trial basis, MEG alpha-band power as reflecting either deceptive responses or truthful responses. Logistic regression is well suited to this task because it is able to regress a continuous 'predictor' variable (alpha power) against a binary 'response' variable (truth or lie). To provide a robust assessment of the predictive accuracy of MEG alpha under logistic regression, we applied N-fold cross-validation. Data from each subject were divided into N subsets, and a logistic classifier was trained N times, each time leaving out one of the subsets from training, but using only the omitted subset to test the accuracy of the classifier ($N=5$ in the present case). We emphasize the use of cross-validation since classification performance on data used to train a classifier may not accurately reflect its ability to generalize to novel data (Browne, 2000; Hastie et al., 2001).

Using the above cross-validation analysis, we compared the predictive accuracy of a logistic classifier based on MEG alpha-band power, with that achieved by an equivalent classifier based on simultaneously recorded SCR signals, which are known to perform well in the GKT (Ben-Shakhar and Elaad, 2003) and which comprise part of the polygraph method. For all subjects, classification using MEG outperformed classification using SCR signals.

Materials and methods

Participants

All twenty subjects that participated in the experiment (14 males and 6 females, aged 19–39) had normal or corrected-to-normal vision. The experimental protocol was approved by the Institutional Review Board, and all subjects gave written informed consent.

Deception task

The present design combines three desirable features for a deception experiment (see Table 1): (i) the potential to contrast spontaneous deception with instructed deception, (ii) a comparatively even balance between deceptive responses and truthful responses and (iii) simplicity. Subjects were presented with a series of pictures of items varying in monetary value, and for each item, they decided whether to 'declare' (tell the truth) or 'smuggle' (lie). The goal for each subject was to retain as much money as possible out of an initial allocation of $100. There were four items: a suitcase, a bottle of whisky, a necklace and a gun (Fig. 1). The suitcases had no value whether declared or smuggled. Declaration of whisky resulted in a 'duty' of $0.70, whereas failed smuggling resulted in a 'fine' of $4. The duty for a necklace was $2 and the corresponding fine was $4.50. Guns were always to be smuggled: the penalty for accidentally declaring a gun ($10) was higher than corresponding fine ($7). For each subject, a pseudorandom sequence of objects was composed, consisting of 21 suitcases, 18 whisky bottles, 18 necklaces and 13 guns.

The outline of an experimental trial is shown in Fig. 1. Each trial began with a visual fixation point, followed by visual presentation of an object from the pseudorandom sequence. After a delay period (the 'pre' period), the visual prompt "Anything to

Table 1
Three dimensions along which deception experimental designs can be compared

Design	Complexity of design	Frequency of deception	Spontaneity of deception
'Smuggle'	Low	~50% of trials	~50% of trials
GKT	Low	Low	Low

In contrast to the widely utilized 'guilty knowledge task' (GKT), the present 'smuggle' design encourages an even balance between deceptive responses and truthful responses and allows spontaneous deception and instructed deception to be contrasted within the same experimental session.

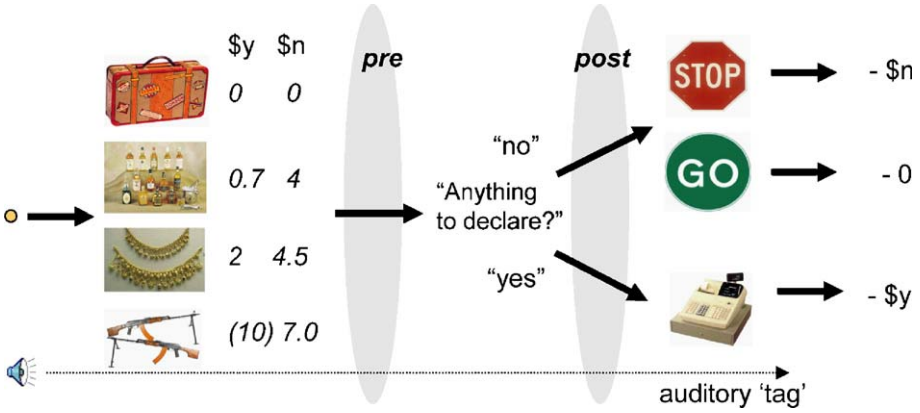

Fig. 1. Each trial began with a visual fixation point (2.5 s), followed by visual presentation of 1 of 4 objects, drawn from a pseudorandom sequence (2 s; gun $P = 0.18$, necklace and whisky $P = 0.26$, suitcase $P = 0.30$). There followed a delay period ('pre' period; 5 s). Upon appearance of the prompt "Anything to declare?" (2 s), the subject made a verbal response ('yes' or 'no'), which was noted by the experimenter and entered into the computer that was running the experiment. There followed another delay ('post' period, 5 s), after which, if the subject had responded 'yes', a cash register appeared (2 s), signifying a 'duty' corresponding to the amount in the column 'y' for the declared object. If the subject had responded 'no', then *either* a red 'stop' sign appeared (2 s; $P = 0.66$) signifying a fine corresponding to the amount in column 'n', *or* a green 'go' sign (2 s; $P = 0.33$) appeared, signifying no financial loss. Visual feedback was given of the amount of money lost during the trial, as well as of the total amount remaining (2 s). During each block of 7 trials, subjects heard an auditory 'tag' (an 800 Hz pure tone, amplitude modulated at 45 Hz). Each trial lasted ≈20.5 s.

declare?" appeared, to which the subject gave a verbal response ('yes' or 'no'). There followed another delay period (the 'post' period), after which there were two possibilities. If the subject had responded 'yes' (declared), he or she had to pay the corresponding 'duty' (if any). If the subject had responded 'no', then either a red 'stop' sign ($P = 0.33$) or a green 'go' sign ($P = 0.67$) appeared. In the former case, the subject had to pay the corresponding fine (if any); in the latter case, there was no financial loss. At the end of each trial, the subject was given visual feedback of the amount of money lost during the trial and of the total amount of money remaining. Each trial lasted ≈20.5 s, and the experiment consisted of 10 blocks of 7 trials each. Subjects were allowed short rest periods between each block and were given one 'practice' block in order to familiarize themselves with the procedure.

Instruction protocol

Subjects were given $100 and informed that they would be able to keep any money that remained at the end of the experiment. They were told that they would see pictures of suitcases, whisky, necklaces and guns, and that they must either declare these objects or deny having them, as if they were passing through a customs checkpoint. They were instructed always to declare suitcases and always to deny having guns. For necklaces and whisky, subjects were informed that it was "up to them" whether to lie (smuggle) or tell the truth (declare). Trials involving guns or suitcases therefore involved instructed deception, whereas trials involving whisky or necklaces involved spontaneous deception. Subjects were instructed to respond verbally at the appropriate point in each trial and to make their responses as quickly as possible, but without hurrying.

Although subjects were informed of the existence of duties (for declaration) and fines (for smuggling) for each object type, they were not informed of the magnitudes of these duties and fines nor of the likelihood with which a smuggling attempt would be challenged. Critically, subjects were not informed of the optimal strategy, which was to always smuggle guns and necklaces, and to always declare suitcases and whisky. They were told only that they would be likely to lose money quickly if they lied all the time, or if they told the truth all the time. If, at the end of the experiment, a subject had a balance of less than $20, they were given $20 in lieu of their remaining balance (if any).

Data acquisition

Neuromagnetic signals were collected using a Magnes 2500 Wh MEG system from 4D Neuroimaging (San Diego). The MEG sensors consisted of 148 magnetometer coils (1-cm diameter) covering the whole head with 3-cm spacing. MEG recordings were collected in a magnetically shielded room. Stimuli were generated outside the MEG recording room using a desktop PC and were projected [using a Proxima (San Diego) 2000 projector] onto a screen in front of the subject via a porthole and mirror. When projected, visual stimuli occupied a square field with a visual angle of ≈13°. Several photodiodes were placed on the computer screen and recorded in real time the occurrence of the 'pre' and 'post' periods in each trial.

During each block, subjects were exposed to a continuously present amplitude-modulated auditory tone (carrier frequency 800 Hz, 45-Hz modulation). This tone was delivered to the subject via MEG-compatible tubephones (model ER30; Etymotic Research, Elk Grove Village, IL) and was adjusted for each subject to be as loud as possible while remaining comfortable. It has been shown previously that such a 'frequency tag' evokes a readily identifiable auditory steady-state response (aSSR) in the MEG signal (Patel and Balaban, 2000). Frequency tags, whether auditory or visual, can be used as 'probes' into cortical dynamics (Silberstein, 1995), and we have previously shown that such tags can be modulated by consciousness (Srinivasan et al., 1999) and attention (Chen et al., 2003; Iversen et al., 2003). We included an auditory frequency tag in the present design to explore whether the corresponding aSSR would be modulated by deception in this task.

In addition to recording neuromagnetic signals, we simultaneously recorded physiologic SCRs using a single-channel SCR device (Autogenics Systems, Wood Dale, IL) which received signals, via custom-made MEG-compatible shielded cabling, from

non-magnetic adhesive pads which were attached to the subjects' non-dominant index and middle fingers (Autogenics Systems, Wood Dale, IL). Subjects were also fitted with several MEG-compatible electrodes (Grass-Telefactor, West Warwick, RI) in order to record eye movements, blinks, swallowing movements and electrocardiogram.

MEG signals were recorded continuously for each block with a sampling rate of 508 Hz. Signals were band-pass filtered at 1–100 Hz. For each trial, MEG signals, SCR signals and electrode signals corresponding to the 'pre' and 'post' periods were extracted for further analysis. The electrode signals were inspected visually for signs of eyeblinks, other muscular artifacts or heart-rate artifacts. 'Pre' or 'post' periods in which any artifacts were present were excluded from further analysis. The subjects' verbal responses were confined to the period in between the 'pre' and 'post' periods and so did not affect the analyzed signals.

Data analysis

For each trial, and for each MEG channel, we used a Fourier transform (MATLAB, Mathworks, Natick, MA) to calculate the power of the aSSR, and the power within the theta (4–8 Hz), alpha (8–12 Hz), beta (15–30 Hz) and gamma (30–50 Hz) frequency ranges.[1] To characterize the time course of power during each trial, we applied a fast Fourier analysis within sliding windows (2 s; step of 0.5 s) for each MEG sensor separately. For each window, we calculated the aSSR power, as well as the power within each of the above frequency bands. SCR signals were characterized by their 'peak-to-peak' value (the absolute difference between the maximum and minimum amplitude) during 'pre' and 'post' periods (Olsson and Phelps, 2004).

For trial-by-trial discrimination of truthful responses from deceptive responses, we used the alpha power of MEG signals during both the 'pre' period and (separately) the 'post' period (see Results). For each subject, and for each period, we composed a dataset of size $Nt \times Nc$, consisting of alpha power values from Nc MEG channels for Nt separate trials ($Nc = 148$, $Nt \leq 70$ depending on exclusion of trials due to artifacts, see Data acquisition above). Within each dataset, we assessed the predictive accuracy of each MEG channel separately (for discrimination of truthful trials from deceptive trials), using N-fold cross-validation with logistic regression as the classification algorithm. As previously mentioned (see Introduction), we employed N-fold cross-validation to ensure accurate characterization of the performance of the logistic classifier. Each dataset is divided into N subsets and a classifier is trained N times, each time leaving out one of the subsets from training, but using only the omitted subset to test the accuracy of the classifier. We note that N-fold cross validation is quite different from the 'split-sample' or 'hold-out' method, in which only a single subset (the validation set) is used to assess the accuracy of a trained classifier, and in which each data point is used only once, either as part of the training set, or as part of the validation set.

Importantly, the present usage of N-fold cross-validation extended the above procedure to take into account the presence of multiple competing classifiers (the different MEG channels). Specifically, in order to determine the final predictive accuracy of MEG for a given dataset, it was necessary to select a single 'best' channel. While high accuracy may be observed by examining performance on training set data only, or by selecting the 'best' channel post hoc based on performance on test set data, such observations may provide inaccurate assessments of predictive accuracy because they do not indicate generalization of performance to new data. To avoid this problem, we assessed the predictive accuracy of each channel on test set data, but we selected the 'best' channel based on accuracy on training set data.

The details of the present procedure are as follows. Within each dataset, alpha power values were pretreated by standardizing. For each channel, 5 repetitions of the cross-validation procedure were carried out. In each repetition, a different 20% of trials were set aside (the test set) and a logistic regression model was fitted to the remaining 80% (the training set). The independent (predictor) variable was the mean alpha power, and the dependent (response) variable was whether the subject had lied or told the truth during the corresponding trial. The accuracy of the resulting logistic classifier was assessed on the training set itself, as well as on the test set. After all 5 cross-validations had been carried out for a given channel, average accuracies were calculated for training set data and, separately, for test set data. After all MEG channels had been assessed, a single channel was chosen based on its average accuracy on training set data. The final predictive accuracy of MEG for the dataset was then given by the average accuracy of this channel on test set data.

To contrast predictive accuracy for instructed deception versus spontaneous deception, we repeated the above cross-validation analysis on two separate subsets of the data. The subset of data corresponding to instructed deception contained only those trials in which the object was either a gun or a suitcase (condition I). The subset of data corresponding to spontaneous deception contained only those trials in which the object was either a necklace or a whisky bottle (condition S). We refer to the original analysis incorporating all trials as condition A.

We also assessed the predictive accuracy of the SCR signal for each subject, in all three conditions. We used the same cross-validation procedure as described above; however, in this case, the dependent (predictor) variable was the standardized peak-to-peak SCR signal during either the 'pre' or 'post' period. Since there was only a single SCR channel, there was no need to choose the 'best' channel. Instead, SCR predictive accuracy was assessed simply by average accuracy on test set data.

Results

Behavior

All subjects successfully completed the experiment. As instructed, they all consistently gave truthful responses (declared) for suitcases and consistently lied (smuggled) for guns (with a single exception on one trial for one subject). However, subjects varied with regard to their responses for whisky and necklaces. 5/20 subjects approximated the optimal strategy of declaring whisky and smuggling guns (subjects A, F, H, O and Q, see Fig. 3). 4/20 subjects lied on almost every trial that did not involve a suitcase (subjects E, R and T). The remaining subjects played a mixed strategy in which they alternated between smuggling or

[1] Power was calculated as the sum of the power spectral density within the corresponding frequency range (Hartmann, 1998). In order to control for possible influences of the frequency tag stimulus on gamma power, when calculating gamma power, we used a notch filter to exclude the aSSR frequency (45 Hz).

declaring both whisky and necklaces. In general, subjects lied more frequently than would have been optimal, which accords with behavioral evidence from both humans and other animals showing risk proneness for potential losses (Kahneman and Tversky, 2000). Further details about the responses of each subject are given in Supplementary Table T1 and Supplementary Fig. S1.

Neuromagnetic analysis

Figs. 2(a–d) show the time course of power during a trial, in the four frequency bands, averaged across all MEG sensors, for a representative subject. Each panel shows the power within a particular frequency band, averaged across all trials (blue line), averaged across trials in which the subject made truthful responses (green line) and averaged across trials in which the subject made deceptive responses (red line). Also shown are (e) the time course of aSSR power and (f) the time course of the peak-to-peak SCR signal. Of the MEG signals, only alpha power was consistently modulated by the structure of the experiment, showing clear elevation during both 'pre' and 'post' periods. In this subject, the SCR signal was highly modulated during the 'post' period, however, when considering all subjects, SCR modulation was much more variable than was alpha power modulation.

Fig. 3 shows differences in the power of the MEG signals between deceptive and truthful responses during the 'pre' period for each subject separately, averaged across all trials and all MEG channels. Panels (a–e) show the power within four different frequency bands as well as the aSSR power, and panel (f) shows average differences in the peak-to-peak SCR response. Data were converted to z scores to allow cross-panel comparisons.

For the 15 subjects showing significant differences in alpha power between truthful responding and deceptive responding (two-tailed t test, $P < 0.05$), 13 showed greater alpha power during truthful responding. The two exceptions in which alpha power significantly decreased (subjects E and J) both had atypical alpha patterns in which power was concentrated in the right-anterior sensor quadrant [data not shown; most other subjects had alpha distributions centered on occipital regions, see Rodin and Rodin, 1995 for a discussion of inter-subject variation in alpha patterns].

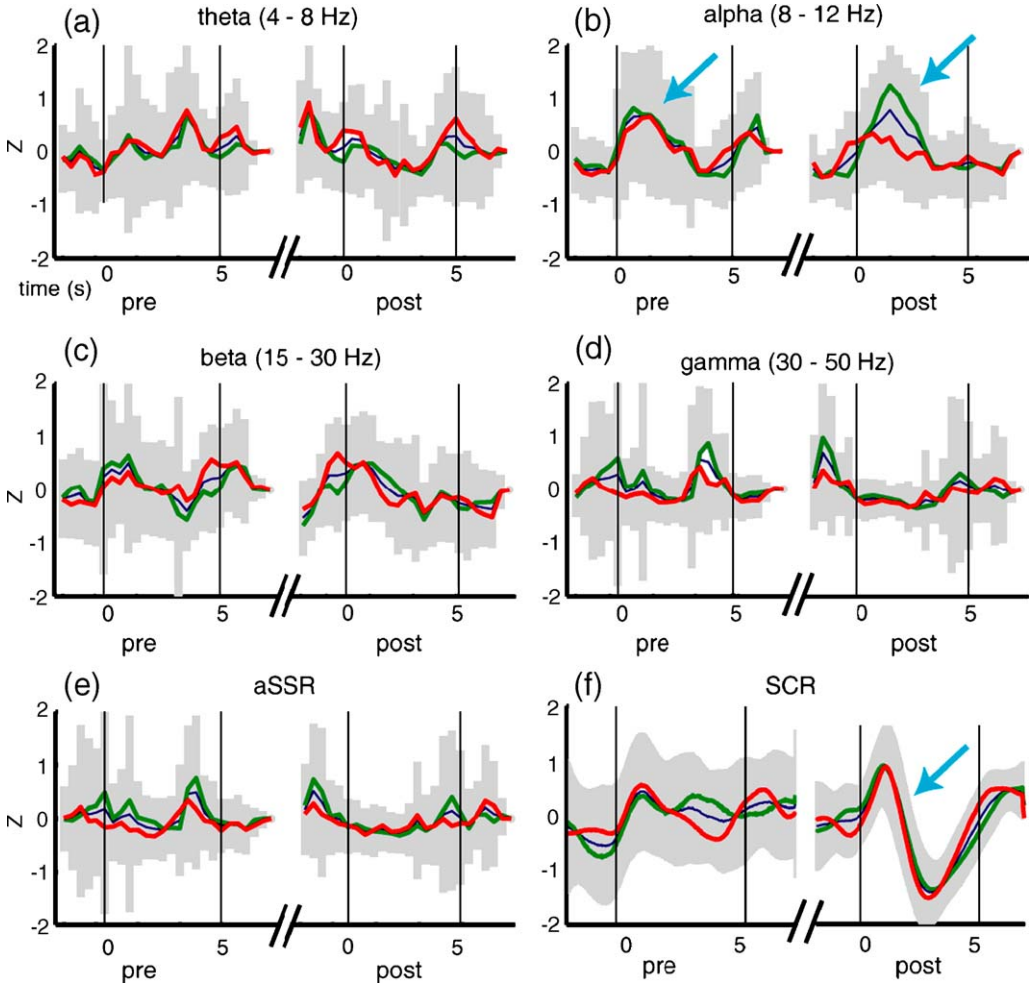

Fig. 2. (a–d) Time course of power in each frequency band, averaged across all frequencies within each band, across all MEG sensors and across all trials (blue line), deception trials only (red line) and truth trials only (green line), for a representative subject. (e). Time course of the aSSR response. (f). Time course of the peak-to-peak SCR signal. Standard errors (for the blue line) are shown in grey, and time series were standardized before averaging. Both the 'pre' period and the 'post' periods are shown. The vertical lines demarcate the onset and offset of each 5-s period. Of the MEG signals, only alpha signals were clearly modulated by the 'pre' and 'post' periods in this subject (blue arrows). The SCR signals, at least during the 'post' period', also showed large modulation (blue arrow). Similar time courses were observed for all 20 subjects (except subject E).

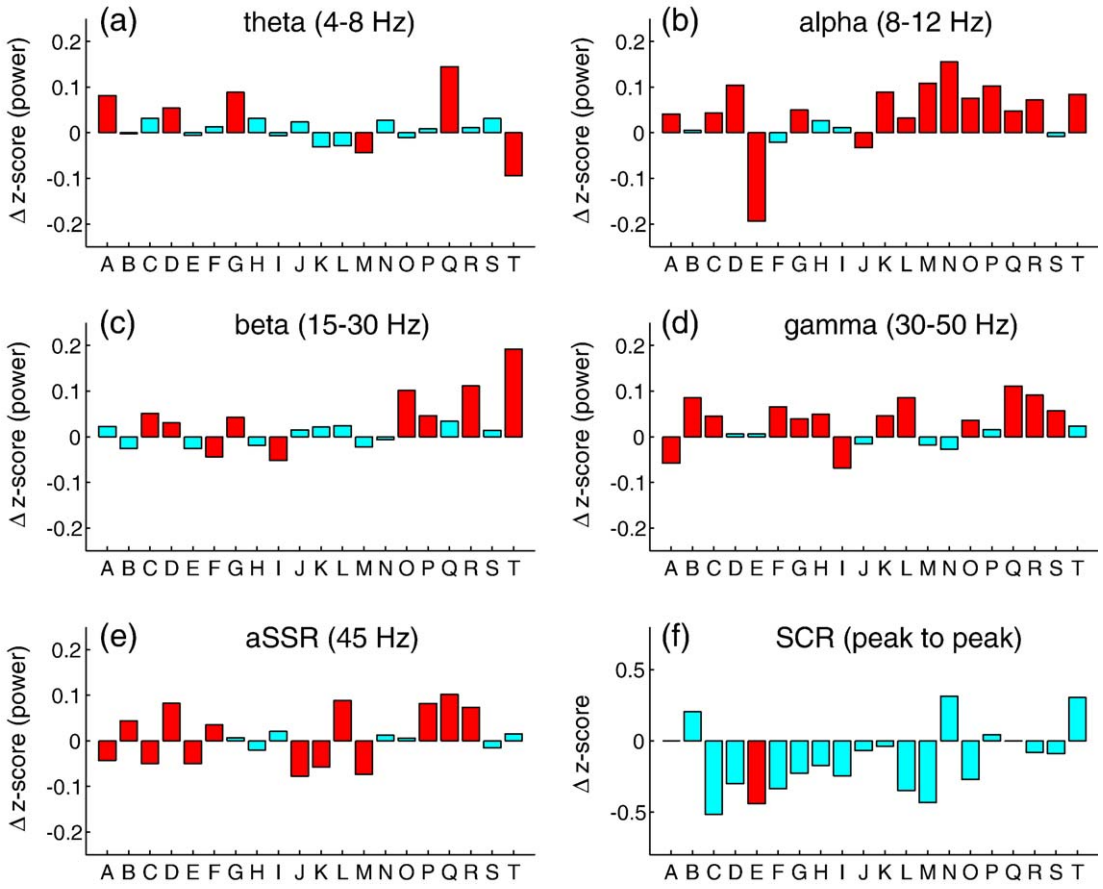

Fig. 3. (a–d) Differences in MEG power between truthful and deceptive responses, averaged across all MEG channels, for each subject ('pre' period). Each panel shows average differences (truth minus deception) within a specific frequency band. Also shown are (e) power differences in the aSSR and (f) average differences in the peak-to-peak SCR response. Data have been converted to z scores to allow cross-panel comparisons (MEG data are plotted in a smaller range [−0.25 to 0.25] than SCR data [−0.5 to 0.5]). Significant differences are shown in red, non-significant differences in blue (two-tailed t test, $P < 0.05$).

Consistent differences in signal power were not seen for any frequency band other than alpha nor for the aSSR. For 14/19 subjects, peak-to-peak SCR signals were lower during truthful responding than during deceptive responding (SCR signals were not available for subject A). Although these differences reached statistical significance only for subject E, this may be attributed to the comparatively small number of data points available for SCR signals (one per 'pre' period or 'post' period) as compared to MEG signals (148 per period).

Taken together, the results in Figs. 2 and 3 suggest that alpha power may be a useful signal for discriminating truthful from deceptive responses in this experiment. To explore this signal further, Fig. 4 shows the topographical distributions of differences in alpha power for three representative subjects, during both 'pre' and 'post' periods. Consistent with Fig. 3, alpha power was higher during truthful responding than during deceptive responding. For these subjects, these differences were most apparent in posterior brain regions. Overall alpha power is also shown for both periods; this power was also concentrated in posterior brain regions. It is notable that there was some variation in these distributions from subject to subject, but that the distributions were consistent within each subject across 'pre' and 'post' periods. For overall alpha power, within-subject consistency was high for all subjects, suggesting that the MEG signals reflected a stable alpha source, or sources.

Single-trial discrimination

Based on the above results, we chose to use alpha power in order to discriminate truthful and deceptive responses on a trial-by-trial basis. We applied N-fold cross-validation (see Materials and methods) separately for each subject and separately for both 'pre' and 'post' periods. Figs. 5(a–b) show predictive accuracy of the MEG signal for each subject, using mean alpha power during the 'pre' period (a) and the 'post' period (b). Predictive accuracy is defined as the average percentage of trials, in the test set, that are correctly classified. Results are shown separately for each subject, as well as in group averages. Due to technical problems, data from 2 subjects were not available for the 'post' period. The different colors reflect different subdivisions of the data. As described in Materials and methods, we assessed predictive accuracy across all trials (condition A; black), across trials involving instructed deception only (condition I; red) and across trials involving spontaneous deception only (condition S; blue).

Across all subjects and including all trials, the mean predictive accuracy was 71.8% (SE = 6%) for the 'post' period and 68.5% (SE = 3.2%) for the 'pre' period. In both cases, predictive accuracy was significantly above chance, and accuracy during the 'post' period was significantly higher than that during the 'pre' period (Wilcoxon signed rank test, $P <$ 0.05). Predictive accuracy for instructed deception was 76.9%

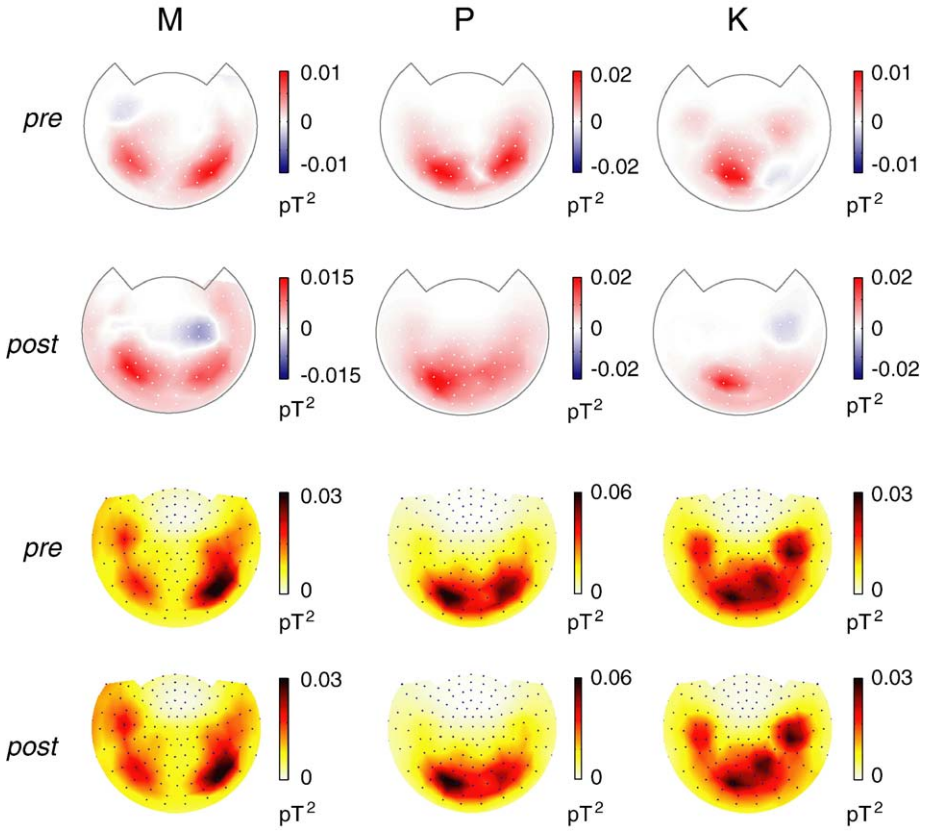

Fig. 4. Topographic distribution of alpha power differences (top two rows) and overall alpha power (bottom two rows) during the 'pre' and 'post' periods, for three example subjects (M, P, K). For the difference plots, red color indicates higher power during truth than during deception, blue color indicates the reverse. Data are averaged across all experimental trials. (1pT = 10^{-12} T).

(SE = 6.3%) for the 'post' period and 75.1% (SE = 4.6%) for the 'pre' period. These accuracies are significantly higher than those attained when all trials were included (Wilcoxon signed rank test, $P < 0.01$). Predictive accuracy for spontaneous deception was 78.3% (SE = 7.2%) for the 'post' period and 73.9% (SE = 4.7%) for the 'pre' period. These accuracies are also significantly higher than those attained when all trials were included (Wilcoxon signed rank test, $P < 0.01$). For 7/18 subjects, using 'post' period data, predictive accuracy for spontaneous deception exceeded 80%, and for subjects F and P, predictive accuracy reached or exceeded 90%.

Table 2 summarizes the above results, showing grand means for predictive accuracy as well as the corresponding levels of sensitivity (the percentage of deceptive responses correctly classified) and specificity (the percentage of truthful responses correctly classified), for all conditions and for both 'pre' and 'post' periods. In most cases, sensitivity was greater than specificity, suggesting that prediction errors were more likely to reflect misclassifications of truthful responses than misclassifications of deceptive responses.

Fig. 6 shows an example of the cross-validation analysis for a representative subject, for conditions A, I and S ('post' period). Each panel shows the average predictive accuracy of each MEG channel. The blue cross marks the 'best' channel as identified from classification performance on training set data. It is the average performance of this channel on test set data that represents the overall predictive accuracy of the MEG signal for this subject (Fig. 5). For conditions A and S, this channel was also the channel that performed best on the test set data. It is important to emphasize that this is not necessarily always the case. For example, in condition I, there is a channel (marked by the green cross) that performed better on test set data than the channel selected on the basis of performance on training set data. However, this channel cannot be used to represent overall predictive accuracy because it could not have been identified on the basis of performance on training set data.

In all three conditions, there were many channels that classified test set data with comparable accuracy to the selected channel. This suggests that the observed predictive accuracy was not based on the 'fluke' activity of a single channel; rather, it reflects neural dynamics that were well represented across large brain regions. On the other hand, there were many channels which were at chance level with respect to test set performance, suggesting that the underlying neural mechanisms may be anatomically or dynamically localized.

In the example subject depicted in Fig. 6, most of the channels with high predictive accuracy were located over posterior brain regions, with the exception of condition I, in which a group of anterior channels also offered high accuracy. However, taking into account these distributions in all 20 subjects, there were no clear differences in distributions between conditions A, I and S. Indeed, there was considerable variability between subjects in the distribution of predictive accuracy. In particular, we did not find any simple relationship between the topographic distributions of predictive accuracy and the topographic distributions of differences in alpha power (Fig. 4), suggesting that differences in alpha power

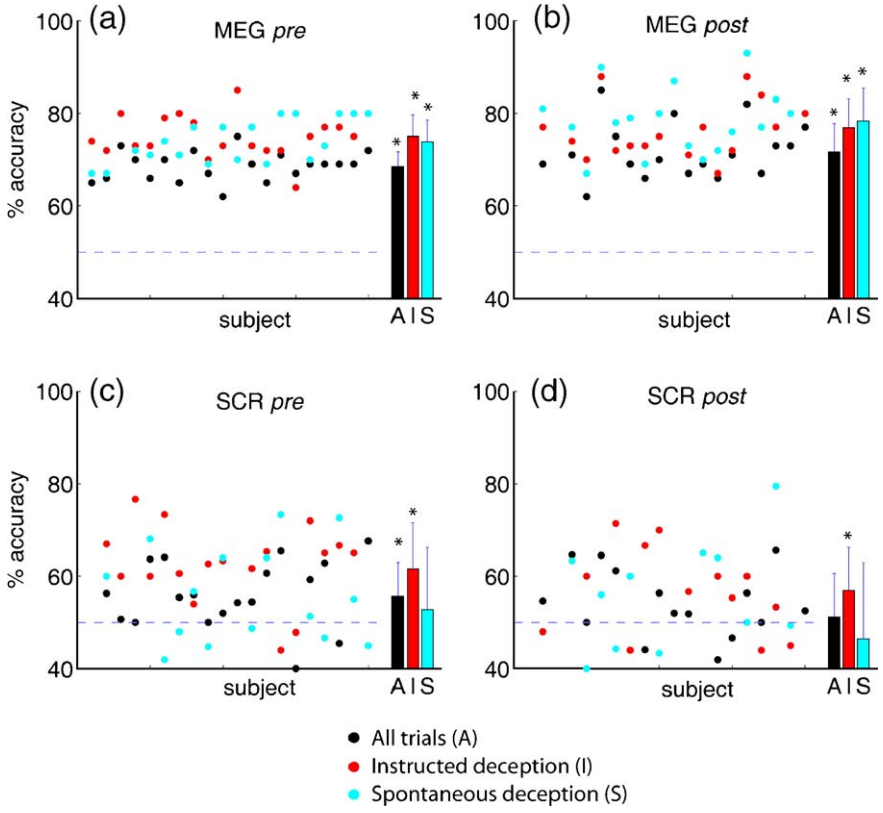

Fig. 5. Predictive accuracy of MEG signals for test set data, for 'pre' period (a) and 'post' period (b). Predictive accuracy of SCR signals for test set data, for 'pre' period (c) and 'post' period (d). Each panel shows predictive accuracies for each subject separately, across all trials (A, black), across trials involving instructed deception only (I, red) and across trials involving spontaneous deception only (S, blue). Group averages are shown in bar charts. Asterisks denote group average accuracies that are significantly above chance level, which is shown by the dashed line ($P < 0.01$, two-tailed t test).

per se may not support accurate trial-by-trial discrimination of deceptive from truthful responses.

Fig. 7 shows the topographic locations of the 'best' channel for all subjects and for both 'pre' and 'post' periods (condition A). While there may be some tendency for channels to cluster in medial regions, this tendency is not strong, and there remains considerable variation in the location of the 'best' channel, both between subjects and within each subject (between 'pre and 'post' periods). This suggests that it may not be possible to identify a single MEG channel as reflecting a difference between truth and deception in general.

Table 2
Grand means for predictive accuracy, sensitivity (percentage of deceptive responses correctly classified) and specificity (percentage of truthful responses correctly classified)

	Condition	Predictive accuracy (%)	Sensitivity (%)	Specificity (%)
Pre	All (A)	68.5	70.0	65.4
	Instructed (I)	75.1	82.7	71.7
	Spontaneous (S)	73.9	73.5	77.3
Post	All (A)	71.8	77.8	64.3
	Instructed (I)	76.9	82.3	75.7
	Spontaneous (S)	78.3	80.1	77.3

In most cases, sensitivity was greater than specificity, suggesting that prediction errors were more likely to reflect misclassifications of truthful responses (false positives) than misclassifications of deceptive responses (false negatives).

Comparison of MEG signals with SCR signals

As shown in Fig. 3, for most subjects, there was a larger peak-to-peak SCR deflection during deceptive responding than during truthful responding (technical problems prevented recording of SCR signals from subject A and, in the 'post' period only, from subject C). Also, the SCR signal appeared to be modulated by the task structure, at least for some subjects (Fig. 2). How well can the SCR signal discriminate deceptive responses from truthful responses on a single-trial basis? As described in Materials and methods, we repeated the cross-validation analysis using peak-to-peak SCR as the predictor variable, retaining logistic regression as the classification algorithm. Figs. 5(c–d) show predictive accuracy of the SCR signal for each subject, for both 'pre' and 'post' periods, and for conditions A, I and S. When averaged across all subjects, predictive accuracy was above chance for conditions A (55.7%, SE = 7.3%) and I (61.6%, SE = 10.0%) in the 'pre' period and in the 'post' period, for condition I only (56.9%, SE = 9.3%) (two-tailed t tests, $P < 0.05$). However, predictive accuracy using the peak-to-peak SCR signal was much lower than the corresponding accuracies achieved using the MEG signal, for all conditions, in both 'pre' and 'post' periods (Wilcoxon signed-rank test, $P < 0.01$). This difference was apparent also on a single-subject basis: Fig. 8 shows the difference between predictive accuracy for MEG and SCR for each subject separately, for both 'pre' and 'post' periods (condition A). For all subjects, predictive accuracy was higher for MEG signals than for SCR. During the 'pre' period, the average increase in predictive accuracy was

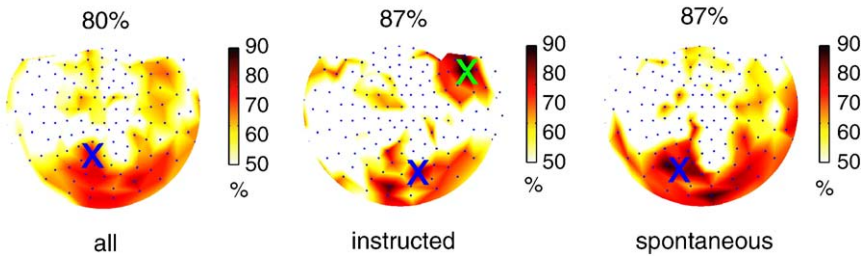

Fig. 6. Example of predictive accuracy for subject K ('post' period). Separate analyses were performed for all trials (condition A), instructed deception only (condition I) and spontaneous deception only (condition S). On each panel, the color scale shows the average predictive accuracy on test set data of each MEG channel. The blue cross identifies the best channel as chosen by performance on training set data; the corresponding predictive accuracy is shown at the top of each plot. For conditions A and S, this channel was the same as the best channel as chosen post hoc by performance on the test data. For condition I, the best channel based on training set performance is *not* the same as the best channel identified post hoc (green cross). In each case, there are several channels that are able to discriminate with high predictive accuracy.

13.1%, and during the 'post' period, the average increase was 20.5%.

Discussion

Functional significance of alpha oscillations for deception

We found that power within the alpha range was modulated by deception with greater consistency across subjects than power within other frequency bands or power of the aSSR. For most subjects, mean alpha power (across the MEG sensor array) was lower during deceptive responding than during truthful responding.

The functional role (or roles) of alpha oscillations remains unresolved. The view that increased alpha reflects cortical 'idling' (Pfurtscheller et al., 1996) has been challenged by multiple studies that show the modulation of alpha oscillations by cognitive functions. For example, frontal alpha power was found to be enhanced during high working memory load (Jensen et al., 2002). In a visual attention task, allocation of attention resulted in increased alpha-band activity in the calcarine and parieto-occipital regions of visual cortex (Yamagishi et al., 2003). However, in the present study, neither working memory nor attention was obviously differentially tasked by deception. It is therefore unlikely that the observed modulations of alpha by deception reflect changing demands on either of these cognitive functions. Nor does the 'idling' hypothesis provide a satisfactory interpretation of our results, since subjects were actively engaged in the task and were experiencing salient financial events, during both truthful responding and deceptive responding.

An alternative view is suggested by the observation that deceptive responding in the present study incurs a higher level of financial risk than truthful responding. In a recent study, Oya and colleagues found a relationship between alpha oscillations and risk in a version of the 'Iowa gambling task' (IGT) (Oya et al., 2005). In the IGT, which was originally designed to assess decision-making impairments in neurological patients (Bechara et al., 1994), subjects pick cards from either 'risky' decks or 'non-risky' decks. Selection from 'risky' decks leads to long term financial loss, whereas selection from 'non-risky' decks leads to long-term financial gain (note that there is no element of deception in the IGT). Oya et al. used intracranial electrodes to record ERPs from the medial prefrontal cortex of a neurosurgical patient, while the IGT was administered.[2] They reported an alpha-band component of ERPs that reflected the mismatch between expected outcomes and actual outcomes but only in trials in which subjects selected from the 'risky' decks.

Recent fMRI studies of deception suggest that deceptive responses may be associated with increased activity of executive areas. While no clear consensus can be expected regarding a unique set of areas, due to considerable variations in the paradigms used, a common finding is increased activity in anterior cingulate cortex (ACC) during deception (Ganis et al., 2003; Langleben et al., 2002; Nunez et al., 2005; Spence et al., 2001). Interestingly, a recent study has implicated the ACC in signaling the predicted likelihood of error (Brown and Braver, 2005), which may reflect risk management. Alternatively, increases in ACC activity may reflect conflict monitoring (Ruff et al., 2001) and inhibition of conflicting responses (Braver et al., 2001) during deception (Nunez et al., 2005). Deceptive responses in the present study may also involve conflict monitoring and response inhibition, raising the possibility that these 'cognitive control' functions may be associated with reduced alpha power. While this relationship has not been examined directly, evidence from neurological patients suggests that the inability to sustain normal patterns of alpha desynchronization can lead to poor performance in a Go/NoGo response inhibition task (Roche et al., 2004).

These lines of evidence together suggest that the present results are unlikely to arise from modulations of working memory, attention or general cortical 'idling'. Instead, alpha modulation during deception in our study may arise from mechanisms of risk management and/or cognitive control. It is also possible that effective risk management may itself require the inhibition of responses to either overly safe or overly risky options. Further studies, directly probing the modulation of alpha oscillations by risk and by cognitive control will be necessary to distinguish among these possibilities.

Kinds of lies

An important objective for experimental studies of deception is to contrive situations in which deceptive responses can be emitted spontaneously by subjects, without instruction. Not only is spontaneous deception more likely to reflect natural behavior, spontaneous responses may also depend on different neural

[2] The patient behaved normally on the IGT, and magnetic resonance scans showed no evidence of damage to the prefrontal cortex.

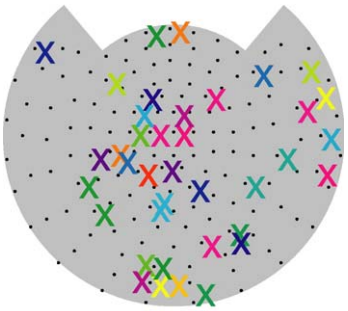

Fig. 7. Crosses mark the 'best' channels (chosen on the basis of performance on training set data in condition A) for each subject separately; different colors indicate different subjects. Both 'pre' and 'post' periods are shown. There was a wide variety of 'best' channels and no consistency within a subject across 'pre' and 'post' periods.

mechanisms than instructed responses (Walton et al., 2004). The present design contrasts spontaneous deception with instructed deception in a transparent fashion by allowing subjects to decide for themselves whether to lie or tell the truth for certain stimuli.

Previous brain-imaging studies of deception have involved elements of both spontaneous deception and instructed deception. Kozel et al. (2004a,b) and Davatzikos et al. (2005); Langleben et al. (2002, 2005)) used modified versions of the 'guilty knowledge task' (GKT) in which subjects were instructed to lie. Spence and colleagues instructed subjects to lie in response to the presence of a particular color (Spence et al., 2001), and Lee and colleagues instructed subjects to feign poor performance on a series of simple tests (Lee et al., 2002). Both Rosenfeld et al. (1999) and Johnson et al. (2005) instructed their subjects to lie, at times of their own choosing, on about 50% of trials. Finally, Ganis and colleagues utilized a rich design distinguishing 'memorized' lies, which involved only the retrieval of stored information, from 'spontaneous' lies, which were constructed on-the-fly. We note, however, that in the 'spontaneous' condition of the Ganis et al. study, subjects were still instructed to lie, even if they decided for themselves the content of each lie.

An additional feature of the present design is that deceptive responses and truthful responses were made with comparable frequencies. This feature, which is not unique to the present study, avoids a potential confound of the occurrence of deceptive responses with the appearance of rare, meaningful stimuli in long sequences of distractors. In the standard GKT, for example, subjects make deceptive responses only for a small subset of stimuli, which may facilitate their detection via the 'oddball' ERP response (Donchin and Coles, 1988; Rosenfeld, 2001).

Single-trial discrimination

We found that power within the alpha band could support trial-by-trial detection of deceptive responses. We used a cross-validation method in which a logistic classifier was trained on 80% of the data, and its predictive accuracy was assessed on the remaining 20%. While predictive accuracy including all data was 71.8%, accuracy was significantly higher when the data set for each subject was restricted to trials involving only instructed deception (76.9%) or only spontaneous deception (78.3%) (all results from the 'post' period). For all subjects, predictive accuracy was well above chance, and for 7/18 subjects, predictive accuracy for spontaneous deception exceeded 80%. These results indicate that alpha-band signals may potentially be useful for detecting deception in forensic and clinical situations (see Practical application below).

Why should predictive accuracy be higher for spontaneous or instructed deception (conditions S and I) than when all trials are included (condition A)? One possibility is that restriction of the data set, to two out of four possible objects, may reduce any variation in the brain response that is due to differences among the objects themselves. This reduction in 'irrelevant' variation may enable enhanced predictive accuracy for detecting deceptive responses. Another possibility is that spontaneous deception and instructed deception may recruit different cognitive and neural mechanisms. In our data, however, a contrast of spontaneous deception with instructed deception revealed no consistent differences in the topographical distributions of alpha power, of distributions of differences in alpha power, or in the locations of MEG channels providing high predictive accuracy.

Several other studies have attempted to use brain signals to detect deceptive responses on a single-trial and/or single-subject basis. ERP methods, while commonly applied to single subjects, typically require averaging over multiple trials in order to identify reliable 'oddball' responses in the ERP waveform [Allen et al., 1992; Farwell and Donchin, 1991; Rosenfeld et al., 1988, 1991; see Rosenfeld, 2001 for a review]. In two recent studies, machine learning techniques were applied to multivariate fMRI signals to

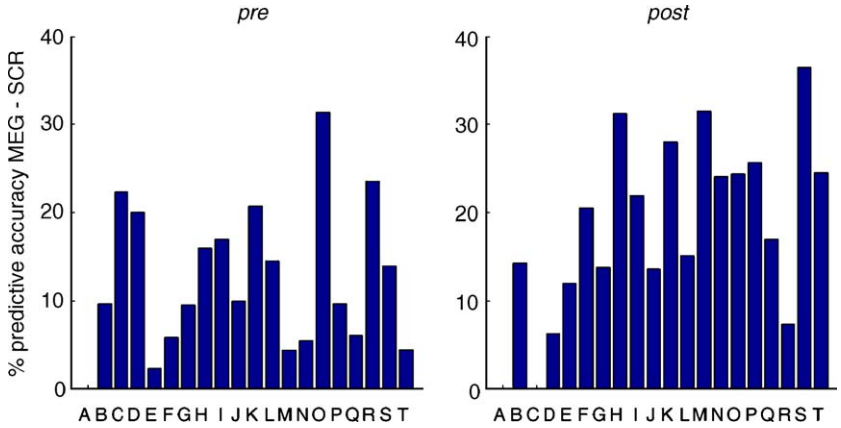

Fig. 8. Predictive accuracy was greater for MEG signals than for SCR signals, for all subjects (condition A). Each panel shows the difference in average predictive accuracy on test set data for each subject separately. During the 'pre' period (left panel), the average increase in predictive accuracy was 13.1%, and during the 'post' period (right panel), the average increase was 20.5%.

detect deceptive responses on single-subject and single-trial basis, using a modified version of the GKT (Davatzikos et al., 2005; Langleben et al., 2005). In Langleben et al. (2005), stepwise multivariate logistic regression was used to detect deception. Predictor variables were chosen by identifying voxels that showed significantly different activity between truth and deception conditions, and predictive accuracy was validated by analyzing four additional subjects. In the second study, the same data were reanalyzed using a support-vector machine approach (Davatzikos et al., 2005), with predictive accuracy assessed by a version of cross-validation in which training sets were constructed from 99% of the data, with test sets consisting of the remaining 1%.

There are two main differences between the analyses of Davatzikos et al. (2005); Langleben et al. (2005) and the present study (besides the obvious contrast between fMRI and MEG). First, as mentioned above, there are several differences between the GKT and the present experimental design. Second, both of the above analyses used multivariate statistics, in contrast to the present approach of selecting a single MEG channel by univariate logistic regression. To test the utility of combining signals from multiple channels in the present experiment, we attempted several multivariate analyses, using both multivariate logistic regression and support vector machines. Although we were able to achieve better classification accuracy on training set data, in some cases attaining 100% (data not shown), in no case did performance on test set data improve to levels beyond those reported here (see Results).

We have emphasized the use of cross-validation as a well-established means of estimating the generalization ability of a classifier (Hastie et al., 2001), as well as for achieving single-trial discrimination (see Materials and methods). However, along with Davatzikos et al. (2005), we stress that cross-validation does not insure against training on data that are not fully representative of a statistical distribution. A cross-validated classifier may perform poorly on novel data drawn from different distributions (e.g., from different subjects, imaging equipment or experimental paradigms). This caveat reinforces the notion that brain-based deception detection methods are likely best applied on a subject-by-subject and experiment-by-experiment basis.

Comparison with SCR signals

In order to assess whether brain-based methods provide enhanced predictive accuracy over alternative methods, it is important to compare the predictive accuracy of different methods within the same experimental session. To our knowledge, this has not previously been done. In the present study, we compared the predictive accuracy of MEG signals with that of simultaneously recorded SCR signals. For all subjects, higher predictive accuracy was obtained using MEG signals. During the 'pre' period, the average increase in predictive accuracy was 13.1%, and during the 'post' period, the average increase was 20.5%. We note that predictive accuracy using SCR signals in the present design was lower than accuracies reported for SCR signals in the GKT (Ben-Shakhar and Elaad, 2003), which might be accounted for by the relative rarity of deceptive responses in the GKT (Elton et al., 1983), as compared to the present design (see Kinds of lies above).

Practical application

Although it was not our intention to develop a practical system for deception detection, our results have several implications for possible practical applications. (1) By contrasting spontaneous deception with instructed deception within a simulated 'customs' environment involving real financial risk, we approximated certain naturalistic environments in which deceptive responses occur spontaneously and frequently (quite unlike the GKT), and we showed a robust neural correlate of deception in this context (reduced alpha power). (2) A within-experiment comparison of alpha-band signals with SCR signals showed that alpha signals provided significantly higher predictive accuracy than the physiological measure. (3) By using logistic regression within a cross-validation analysis, we were able to demonstrate successful single-subject, single-trial discrimination. (4) It may well be possible to utilize alpha-band signals recorded with (portable and cheap) EEG to achieve similar results.

Conclusions

The results of our study represent two main contributions to experimental studies of deception. First, they derive from an ecologically valid experimental paradigm in which deceptive responses were given both spontaneously and as a result of instruction and in which these responses were balanced in frequency with truthful responses. Second, MEG signals – specifically power within the alpha band – were used to support trial-by-trial discrimination of truthful and deceptive responses with high accuracy. By cross-validation, we were able to show that the corresponding predictive accuracy significantly exceeded that provided by a simultaneously recorded physiological SCR signal. Further studies are required to ascertain whether the observed reductions in alpha power during deception reflect neural mechanisms of risk management or cognitive control.

Acknowledgments

For the useful discussions, we thank Drs. Aniruddh Patel, Joseph Gally, Eugene Izhikevich Andy Philippides, and Bruno van Swinderen. We are also grateful for the comments of two anonymous reviewers on an earlier draft of this article. We thank Lacey Kurelowech for expert contribution in data collection and technical support. This research was partially supported by funding from the Central Intelligence Agency and by the Neurosciences Research Foundation.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2006.02.050.

References

Allen, J.J., 2002. The role of psychophysiology in clinical assessment: ERPs in the evaluation of memory. Psychophysiology 39, 261–280.

Allen, J.J., Iacono, W.G., Danielson, K.D, 1992. The identification of concealed memories using the event-related potential and implicit behavioral measures: a methodology for prediction in the face of individual differences. Psychophysiology 29, 504–522.

Bechara, A., Damasio, A.R., Damasio, H., Anderson, S.W., 1994. Insensitivity to future consequences following damage to human prefrontal cortex. Cognition 50, 7–15.

Ben-Shakhar, G., Elaad, E., 2003. The validity of psychophysiological

detection of information with the guilty knowledge test: a meta-analytic review. J. Appl. Psychol. 88, 131–151.
Braver, T.S., Barch, D.M., Gray, J.R., Molfese, D.L., Snyder, A., 2001. Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. Cereb. Cortex 11, 825–836.
Brown, J.W., Braver, T.S., 2005. Learned predictions of error likelihood in the anterior cingulate cortex. Science 307, 1118–1121.
Browne, M.W., 2000. Cross-validation methods. J. Math. Psychol. 44, 108–132.
Chen, Y., Seth, A.K., Gally, J., Edelman, G.M., 2003. The power of human brain magnetoencephelographic signals can be modulated up or down by changes in an attentive visual task. Proc. Natl. Acad. Sci. U. S. A., 3501–3506.
Davatzikos, C., Ruparel, K., Fan, Y., Shen, D.G., Acharyya, M., Loughead, J.W., et al., 2005. Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. Neuroimage 28, 663–668.
Donchin, E., Coles, M., 1988. Is the P300 component a manifestation of context updating? Behav. Brain Sci. 11, 357–374.
Elton, M., Schandry, R., Sparrer, B., 1983. A comparative investigation of ERP components and the SCR in a habituation and dishabituation paradigm. Int. J. Neurosci. 22, 55–62.
Farwell, L.A., Donchin, E., 1991. The truth will out: interrogative polygraphy ("lie detection") with event-related brain potentials. Psychophysiology 28, 531–547.
Ganis, G., Kosslyn, S.M., Stose, S., Thompson, W.L., Yurgelun-Todd, D.A., 2003. Neural correlates of different types of deception: an fMRI investigation. Cereb. Cortex 13, 830–836.
Hartmann, W.M., 1998. Signals, Sound, and Sensation. Springer, New York, NY.
Hastie, T., Tibshiranis, E., Friedman, J., 2001. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer-Verlag, New York.
Iversen, J.R., Chen, Y., Seth, A.K., Edelman, G.M., Patel, A.D., 2003. Attention Increases Long Distance Phase Coherence, But Not Power, of the Cortical Auditory Steady-State Response in a Bimodal Paradigm. Society for Neuroscience, San Diego, CA.
Jensen, O., Gelfand, J., Kounios, J., Lisman, J.E., 2002. Oscillations in the alpha band (9–12 Hz) increase with memory load during retention in a short-term memory task. Cereb. Cortex 12, 877–882.
Johnson Jr., R., Barnhardt, J., Zhu, J., 2005. Differential effects of practice on the executive processes used for truthful and deceptive responses: an event-related brain potential study. Brain Res. Cogn. Brain Res. 24, 386–404.
Kahneman, D., Tversky, A. (Eds.), Choices, Values, and Frames. Cambridge University Press, Cambridge, UK.
Kozel, F.A., Padgett, T.M., George, M.S., 2004a. A replication study of the neural correlates of deception. Behav. Neurosci. 118, 852–856.
Kozel, F.A., Revell, L.J., Lorberbaum, J.P., Shastri, A., Elhai, J.D., Horner, M.D., et al., 2004b. A pilot study of functional magnetic resonance imaging brain correlates of deception in healthy young men. J. Neuropsychiatry Clin. Neurosci. 16, 295–305.
Langleben, D.D., Schroeder, L., Maldjian, J.A., Gur, R.C., McDonald, S., Ragland, J.D., et al., 2002. Brain activity during simulated deception: an event-related functional magnetic resonance study. Neuroimage 15, 727–732.
Langleben, D.D., Loughead, J.W., Bilker, W.B., Ruparel, K., Childress, A.R., Busch, S.I., et al., 2005. Telling truth from lie in individual subjects with fast event-related fMRI. Hum. Brain Mapp. 26, 262–272.
Lee, T.M., Liu, H.L., Tan, L.H., Chan, C.C., Mahankali, S., Feng, C.M., et al., 2002. Lie detection by functional magnetic resonance imaging. Hum. Brain Mapp. 15, 157–164.
Lykken, D.T., 1959. The GSR in the detection of guilt. J. Appl. Psychol. 43, 385–388.
Lykken, D.T., 1991. Why (some) Americans believe in the lie detector while others believe in the guilty knowledge test. Integr. Physiol. Behav. Sci. 26, 214–222.
Nunez, P., 1995. Neocortical Dynamics and Human EEG Rhythms. Oxford University Press, Oxford.
Nunez, J.M., Casey, B.J., Egner, T., Hare, T., Hirsch, J., 2005. Intentional false responding shares neural substrates with response conflict and cognitive control. Neuroimage 25, 267–277.
Office of Technology Assessment, 1990. The Use of Integrity Tests for Pre-Employment Screening. U.S. Office of Technology Assessment, Washington, DC.
Olsson, A., Phelps, E.A., 2004. Learned fear of "unseen" faces after Pavlovian, observational, and instructed fear. Psychol. Sci. 15, 822–828.
Oya, H., Adolphs, R., Kawasaki, H., Bechara, A., Damasio, A., Howard, M.A. III, 2005. Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex. Proc. Natl. Acad. Sci. U. S. A. 102, 8351–8356.
Patel, A.D., Balaban, E., 2000. Temporal patterns of human cortical activity reflect tone sequence structure. Nature 404, 80–84.
Pfurtscheller, G., Stancak, A. Jr., Neuper, C., 1996. Event-related synchronization (ERS) in the alpha band-an electrophysiological correlate of cortical idling: a review. Int. J. Psychophysiol. 24, 39–46.
Roche, R.A., Dockree, P.M., Garavan, H., Foxe, J.J., Robertson, I.H., O'Mara, S.M., 2004. EEG alpha power changes reflect response inhibition deficits after traumatic brain injury (TBI) in humans. Neurosci. Lett. 362, 1–5.
Rodin, E.A., Rodin, M.J., 1995. Dipole sources of the human alpha rhythm. Brain Topogr. 7, 201–208.
Rosenfeld, J.P., 2001. Event-related potentials in the detection of deception. In: Kleiner Jr., M (Ed.), Handbook of Polygraphy. Academic Press, New York.
Rosenfeld, J.P., Cantwell, B., Nasman, V.T., Wojdac, V., Ivanov, S., Mazzeri, L., 1988. A modified, event-related potential-based guilty knowledge test. Int. J. Neurosci. 42, 157–161.
Rosenfeld, J.P., Angell, A., Johnson, M., Qian, J.H., 1991. An ERP-based, control-question lie detector analog: algorithms for discriminating effects within individuals' average waveforms. Psychophysiology 28, 319–335.
Rosenfeld, J.P., Ellwanger, J.W., Nolan, K., Wu, S., Bermann, R.G., Sweet, J., 1999. P300 scalp amplitude distribution as an index of deception in a simulated cognitive deficit model. Int. J. Psychophysiol. 33, 3–19.
Ruff, C.C., Woodward, T.S., Laurens, K.R., Liddle, P.F., 2001. The role of the anterior cingulate cortex in conflict processing: evidence from reverse Stroop interference. Neuroimage 14, 1150–1158.
Silberstein, R., 1995. Steady-state visually evoked potentials, brain resonances, and cognitive processes. In: Nunez, P (Ed.), Neocortical Dynamics and Human EEG Rhythms. Oxford University Press, Oxford, pp. 272–303.
Spence, S.A., Farrow, T.F., Herford, A.E., Wilkinson, I.D., Zheng, Y., Woodruff, P.W., 2001. Behavioural and functional anatomical correlates of deception in humans. NeuroReport 12, 2849–2853.
Spence, S.A., Hunter, M.D., Farrow, T.F., Green, R.D., Leung, D.H., Hughes, C.J., et al., 2004. A cognitive neurobiological account of deception: evidence from functional neuroimaging. Philos. Trans. R. Soc. Lond., B Biol. Sci. 359, 1755–1762.
Srinivasan, R., Russell, D.P., Edelman, G.M., Tononi, G., 1999. Increased synchronization of neuromagnetic responses during conscious perception. J. Neurosci. 19, 5435–5448.
Thornton, K.E., 1995. The anatomy of a lie: an exploratory investigation via the quantitative EEG. J. Offender Rehabil. 22, 179–210.
Walton, M.E., Devlin, J.T., Rushworth, M.F., 2004. Interactions between decision making and performance monitoring within prefrontal cortex. Nat. Neurosci. 7, 1259–1265.
Yamagishi, N., Callan, D.E., Goda, N., Anderson, S.J., Yoshida, Y., Kawato, M., 2003. Attentional modulation of oscillatory activity in human visual cortex. Neuroimage 20, 98–113.