Identifying Spatially Overlapping Local Cortical Networks with MEG

Keith Kawabata Duncan,[1] Avgis Hadjipapas,[2,3] Sheng Li,[4] Zoe Kourtzi,[4] Andy Bagshaw,[4] and Gareth Barnes[5]*

[1]*Institute of Cognitive Neuroscience, UCL, London WC1E 6BT*
[2]*School of Life and Health Sciences, Aston University, Birmingham B4 7ET*
[3]*Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands*
[4]*School of Psychology, University of Birmingham, Birmingham UK. B15 2TT*
[5]*Wellcome Trust Centre for Neuroimaging, UCL, London, WC1N 3BG*

◆ ══════════════════════════════ ◆

Abstract: Recent modelling studies (Hadjipapas et al. [2009]: Neuroimage 44:1290-1303) have shown that it may be possible to distinguish between different neuronal populations on the basis of their macroscopically measured (EEG/MEG) mean field. We set out to test whether the different orientation columns contributing to a signal at a specific cortical location could be identified based on the measured MEG signal. We used 1.5deg square, static, obliquely oriented grating stimuli to generate sustained gamma oscillations in a focal region of primary visual cortex. We then used multivariate classifier methods to predict the orientation (left or right oblique) of the stimuli based purely on the time-series data from this one location. Both the single trial evoked response (0–300 ms) and induced post-transient power spectra (300–2,300 ms, 20–70 Hz band) due to the different stimuli were classifiable significantly above chance in 11/12 and 10/12 datasets respectively. Interestingly, stimulus-specific information is preserved in the sustained part of the gamma oscillation, long after perception has occurred and all neuronal transients have decayed. Importantly, the classification of this induced oscillation was still possible even when the power spectra were rank-transformed showing that the different underlying networks give rise to different characteristic temporal signatures. *Hum Brain Mapp 31:1003–1016, 2010.*   © 2009 **Wiley-Liss, Inc.**

Key words: MEG; mean field; gamma; classifier; orientation; grating; beamformer

◆ ══════════════════════════════ ◆

INTRODUCTION

A number of invasive [Gail et al., 2000; Henrie and Shapley, 2005; Rols et al., 2001] and non-invasive [Adja-mian et al., 2004; Hall et al., 2005; Muthukumaraswamy and Singh, 2008] studies have characterised the gamma band response to simple (grating) visual stimuli. This response has an early evoked transient (0–300 ms), followed by a sustained oscillation in the 30–70 Hz band which persists for the duration of the (static) stimulus and is dominant in, and retinotopically mapped to, primary visual cortex. The role of this sustained gamma activation remains unclear, as it persists long after neuronal transients have adapted, and the perception of the stimulus has occurred. It has been speculated that it represents some non-specific state of cortical activation [Henrie and Shapley, 2005]. On the other hand, studies examining the synchronization properties of gamma oscillations across local cortical networks suggest quite the opposite, namely that

gamma oscillations are highly stimulus/object-specific (see Frien and Eckhorn, 2000a, Gail et al., 2000, Eckhorn et al., 2004 for a review).

We hypothesized that exciting orientation specific subpopulations of neurons (within the same cortical area) by the same coherent visual object (coherent texture of the grating stimulus) would give rise to a locally synchronous stimulus-specific network in the gamma band. Following from this, stimuli of different orientation are expected to give rise to differential collective (synchronous) modes in a largely spatially overlapping cortical network. This theoretical prediction is supported by a recent modelling study looking at networks of coupled systems (representing a rough analogy to networks of coupled cortical modules), which was then observed by a macroscopic spatially aggregate signal, in rough analogy to the biophysical generation of MEG/EEG signals [Hadjipapas et al., 2009]. We showed that information about collective modes expressed in this spatially identical microscopic network of coupled systems can be recovered from temporal properties of the macroscopic aggregate signal. Hence, temporal properties of the macroscopic aggregate can differentiate between distinct collective modes (distinct dynamics) in the underlying network. Consequently, one should be still able to differentiate responses to different stimulation by the temporal properties of aggregate MEG signals, even if these responses arise from the same area of cortex. Invasive studies of macaque LFPs which offer superior spatial resolution seem to be in agreement with this idea: the peak frequency of fast (gamma) oscillations depends on stimulus orientation, such that stimuli of different orientation show a small but significant difference in peak gamma frequencies [Frien and Eckhorn, 2000a,b]. Our previous study [Hadjipapas et al., 2007] provided some empirical evidence supporting this idea in non-invasive MEG measurements in man: it was shown that different spatial frequency grating stimuli could be distinguished through the spectral shape of the stimulus related gamma response. A curious finding was that the stimulus-discrimination ability improved as the analysis window extended 4s beyond the initial transient, even though the stimuli remained static; implying that the sustained portion of the oscillation contained stimulus-related information.

Previous studies have shown the dependence of univariate features of the gamma spectra (such as peak power or frequency) on stimulus features such as motion [Siegel et al., 2007], contrast [Hall et al., 2005; Henrie and Shapley, 2005], colour [Adjamian et al., 2008] or spatial frequency [Adjamian et al., 2004]. To test whether we could identify distinct populations of neurons based on the macroscopic signal properties we sought to avoid stimuli modulated along these dimensions. Additionally we were keen to be able to ascribe any changes in signal to a relatively small functionally homogeneous area of cortex, removing any potential confound of changing source location from the analysis. For example, in a previous study [Hadjipapas et al., 2007] we used centrally fixated gratings of 6° eccentricity. It could be argued that the different power spectra

observed were due to the mixing of signals from spatially distinct (perhaps differentially spatial frequency sensitive) regions of a large area of visual cortex (as opposed to the interaction of different neuronal populations within the same area). Thus we used small (1.5°) oblique grating stimuli differing only in orientation (Fig. 1A). Such (orthogonal) stimuli have similar psychophysical detection thresholds [Campbell et al., 1966]. Orientation columns are arranged in a regular periodic pattern throughout V1 and (in the primate) are known to repeat every 0.5 mm [Hubel and Wiesel, 1974; McLoughlin and Schiessl, 2006]. That is, for any two stimuli we aimed to excite different populations of cells within the same area of cortex. Given that the estimated cortical area stimulated was relatively small (between 9*9 mm and 20*20 mm, based on linear extrapolation about the 0.75° point [Schira et al., 2007] and integration under the foveal cortical magnification curve presented in Schira et al. [2009] respectively), and that the sustained portion of the gamma oscillation is known to be constrained to primary visual cortex [Leopold et al., 2003; Rols et al., 2001], together with the fact that these oscillations are coherent over cortical distances of only a few millimetres [Eckhorn et al., 1993; Leopold et al., 2003] made the signal an ideal candidate for MEG beamformer analysis [Hillebrand et al., 2005; VanVeen et al. 1997].

We were particularly interested in whether differences could be reliably detected in the stimulus related spectral distributions when only utilizing the post-transient portion of the gamma oscillation. We were also interested to see whether any differences between the stimuli could be discerned from the single-trial evoked response, as based on single neuron recordings in the primate, orientation selectivity seems to be developed within the first 150 ms [Celebrini et al., 1993; Muller et al., 2001].

METHODS

All participants had normal or corrected to normal vision (four males, two females, aged between 23 and 52) and were instructed to passively view the fixation spot. The stimuli were generated on a VSG 3/2 graphics card, presented on a gamma corrected Eizo Flexscan T562-T CRT monitor with a refresh rate of 100 Hz. Stimuli were stationary square wave gratings presented either to the lower right or left of fixation (1.5 degrees in extent, 3cpd, 80% contrast) obliquely oriented at either +45°(right) or −45° (left oblique) to vertical (Fig. 1A). Randomly intermixed (left or right oblique) stimuli were presented for 2.5 s and followed by a 2.5-s period of a uniform field of the same mean luminance (baseline period). Subjects were instructed to simply look at the fixation spot but no other instructions were given (i.e. there was no attentional task). The experiment comprised two (left and right visual field) 800-s runs each of which contained 80 presentations of each stimulus type. MEG data were recorded on a 275 channel whole head system (CTF Systems Inc.) at a sample rate 600 Hz with 150 Hz low-pass filter.
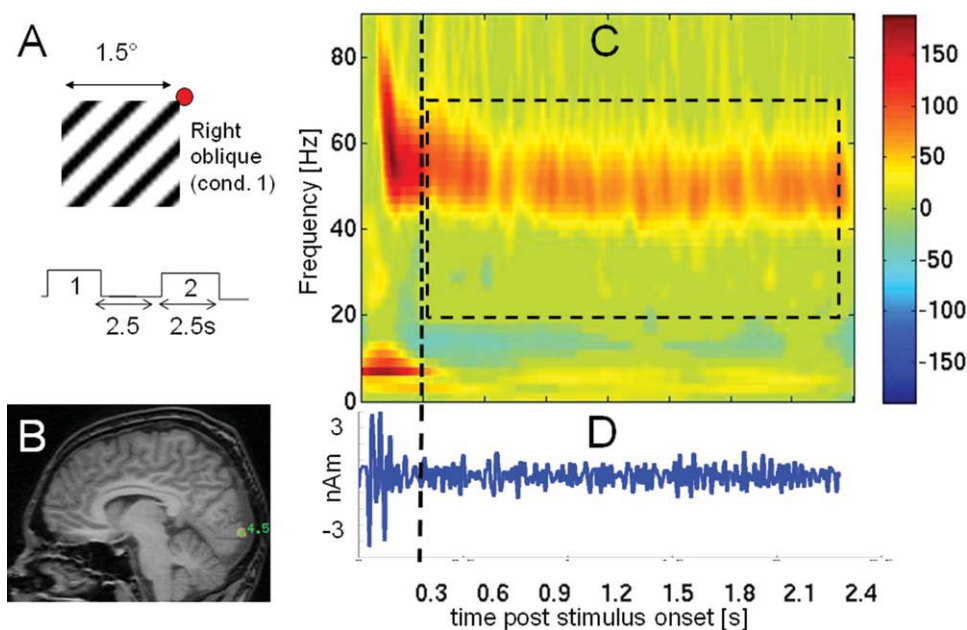
Figure 1.

A: An example of the right oblique stimulus (a square, square wave 3cpd grating of 1.5° side length) presented in the lower left visual field. Subjects passively viewed the fixation spot (red dot) whilst the stimuli were presented statically (without any temporal modulation) from a background of the same mean luminance for 2.5 s in random sequence. B: Sagittal T1 image with stimulus induced gamma activity superimposed for a single dataset (hemisphere 1). Red/orange colour shows power changes (pseudo T, peak = 4.5) between stimulus (right oblique, left visual field) and baseline in primary visual cortex. C: Time frequency representation of the response in visual cortex, non-significant changes set to zero (see methods). Note that the response starts with a broadband transient (0–300ms) corresponding to the evoked response (D), and then evolves in a gamma oscillation (30–60 Hz), which is sustained for the entire duration of visual stimulation. D: Average filtered (20–70 Hz) evoked response to stimulus, note that the stimulus onset transient has largely adapted after the first 300 ms.

Head motion was monitored using three coils attached to the head and energised during the start and end of recording. Experimental runs with head movement of greater than 5 mm were discarded. Each individual subject's data was coregistered with their anatomical MRI by a matching a digitized head surface (including the three coils) consisting of typically 500 points (Polhemus) to the surface of the MRI.

Source Estimation

A beamformer spatial filtering algorithm [Robinson, 1999; VanVeen et al., 1997] was used to generate maps of electrical power change on a 1 mm grid throughout the brain. We compared the power in the 20–70 Hz band between the baseline (−2.5 to 0s) and stimulus (0 to 2.5 s) periods (Fig. 1B). The locations of the image peaks due to both individual (e.g. left oblique vs. baseline) and combined (left and right oblique vs. baseline) contrasts were noted. The combined vs. baseline condition was simply the contrast between all stimuli and baseline data segments. The orientation of the underlying dipolar source in the beamformer scan was set to that which maximised projected power as per [Sekihara et al., 2004]. We then constructed a time-varying estimate of the electrical activity for the peak location of the image based on a spatial filter constructed from (a covariance matrix comprising) baseline and combined conditions [Barnes and Hillebrand, 2003]. That is the covariance window was −2.5 to 2.5 s, 20–70 Hz. In other words, there was no bias in the spatial filter construction towards either orientation condition. These estimates were obtained independently for right and left visual field presentation in each of the six subjects. All subsequent analyses, including spectral estimates and classification, were based on these time series data. Time frequency representations at the location of interest (as in Fig. 1B) were constructed as follows: single trial data were convolved with Morlet wavelets (cycle width = 8 and resolution 2 cycles/s) and the resultant wavelet energies for each time-frequency bin were averaged across epochs. For each post-stimulus time frequency bin a percentage change from the mean of the pre-stimulus window (−2.3 to 0 sec) at that frequency was calculated. Significant changes from the prestimulus to the stimulus time window were assessed using a bootstrap technique (500 iterations) and thresholded whereby non-significant ($P > 0.05$) changes were set to 0.

Figure 2.

Spectral power (top) and ranked power (lower panels) estimates due to the post onset-transient sustained gamma activity (300–2,300 ms) for three datasets (1,12,3). Mean spectra (across trials) for condition 1 (right oblique stimuli) are in red/solid, spectra for condition 2 (left oblique stimuli) are in blue/dotted, and spectra for the baseline (−2,300 to −300 ms) are in green/dashed. Error bars denote one standard error. Although the baseline spectra fall monotonically with frequency, both left and right obliquely orientated stimuli give rise to increases in spectral power in the gamma band (30–60 Hz, see inset), which show slightly different deviations from the monotonic decay. Lower panels show the same data plotted according to the ranked spectral bins (i.e. bin with smallest magnitude ranked 1) on an epoch by epoch basis and then averaged. For **A, B** the difference between these distributions suggests that the two spectra have characteristically different shape (i.e. peak at different frequencies) and do not differ solely in magnitude. This is not the case for **C** where the differences observed in the power spectra could have been due to simply greater power due to one stimulus condition over the other.

Classification

To investigate whether gamma oscillations reflect a stimulus specific cortical state, we tested whether stimuli of different orientation induce gamma oscillations with distinct post-transient spectral distributions. For each epoch we calculated the spectral density in post-onset stimulus (300–2,300 ms) periods. Spectral estimates were based on the Matlab psd function using Welch's average periodogram method with a Tukey window function. We used the maximum segment length (NFFT in Matlab) possible, (i.e. spectral resolution at the expense of accuracy) using 1,200 and 180 samples for the post-transient and transient periods, respectively. This data was analysed in event-related form by subtracting the pre-stimulus spectral density.

Classifier input data therefore consisted of epochs of power spectral density or ranked spectral density within the 20–70 Hz band. The rank transform simply replaced the largest power frequency bin with the largest available rank continuing to the lowest power bin which was assigned rank 1. Each epoch was labelled according to stimulus type (left or right oblique). To predict the stimulus orientation (left or right oblique) based on the spectral characteristics of the individual trials we used a linear support vector machine (SVM, [Vapnik, 2009]) specifically the software package (http://svmlight.joachims.org/) by Joachims et al. [1999]. Briefly, based on the data labelled with two stimulus conditions, the SVM classifier defines a separating hyperplane from training samples by maximizing the margin of separation (i.e. the distance of the nearest sample to the separating hyperplane) between two conditions. This optimally trained classifier is used to predict the condition membership of other unlabelled data samples. Significantly higher than chance accuracy of this prediction indicates the discriminability of recorded signal between two stimulus conditions in the multidimensional pattern space.

There were 160 recorded epochs of data giving approximately 80 epochs of each stimulus type. In the cases of unequal numbers of stimulus type (due to the random ordering) the largest possible equal number of trials from each condition was used. These trials were divided into 10

Figure 3.

Classification accuracy as a function of the number of features (frequency bins) used for classification in both power-based (blue/diamonds) and ranked (green/squares) post-transient spectra with corresponding 95% confidence intervals (dotted) for hemisphere 1. In this case, the maximum accuracy for classification of the power spectrum is reached after 22 features at 72% whereas for rank it is reached at 24 features at 68% (compare Table I).

non-overlapping groups, each containing a maximum of eight epochs of a given stimulus orientation. To estimate classifier accuracy, nine groups were used to train and the final group was used to test, and the training and test groups were rotated (giving 10 test groups in total). The order in which the data entered the classifier had an effect on the peak achievable accuracy. At any iteration, the order in which the features (frequency bins) entered the classifier was determined by the training set (i.e. 90% of the data). The features were ordered according to the maximum magnitude of the combined (left and right oblique trials) average (across trials) training set. For example this would mean that if the average power spectra of both stimulus types across epochs peaked at 30 Hz, the 30-Hz bin would be the first classifier feature. The ordering of the input essentially corresponds to a feature selection step, which in our case is based on the criterion of reactivity of the spectrum to the grating stimuli. The asymptotic (all feature accuracy), which was independent of ordering, was also recorded. Accuracy values were then averaged across all ten test groups to get an estimate of mean classification accuracy (based on 160 test trials) as a function of the number of features used.

To test how significantly different the percent correct classification was from chance, the data epochs were randomly shuffled (meaning that they were assigned incorrect labels) 100 times, the training and testing repeated, and a null distribution of mean classification accuracy created

(the dotted lines in Fig. 3 indicate the 95th percentile of the null distributions).

Finally, we examined whether stimulus onset transients (0–300 ms) could be classified according to the orientation of the stimulus viewed. We did this in both spectral and time-domains. The spectral methods were as above except using a shorter length data segment (NFFT = 180). The time-series approach involved replacing the single trial spectra used above with single trial filtered (20–70 Hz) and dc corrected (based on pre-stimulus time) amplitude time series. The classifier input order was again determined by maximum absolute amplitude and asymptotic (order independent) classifier performance was computed. Training and testing of the classifiers, calculation of their mean accuracy and randomization procedures to obtain the null distributions were identical to that undertaken for classification of the spectral power and spectral rank.

RESULTS

Source Estimation

In all subjects electrical power change peaked in the contralateral visual cortex for all stimulus contrasts. Figure 1B shows for one subject the sagittal slice containing the peak voxel when contrasting all epochs (combined condition) to baseline. The difference in location between the image peaks due to either stimulus in isolation (e.g. left vs.

Figure 4.

The features (spectral bins in **A**, time points in **B**) used to classify each subject's data (until peak accuracy was reached). The number of features corresponds to the values given in Tables I, 3 for A and B, respectively. The colour of each bin represents and F test between conditions (normalised to I for each dataset). Features not used by the classifier are set to zero. A: Note that each pair of subsequent hemispheres are from the same subjects and this can be broadly seen from where the spectral difference between stimuli is expressed (see subject 2, hemispheres 3,4 where the frequency drops). Note also however that the spectral features used to make the classification, although in broadly similar bands, are by no means identical across hemispheres. B: Time samples used for the single trial evoked response classification. Note that these tend to cluster between 50–150 ms.

TABLE I. Accuracy of classification for post-transient (300–2,300 ms, 20–70 Hz) power (left) and rank spectra (right) at optimal number of features (max accurate) with corresponding significance level (P_{max}) from randomisation testing and number of features (spectral bins) used

Hem	Power spectrum, post transient					Rank spectrum, post transient				
	Max accurate	P_{max}	N. Features (of 99)	Asym' accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 99)	Asym accurate	P_{asym}
1	**72**	**0**	**22**	**68**	**0**	68	0	24	66	0
2	**73**	**0**	**31**	**70**	**0**	68	0	75	66	0
3	**67**	**0**	**56**	**62**	**0.0101**	56	0.0808	6	50	0.5152
4	**64**	**0**	**6**	56	0.1111	**61**	**0**	**1**	52	0.3434
5	**63**	**0**	**71**	**60**	**0.0202**	62	0	69	**59**	**0.0202**
6	55	0.1212	65	45	0.7172	54	0.2323	8	47	0.6768
7	**71**	**0**	**24**	**69**	**0**	**58**	**0.0303**	95	57	0.0606
8	**65**	**0**	**91**	**63**	**0**	59	0	4	58	0.0808
9	56	0.0505	**6**	52	0.3434	58	0.0606	27	56	0.1212
10	**64**	**0.0101**	**12**	59	0.0505	58	0.0707	10	48	0.7172
11	**59**	**0.0202**	**17**	58	0.0505	**63**	**0**	**22**	55	0.0808
12	**65**	**0**	**54**	**61**	**0.0101**	67	0	7	**65**	**0**
Summary	64.5	10/12	38	60.25	7/12	61	8/12	29	57	3/12

Significant ($P < 0.05$) classifications in bold font. Also shown are asymptotic (i.e. using all frequency bins) accuracy and significance levels (P_{asym}). Bottom row contains a summary statistic which is the mean value for accuracy and feature columns and the fraction of significant ($P < 0.05$) hemispheres for significance levels.

baseline) to the peak formed when they were combined (left and right vs. baseline) was on average 1.3 mm ($\sigma = 0.9$ mm); that is, typically a displacement of 1 voxel (1 mm) in two of the three dimensions. On average peak pseudo-t values ($\mu = 4.1$, $\sigma = 1.87$) in either condition were 97% ($\sigma = 14\%$) of the combined; there was no significant difference ($t = -0.56$, df = 11, $P = 0.58$) between conditions. In standard space the mean ($\pm$standard deviation) location (in mm) of the combined peak was $x = -11.5 \pm 13.7$, $y = -91.5 \pm 3.1$, $z = -11 \pm 10$ and $x = 18.5 \pm 7.7$, $y = -92 \pm 7.0$, $z = -11 \pm 8.2$ for left and right visual field stimulation respectively (i.e. was consistent with contra-lateral primary visual cortex). Mean full width half maximum (FWHM) at these voxels, as determined by correlations between neighbouring weight vectors [Barnes and Hillebrand, 2002], was 11.88 mm with standard deviation 2.83 mm.

Time Series

Figure 1C shows the characteristic (compare with invasive recordings in primate cortex [Rols et al., 2001]) time-frequency response to these stimuli at the combined peak location in one hemisphere of a representative subject. The colour code represents percent change from pre-stimulus baseline. The response starts with a phase-locked stimulus onset transient (the evoked response, see 1D) extending from 10–80 Hz and evolves into a sustained gamma oscillation which persists for the duration of the stimuli. The highlighted box indicates the time-frequency window used in all post-transient analysis.

The top panels in Figure 2A–C shows the mean spectra for the baseline and the post-onset transient sustained

gamma activity for both conditions projected through the same (combined) spatial filter for three subjects (hemispheres 1,12,3, respectively). Although the baseline spectra fall off with frequency the two stimuli produce slightly different modulations of the gamma band spectra. The lower panels show the ranked spectra for the same data. The rank spectra remove all amplitude variations and so only differences due to spectral shape are apparent. The rank spectra quantify the spectral shape of the induced response, which can be thought of as a measure of the relative reactivity of the signal spectrum to the stimulus across the frequency space [Hadjipapas et al., 2007].

In dataset C, for example, there is no difference in the rank spectra and so it is possible that the difference between the two stimulus conditions observed in the top panel is due purely to one stimulus giving a larger gamma band response (peaking at the same characteristic frequencies). In datasets A,B however the two different stimuli give rise to rank spectra with distinct deviations from the monotonically decreasing baseline. The differences between rank spectra suggest that the amplitude differences observed (top panel) are in part due to a change in the shape of the spectral response.

Classification

Post-transient

We attempted to classify single trial epochs (as left or right oblique stimuli) based on single trial post-transient (300–2,300 ms) power and rank spectra. Figure 3 shows the mean classification accuracy of single trial power (blue

TABLE II. Maximum and asymptotic (all feature) accuracy with significance levels of classification for transient (0-300ms,20–70Hz) spectra power (left) and rank (right) spectra

Hem	Power spectrum, transient					Rank spectrum, transient				
	Max accurate	P_{max}	N. Features (of 15)	Asym accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 15)	Asym accurate	P_{asym}
1	53	0.1010	4	52	0.3838	**55**	**0.0303**	**3**	49	0.5859
2	**60**	**0**	**4**	54	0.1818	53	0.1515	2	52	0.3333
3	50	0.3030	1	45	0.7980	51	0.3030	5	43	0.9091
4	**66**	**0**	**4**	58	0.0505	**59**	**0**	**5**	57	0.0606
5	**66**	**0**	**10**	**64**	**0**	54	0.1111	3	51	0.3737
6	52	0.1818	1	46	0.7980	**57**	**0.0404**	**4**	42	0.9293
7	**60**	**0.0101**	**8**	58	0.0606	53	0.0909	1	52	0.3434
8	50	0.4949	2	43	0.8788	55	0.0808	2	42	0.8788
9	58	0.0606	12	58	0.0606	57	0.1010	6	53	0.3030
10	**62**	**0**	**10**	56	0.0707	57	0.1515	13	57	0.1313
11	55	0.0606	1	45	0.7172	55	0.0707	1	44	0.8182
12	**62**	**0**	**7**	**58**	**0.0404**	55	0.0808	2	48	0.6970
Summary	57.8	6/12	5	53	2/12	55	3/12	4	49	0/12

Also shown is the number of frequency bins (features) required to reach maximum accuracy. Significant ($P < 0.05$) classifications in bold font. Bottom row contains summary statistic which is the mean value for accuracy and feature columns and the fraction of significant ($P < 0.05$) hemispheres for significance levels.

diamonds) and rank spectra (green squares) in a single dataset as a function of the number of frequency bins (features) used. Typically classification accuracy improves up until a certain number of features are input and then begins to fall off again to chance levels as the classifier becomes over-trained. The 95th percentile classifier accuracy when the stimulus labels were randomly assigned is indicated by dotted lines. The classification accuracy peaked significantly above chance for the spectral power and rank distributions in 10 and 8 of the 12 datasets,

respectively (Table I). The mean accuracies were 64.5% and 61.0% for power and rank respectively being significantly greater for power ($t = 2.3573$, df $= 11$, $P < 0.038$). The mean number of features (spectral bins) required to reach maximum accuracy were 38 and 29 for power and rank-power respectively, although this difference was not significant ($t = 0.689$, df $= 11$, $P = 0.504$). Given that each feature corresponded to a 0.5 Hz wide spectral bin this meant that both methods required on average approximately 17 Hz (one third of the 50 Hz bandwidth available)

TABLE III. Maximum and asymptotic (all feature) accuracy with significance levels of classification for transient (0–300 ms, 20–70 Hz) evoked response amplitude and ranked amplitude

Hem	Evoked response transient					Rank evoked response transient				
	Max accurate	P_{max}	N. features (of 180)	Asym accurate	P_{asym}	Max accurate	P_{max}	N. features (of 180)	Asym accurate	P_{asym}
1	**61**	**0.0202**	**25**	50	0.4646	**59**	**0**	**22**	52	0.3333
2	**59**	**0.0404**	**70**	53	0.3535	56	0.1414	56	54	0.2323
3	56	0.1414	83	46	0.7677	55	0.2020	16	47	0.7172
4	**68**	**0**	**7**	**62**	**0.0101**	**66**	**0**	**42**	56	0.2121
5	**60**	**0**	**6**	50	0.5354	53	0.0808	3	51	0.4343
6	**60**	**0**	**2**	52	0.3333	57	0.0707	9	50	0.4949
7	**63**	**0.0101**	**25**	56	0.1515	**59**	**0.0303**	**8**	52	0.2929
8	**59**	**0.0404**	**15**	51	0.4343	53	0.2626	20	49	0.5859
9	**64**	**0**	**34**	58	0.0707	**62**	**0.0101**	**127**	58	0.0505
10	**64**	**0**	**80**	53	0.2323	**65**	**0**	**24**	52	0.3131
11	**72**	**0**	**38**	**67**	**0**	**73**	**0**	**6**	**65**	**0**
12	**62**	**0.0101**	**93**	**58**	**0.0303**	**61**	**0.0202**	**8**	**60**	**0.0303**
Summary	62.3	11/12	40	55	3/12	60	7/12	28	53.8	2/12

Also shown is the number of time-samples (features) required to reach maximum accuracy. Significant ($P < 0.05$) classifications in bold font. Bottom row contains summary statistic which is the mean value for accuracy and feature columns and the fraction of significant ($P < 0.05$) hemispheres for significance levels.

to reach optimal classification. We were interested to find out which features the classifier was exploiting in order to reach maximum accuracy. Figure 4A shows the spectral bins which were used to reach maximal classification accuracy (for power) in each subject. The colour code of the bins (red max, blue min) depicts the univariate F statistic between stimulus conditions for each feature (normalised for each subject). For example for the dataset shown in Figure 3 (hem 1) 22 features were used to reach maximal classification accuracy (for power), the position of these features in the frequency spectrum however depends on the input order to the classifier (see methods). Figure 4 shows that these features (i.e. those leading up to the maximum in Fig. 3) were predominantly in the 40–50 Hz band with the 52 Hz bin showing the largest univariate difference between stimuli. Bins not assigned a colour were not used in the optimal classification (i.e. the addition of these features caused the classification accuracy to decline). Each subsequent pair of hemispheres belongs to the same subject. There is some variation in the frequency bins chosen for classification (which were based on the difference between combined stimulus conditions and baseline power) expected from gamma band variations in the general population [Muthukumaraswamy et al., 2009]. In all but one subject (hemispheres 3,4) features within the 40–60 Hz band were used and typically showed the largest mean difference between stimulus conditions. Also, data from any pair of hemispheres appears to show such spectral differences at broadly similar frequencies. However it is worth noting that in any pair of hemispheres the differences between features are not identical; i.e. the spectral signature of the gamma change is specific to that location and not to the individual.

Using the complete bandwidth (P_{asym}, Table I), it was possible to classify 7 out 12 datasets using power spectra and only 3 out 12 with rank spectra; suggesting that at a certain point extra frequency bins (i.e. the white space in Fig. 4) contribute only noise to the classification.

Given the apparent differences (in Fig. 2) we attempted to distinguish between conditions on an epoch by epoch basis by using common univariate spectral metrics of peak frequency, peak power deviation from baseline and total power deviation from baseline. We used the same classifier framework but with a single rather than multiple features per epoch. At best it was possible to classify at maximum one of the 12 datasets at above chance levels. The results are shown in Table AI, Appendix. This is consistent with our aim to remove any stimulus related tuning effects (such as one stimulus causing a larger power/frequency of response).

Transient Power

We went on to examine whether similar amounts of information were contained in the stimulus transient by training new classifiers on both spectral (power and rank-power) and time-amplitude (evoked response) single trials

in the same 20–70 Hz band but crucially using data from the first 300 ms only (see Methods). Note that as there is much less data (300 ms instead of 2 s) in these transient periods the results are not directly comparable with the post-transient analysis.

The results for the stimulus transient classifiers based on spectral power and rank-power are shown in Table II. In this case with only 6 and 3 of 12 hemispheres respectively for power (mean = 57.8%) and rank (mean = 49%) classifying above chance. Both methods required on average five features (at 3.33 Hz per bin approximately 17 Hz) to reach maximum accuracy.

Transient Time-Series

The results of the classifiers based on the filtered (20–70 Hz) single trial evoked responses (in the first 300 ms) are shown in Table III. Significant classification was possible for 11 and 7 out of 12 datasets using amplitude and ranked time-samples, respectively; with corresponding overall mean accuracies of 62.3% (amplitude) and 60% (rank). The improvement in using amplitude rather than ranked time-series data to classify was significant ($t = 3.44$, df = 11, $P < 0.005$). The average number of samples (out of a possible 180) to reach maximum accuracy for amplitude and rank were 40 and 28 respectively, and these were not significantly different ($t = 0.8336$, df = 11, $P = 0.433$). This corresponded to approximately 60 ms of the total 300 ms. In summary, for the transient (0–300 ms) portion of the response, using both amplitude and phase information (in the single trial evoked response) produced significantly greater classifier accuracy than the use of amplitude information alone (in the form of spectral power) ($t = 4.2$, df = 11, $P < 0.0015$).

Figure 4B shows the features (times samples) used in the optimal classification of the single trial evoked response (amplitude). Colour codes show maximal univariate F statistic differences in averaged evoked response amplitude at this latency (normalised for each subject). Note that these samples, based on the difference between the combined average evoked response and baseline, cluster in between 50 and 150 ms. On the basis of this figure we were interested whether similar classification accuracy could be achieved using only the first 100 ms of single trial evoked response data. Indeed (Table AIV) it was possible to distinguish between the single trial responses to either grating significantly above chance in 10 and 8 of the 12 hemispheres for amplitude and ranked data, respectively.

DISCUSSION

We present two main findings. The first is that we can use MEG to probe two spatially overlapping cortical networks through the use of multivariate classifiers in both single trial induced and evoked responses. The second is that the orientation of static visual stimuli can be

distinguished by the induced sustained gamma oscillation (300–2,500 ms) in the absence of any onset transient (0–300 ms).

We confirm that the macroscopic aggregate MEG signals arising from a constrained cortical area can be further broken down to classify different microscopic sub-populations of neurons within that area as suggested by theoretical predictions [Hadjipapas et al., 2009]. This is true not only for the evoked response but also for on-going oscillations. We believe that this could represent an exciting technological step, as by using the information in the time-series we were able to distinguish between two spatially coincident populations of neurons. That, is better exploiting the temporal information in the virtual sensor by using multivariate methods, one can improve the discrimination capability of the MEG measurement beyond source localization by identifying functionally diverse behaviours (distinct response modes of the underlying network) within the same source signal. This approach can be seen as a temporal analogue of approaches used in recent fMRI studies [Haynes and Rees, 2005; Kamitani and Tong, 2005] where stimulus classification is achieved by measuring spatial amplitude changes due to non-uniform spatial sampling of (rather than non-uniform network dynamics between) orientation columns. We should emphasize at this point that such an approach applied to MEG/EEG data has potential to provide insights into the underlying cortical physiology. Importantly, using ranked spectral bins enabled us to confirm that our findings were due to a spectral shape, and not just magnitude, change [Hadjipapas et al., 2007]. Although, we note that the classification was consistently more powerful when the magnitude information was used. Given that there may be a number of ways that network dynamics may differ in response to varying stimulation and also due to the complex relationship between network dynamics and observed aggregate [Hadjipapas et al., 2009], it is also expected that there may be a number of candidate temporal features of the observed aggregate signals, which in turn, may provide richer or poorer descriptors of the underlying network state. Two approaches lend themselves to inform the question of which the relevant temporal features of the aggregate signal may be. Firstly, one could derive a theoretical expectation regarding the differentiation of dynamics in the network response and additionally, derive an expectation regarding which features of the aggregate signal may best capture such dynamics. This approach would require a detailed large-scale biophysical model and also a detailed model of the measurement function relating network dynamics to the observed aggregate signal [Friston et al., 2003; Hadjipapas et al., 2009; Moran et al., 2007, 2009]. Both steps entailed in such an approach are rather challenging, as the amount of detailed knowledge of the biophysics required is still lacking at the large scales investigated here but perhaps such theoretical predictions may become feasible (and would then be immensely informative) in the future. Alternatively, one could utilize a more empirical approach and aim to identify temporal features of the measured aggregate MEG/EEG signal that allow for the most robust discrimination between different response modes in the same local network. Here, we start to attempt this by using a number of different candidate temporal features (power in gamma band, peak frequency, spectral pattern, rank spectrum, phase-locked and non-phase locked, transient or sustained components etc.) and in such way gain more insights in what the relevant temporal signal features may be. Clearly, in this paper we have by far not exhausted the possibilities but in contrast, merely have taken a first step in this direction.

Of interest for the near future is the optimization of network descriptors using higher-order statistical and non-linear properties. Of particular relevance is perhaps the estimation of a phase variable of the gamma responses: this does represent a challenge as gamma responses are distinctly broadband phenomena, i.e. these are not necessarily akin to the classical concept of oscillations characterised by a single peak frequency, and therefore the extraction of a (generalised) phase variable is non-trivial.

We also note that the relatively modest classification accuracies reported here could indeed be improved upon if classification were the only aim. However, the main goal in this article was to elucidate functional aspects of the post-transient sustained gamma oscillation rather than achieve maximal stimulus classification. Furthermore, in the context of our analysis we should note that improvements in the classification will come from identification of the relevant characteristic time-frequency differences between conditions. We should stress that the inclusion of more data does not necessarily help the classification (compare p_asym to p_max in the tables). We also tried including all data (0–0.3 s, 20–70 Hz), but this had little effect and even degraded the accuracy slightly (Table AII). This is related to the issue of a selection of an appropriately sparse dataset for training (see [Ganesh et al., 2008; Yamashita et al., 2008] for a discussion of the fMRI case). Essentially, if we increase the number of features fed to the classifier, we run into the danger of overfitting the training data: the classifier focuses on insignificant parts of the data and will not generalise very well to unseen test data. Analogously, if the input data is has too few dimensions, there is the danger of under fitting (that is the classifier becomes too general and can no longer recover the true mapping between data and labels). This trade off can be seen in Fig. 3. In fact, the most promising results were obtained when we split the post-transient data into low (20–45 Hz) and high (45–70 Hz) gamma, Tables AIII, AIV, respectively. These data showed that using high-gamma features significant classification could be achieved for both power and ranked data in 10/12 hemispheres (i.e. an improvement on the broadband data). Therefore, the feature selection step is critical and in this case rather conservative (as we simply ordered features according to mean departure from baseline).

To what extent could our results have been due to one of the stimuli activating another cortical area whose activity inadvertently mixed in with the signal at this location of interest? All MEG source inversions rely on estimating electrical activity through a weighted contribution of a finite number of sensors. The spatial resolution of the beamformer is high (due to its adaptive weights) compared to other techniques based on weights determined purely by geometry [Hillebrand et al., 2005]. In our data we verified that the contrast of either condition (for two well-matched stimuli for which the mapping to the cortex is well understood) alone (versus baseline) produced peaks in the same location of visual cortex. We are therefore able to say with some confidence that we are classifying the stimuli based purely on temporal effects, rather than mixing of signals from other area of cortex. However, there is certainly a caveat here for future work in that if one of the stimuli gave rise to activity in a different area of cortex this could be inadvertently mixed (via the weight vectors) into the data. Indeed this was one of the reasons we were keen to confirm findings from previous work which used full-field stimuli [Hadjipapas et al., 2007]. One key point here is that there is no reason why the classification and the beamformer image stage should not be combined. This would involve simply substituting the beamformers univariate test for total power change with a multivariate test to look for combinations of frequency bins which covary with the experimental design (similar to framework proposed by Soto et al., [2009]). This would help directly establish whether differences between stimuli existed at any other cortical location.

Although the original motivation for this paper was to test whether post-transient oscillations differed spectrally we note that the best classification accuracy was obtained from the single-trial evoked response data (Fig. 4b). A large body of literature exists showing that the onset transient is information rich [Buracas et al., 1998; Celebrini et al., 1993; Heller et al., 1995; Muller et al., 2001; Panzeri et al., 2001; Reich et al., 2001; VanRullen and Thorpe, 2002]. Celebrini et al. [1993] show that the orientation selectivity at the level of cortical neurons in the awake primate is developed at the start of the neuronal response (varying from 40 to 150 ms) but typically at 50–70 ms. Similarly Muller et al. [2001] show that the orientation selectivity of a neurons in anesthetized primates is fully developed by 150 ms. Our results are broadly consistent with these findings with discrimination between stimuli possible, with both ranked and unranked data, as early as the first 100 ms in 10/12 datasets (Table AV). Analagous with the gamma band explanation we would attribute these subtle multivariate differences in the evoked response (mean field) to different populations of contributing neurons (with different optimal orientation specificity).

Recent EEG studies have shown that saccadic eye movements also give rise to electrical activity in the gamma range [Yuval-Greenberg, 2008, 2009]. In MEG, which is of course reference free, this confusion does not arise; the basic (left-right) retinotopy and the similarity of the time-

frequency plots to invasive primate recordings in V1 [Gail et al., 2000; Rols et al., 2001] are corroborating factors. Also, the gamma activity observed here is of sustained (rather than clearly transient) nature as in the case of ocular EMG due to saccades shown by Yuval Greenberg and colleagues. Here, we did not measure eye movements and therefore we could not exclude the possibility that the gamma activity indeed could be indirectly influenced by micro-saccades in the sense that these could modulate the gamma amplitude by means of refreshing the visual input, as suggested by the recent invasive findings in [Bosman et al., 2009].

The presence of static stimulus related information in the post-transient oscillation was initially surprising to us as most literature on invasive recordings at the neuronal level suggests the stimulus information is contained solely in the earliest parts of the neuronal response [Buracas et al., 1998; Celebrini et al., 1993; Heller et al., 1995; Muller et al., 2001; Panzeri et al., 2001; Reich et al., 2001; Van-Rullen and Thorpe, 2002]. Our results are in apparent contrast with those of [Gray and Singer, 1989] who report that the peak frequency of LFP oscillation and oscillation frequency of single cells in the cat did not change significantly as a result of changing stimulus orientation (or direction of stimulus movement). Furthermore, [Gray et al., 1990] showed that in 13 out of 18 single cell recordings (exhibiting oscillatory responses) in the cat there was no difference (>4 Hz) in oscillation frequency for stimuli of differing orientation. This observation would however be quite consistent with our findings, as we are suggesting that the changes in spectra are due to different collective response modes in networks of interacting neurons, rather than the same set of neurons changing their frequency. Indeed, studies in the macaque Frien and Eckhorn [2000a] and Frien et al. [2000b] found that peak LFP (and multi-unit activity (MUA)) frequency was dependent on moving stimulus orientation in the high (gamma) frequencies. The aforementioned studies showed a small ($\sim$1 Hz) but significant difference in the peak gamma frequencies between an optimal orientation (determined by a previous characterization of the receptive field of the recording site on the basis of peristimulus histograms of the MUA) and its orthogonal orientation. Interestingly, in Frien et al. [2000b], the most pronounced difference in average peak frequencies (across trials) was seen in the sustained gamma band responses, while this difference was much smaller during the stimulus onset transient. Although invasively recorded LFP allows characterization of different populations of neurons at a millimetre scale, it is possible that the considerable spatial summation of many such sub-networks giving rise to the macroscopic MEG/EEG (we estimate that our stimulus subtends an area of 1–2 cm^2 in primary visual cortex) will inevitably result in a broader spectrum in which the information is contained at many rather than a single frequency component. Indeed in this study we show that although univariate spectral measures do not distinguish between stimulus features at this scale, a

multivariate approach quantifying the spectral shape (across a bandwidth of approximately 17 Hz) can be used to successfully predict stimulus orientation above chance at a single-trial level. Why this stimulus specific information should be preserved at a network level long after the initial transient remains an exciting question.

ACKNOWLEDGMENTS

REFERENCES

Adjamian P, Hadjipapas A, Barnes GR, Hillebrand A, Holliday IE (2008): Induced Gamma activity in primary visual cortex is related to luminance and not color contrast: An MEG study. J Vis 8:4–7.

Adjamian P, Holliday IE, Barnes GR, Hillebrand A, Hadjipapas A, Singh KD (2004): Induced visual illusions and gamma oscillations in human primary visual cortex. Eur J Neurosci 20:587–592.

Barnes GR, Hillebrand A, (2003): Statistical flattening of MEG beamformer images. Hum Brain Mapp 18:1–12.

Bosman CA, Womelsdorf T, Desimone R, Fries P (2009): A microsaccadic rhythm modulates gamma-band synchronization and behavior. J Neurosci 29:9471–9480.

Buracas GT, Zador AM, DeWeese MR, Albright TD (1998): Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. Neuron 20:959–969.

Celebrini S, Thorpe S, Trotter Y, Imbert M (1993): Dynamics of orientation coding in area V1 of the awake primate. Visual Neurosc 10:811–825.

Eckhorn R, Frien A, Bauer R, Woelbern T, Kehr H (1993): High frequency (60–90 Hz) oscillations in primary visual cortex of awake monkey. Neuroreport 4:243–246.

Friston KJ, Harrison L, Penny W (2003): Dynamic causal modelling. Neuroimage 19:1273–1302.

Gail A, Brinksmeyer HJ, Eckhorn R (2000): Contour decouples gamma activity across texture representation in monkey striate cortex. Cereb Cortex 10:840–850.

Ganesh G, Burdet E, Haruno M, Kawato M (2008): Sparse linear regression for reconstructing muscle activity from human cortical fMRI. Neuroimage 42:1463–1472.

Gray CM, Engel AK, Konig P, Singer W (1990): Stimulus-dependent neuronal oscillations in cat visual cortex: Receptive field properties and feature dependence. Eur J Neurosci 2:607–619.

Gray CM, Singer W (1989): Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc Natl Acad Sci USA 86:1698–1702.

Hadjipapas A, Adjamian P, Swettenham JB, Holliday IE, Barnes GR (2007): Stimuli of varying spatial scale induce gamma activity with distinct temporal characteristics in human visual cortex. Neuroimage 35:518–530.

Hadjipapas A, Casagrande E, Nevado A, Barnes GR, Green G, Holliday IE (2009): Can we observe collective neuronal activity from macroscopic aggregate signals? Neuroimage 44:1290–1303.

Hall SD, Holliday IE, Hillebrand A, Singh KD, Furlong PL, Hadjipapas A, Barnes GR (2005): The missing link: Analogous human and primate cortical gamma oscillations. Neuroimage 26:13–17.

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

Heller J, Hertz JA, Kjaer TW, Richmond BJ (1995): Information flow and temporal coding in primate pattern vision. J Comput Neurosci 2:175–193.

Henrie JA, Shapley R (2005): LFP power spectra in V1 cortex: The graded effect of stimulus contrast. J Neurophysiol 94:479–490.

Hillebrand A, Singh KD, Holliday IE, Furlong PL, Barnes GR (2005): A new approach to neuroimaging with magnetoencephalography. Hum Brain Mapp 25:199–211.

Hubel DH, Wiesel TN (1974): Sequence regularity and geometry of orientation columns in the monkey striate cortex. J Comp Neurol 158:267–293.

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

Leopold DA, Murayama Y, Logothetis NK (2003): Very slow activity fluctuations in monkey visual cortex: Implications for functional brain imaging. Cereb Cortex 13:422–433.

McLoughlin N, Schiessl I (2006): Orientation selectivity in the common marmoset (Callithrix jacchus): The periodicity of orientation columns in V1 and V2. Neuroimage 31:76–85.

Moran RJ, Kiebel SJ, Stephan KE, Reilly RB, Daunizeau J, Friston KJ (2007): A neural mass model of spectral responses in electrophysiology. Neuroimage 37:706–720.

Moran RJ, Stephan KE, Seidenbecher T, Pape HC, Dolan RJ, Friston KJ (2009): Dynamic causal models of steady-state responses. Neuroimage 44:796–811.

Muller JR, Metha AB, Krauskopf J, Lennie P (2001): Information conveyed by onset transients in responses of striate cortical neurons. J Neurosci 21:6978–6990.

Muthukumaraswamy SD, Singh KD (2008): Functional decoupling of BOLD and gamma-band amplitudes in human primary visual cortex. Hum Brain Mapp.

Muthukumaraswamy SD, Edden RA, Jones DK, Swettenham JB, Singh KD (2009): Resting GABA concentration predicts peak gamma frequency and fMRI amplitude in response to visual stimulation in humans. Proc Natl Acad Sci USA 106:8356–8361. (Epub ahead of print).

Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME (2001): The role of spike timing in the coding of stimulus location in rat somatosensory cortex. Neuron 29:769–777.

Reich DS, Mechler F, Victor JD (2001): Temporal coding of contrast in primary visual cortex: When, what, and why. J Neurophysiol 85:1039–1050.

Robinson SEVJ (1999): Functional Neuroimaging by Synthetic Aperture Magnetometry (SAM). Recent Advances in Biomagnetism. Sendai: Tohoku University Press, pp 302–305.

Rols G, Tallon-Baudry C, Girard P, Bertrand O, Bullier J (2001): Cortical mapping of gamma oscillations in areas V1 and V4 of the macaque monkey. Visual Neurosci 18:527–540.

Schira MM, Tyler CW, Breakspear M, Spehar B (2009): The foveal confluence in human visual cortex. J Neurosci 29: 9050–9058.

Schira MM, Wade AR, Tyler CW (2007): Two-dimensional mapping of the central and parafoveal visual field to human visual cortex. J Neurophysiol 97:4284–4295.

Sekihara K, Nagarajan SS, Poeppel D, Marantz A (2004): Asymptotic SNR of scalar and vector minimum-variance beamformers for neuromagnetic source reconstruction. IEEE Trans Biomed Eng 51:1726–1734.

Siegel M, Donner TH, Oostenveld R, Fries P, Engel AK (2007): High-frequency activity in human visual cortex is modulated by visual motion strength. Cereb Cortex 17:732–741.

Soto JL, Pantazis D, Jerbi K, Lachaux JP, Garnero L, Leahy RM (2009): Detection of event-related modulations of oscillatory brain activity with multivariate statistical analysis of MEG data. Hum Brain Mapp 30:1922–1934.

VanRullen R, Thorpe SJ (2002): Surfing a spike wave down the ventral stream. Vis Res 42:2593–2615.

VanVeen BD, vanDrongelen W, Yuchtman M, Suzuki A (1997): Socalization of brain electrical activity via linearly constrained minimum variance spatial filtering. IEEE Trans Biomed Eng 44:867–880.

Vapnik (2009): The Nature of Statistical Learning. New York: Springer Verlag.

Yamashita O, Sato MA, Yoshioka T, Tong F, Kamitani Y (2008): Sparse estimation automatically selects voxels relevant for the decoding of fMRI activity patterns. Neuroimage 42:1414–1429.

Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY (2008): Transient induced gamma-band response in EEG as a manifestation of miniature saccades. Neuron 58:429–441.

Yuval-Greenberg S, Deouell LY (2009): The broadband-transient induced gamma-band response in scalp EEG reflects the execution of saccades. Brain Topogr 22:3–6. (Epub ahead of print).

APPENDIX

TABLE AI. Classification results using a single classifier feature based on univariate properties of the signal

Hem	Peak frequency	Peak power	Mean power
1	0.9899	**0**	0.7980
2	**0.0404**	0.7172	0.2121
3	0.7980	0.3232	0.6162
4	0.8990	1.0000	1.0000
5	0.6061	0.1010	0.2525
6	0.3131	0.3636	0.1010
7	0.3434	0.1616	1.0000
8	0.8081	0.2828	0.3030
9	0.1111	0.3232	0.3434
10	0.2020	0.7374	0.8081
11	0.4141	0.1212	0.3737
12	0.6061	0.1414	0.8485
Summary	1/12	1/12	0/12

The table shows significance levels for classification based on peak frequency, peak power and total power (deviation from baseline) of the post-transient gamma spectrum (20–70 Hz, 300–2,300 ms) across epochs in each dataset. Significance levels of $P < 0.05$ are in bold font. The summary row gives the total number of hemispheres in which classification performance significantly ($P < 0.05$) above chance was found. Note that in general (at best 1/12 cases) one is not able to discriminate between gamma spectra using these basic univariate descriptors.

TABLE AII. Accuracy of classification for post-stimulus (0–2,300 ms, 20–70 Hz) power (left) and rank spectra (right) at optimal number of features (max accurate) with corresponding significance level (P_{max}) from randomisation testing and number of features (spectral bins) used

	Power spectrum, all data					Rank spectrum, all data				
Hem	Max accurate	P_{max}	N. Features (of 99)	Asym' accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 99)	Asym accurate	P_{asym}
1	**64**	**0.0101**	**24**	**62**	**0.0101**	**65**	**0**	**50**	57	0.0808
2	**73**	**0**	**19**	**72**	**0**	**59**	**0**	**99**	**59**	**0**
3	**60**	**0.0202**	**97**	**60**	**0.0202**	**58**	**0.0303**	**28**	53	0.2626
4	57	0.0909	14	54	0.2020	55	0.0909	11	53	0.2121
5	**73**	**0**	**63**	**68**	**0**	**60**	**0.0303**	**99**	**60**	**0.0303**
6	58	0.0505	60	53	0.2727	58	0.0808	55	56	0.1414
7	**69**	**0**	**16**	**66**	**0**	**61**	**0.0101**	**99**	**61**	**0.0101**
8	**67**	**0**	**6**	61	0.0202	**60**	**0.0202**	**6**	53	0.3030
9	**61**	**0.0101**	**14**	54	0.2222	52	0.2222	3	49	0.5556

TABLE AII. (Continued)

	Power spectrum, all data					Rank spectrum, all data				
Hem	Max accurate	P_{max}	N. Features (of 99)	Asym' accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 99)	Asym accurate	P_{asym}
10	**71**	**0**	**72**	**67**	**0**	**63**	**0**	**14**	**61**	**0.0202**
11	57	0.0606	30	50	0.4848	55	0.1616	16	51	0.3333
12	**65**	**0**	**41**	**59**	**0.0303**	**71**	**0**	**42**	**68**	**0**
Summary	64.5	9/12	38	60.5	8/12	59.75	8/12	43.5	56.75	3/12

Significant ($P < 0.05$) classifications in bold font. Also shown are asymptotic (i.e. using all frequency bins) accuracy and significance levels (P_{asym}). Bottom row contains summary statistic which is the mean value for accuracy and feature columns and the fraction of significant ($P < 0.05$) hemispheres for significance levels. Note that generally performance degraded from when using spectral estimates based on purely post-transient data (Table I).

TABLE AIII. Identical analysis on post-transient spectra (300–2,300 ms) to that of Table I except using only the low (20–45 Hz) gamma band only

	Power spectrum, post-transient 20–45 Hz					Rank spectrum, post-transient 20–45 Hz				
Hem	Max accurate	P_{max}	N. Features (of 49)	Asym' accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 49)	Asym accurate	P_{asym}
1	**59**	**0**	**14**	54	0.2626	**57**	**0**	**4**	49	0.5253
2	**73**	**0**	**6**	**71**	**0**	**61**	**0**	**20**	**61**	**0**
3	**66**	**0**	**18**	**62**	**0**	**58**	**0**	**1**	49	0.5657
4	**64**	**0**	**6**	54	0.1818	**59**	**0**	**6**	55	0.1515
5	56	0.1010	40	55	0.1313	**60**	**0.0202**	**48**	**59**	**0.0303**
6	54	0.0808	2	46	0.7778	52	0.2121	2	42	0.9293
7	**67**	**0**	**25**	**65**	**0**	51	0.3333	5	40	0.9697
8	**62**	**0**	**6**	57	0.0505	**58**	**0**	**6**	53	0.3333
9	51	0.4646	43	49	0.5859	54	0.2323	8	51	0.4343
10	**62**	**0**	**10**	56	0.1717	**60**	**0.0101**	**35**	58	0.0707
11	55	0.1111	9	51	0.3535	53	0.2222	8	51	0.4040
12	**62**	**0**	**32**	**59**	**0.0202**	**66**	**0**	**10**	**62**	**0**
Summary	60.9	8/12	17.6	56.6	4/12	57.4	8/12	12.75	52.5	3/12

Note that the performance is slightly degraded.
Significant classifications (P<0.05) are in BOLD font.

TABLE AIV. Identical analysis on post-transient spectra (300–2,300 ms) to that of Table I except using only the high (45–70 Hz) gamma band only

	Power spectrum, post-transient 45–70 Hz					Rank spectrum, post-transient 45–70 Hz				
Hem	Max accurate	P_{max}	N. Features (of 49)	Asym' accurate	P_{asym}	Max accurate	P_{max}	N. Features (of 49)	Asym accurate	P_{asym}
1	**71**	**0**	**15**	**67**	**0**	**67**	**0**	**32**	**66**	**0**
2	**73**	**0**	**16**	**71**	**0**	**61**	**0.0202**	**6**	57	0.0606
3	55	0.1212	10	51	0.4040	**57**	**0.0303**	**3**	52	0.3636
4	55	0.0909	5	44	0.8485	**59**	**0.0202**	**2**	51	0.4848
5	**65**	**0**	**8**	**61**	**0**	58	0.0606	49	58	0.0606
6	**55**	**0.0202**	**3**	50	0.4949	**56**	**0.0101**	**1**	52	0.3030
7	**69**	**0**	**6**	**65**	**0**	**59**	**0.0202**	**27**	55	0.1717
8	**65**	**0**	**21**	**65**	**0**	**56**	**0**	**2**	54	0.1919
9	**61**	**0.0303**	**49**	**61**	**0.0303**	**59**	**0.0202**	**20**	58	0.0707
10	**64**	**0**	**12**	**59**	**0.0404**	**57**	**0**	**6**	50	0.4444
11	**66**	**0**	**26**	**63**	**0**	**65**	**0**	**49**	**65**	**0**
12	**59**	**0**	**8**	54	0.1919	58	0.0505	37	56	0.0808
Summary	63.2	10/12	14.9	59.3	8/12	59.3	10/12	19.5	56.2	2/12

Note that the performance is marginally enhanced with significant classification for 10/12 (as opposed to 8/12) rank spectra.
Significant classifications (P<0.05) are in BOLD font.

TABLE AV. The same analysis as for Table III but using data from only the first 100 ms of the evoked response transient

Hem	Evoked response transient 0–100 ms					Rank evoked response transient, 0–100 ms				
	Max accurate	P_{max}	N. features (of 62)	Asym accurate	P_{asym}	Max accurate	P_{max}	N. features (of 62)	Asym accurate	P_{asym}
1	58	0.0606	47	54	0.2222	**59**	**0**	**9**	51	0.3535
2	**63**	**0**	**39**	**60**	**0.0202**	**61**	**0.0202**	**36**	58	0.0707
3	52	0.2323	3	47	0.6364	53	0.2727	26	52	0.3535
4	**68**	**0**	**7**	**60**	**0.0404**	**63**	**0**	**6**	**58**	**0.0404**
5	**60**	**0.0202**	6	49	0.6263	56	0.0505	6	49	.5455
6	**60**	**0**	**2**	55	0.2121	58	0.0707	7	56	0.1616
7	**61**	**0.0202**	14	50	0.4949	**62**	**0**	**54**	**57**	**0.0303**
8	**59**	**0.0303**	15	51	0.3333	55	0.1010	17	48	0.5859
9	**62**	**0.0101**	43	58	0.0606	**63**	**0.0101**	**31**	**62**	**0**
10	**64**	**0**	**18**	**62**	**0**	**62**	**0.0101**	**14**	**61**	**0**
11	**63**	**0.0101**	12	58	0.0707	**66**	**0**	**26**	58	0.1212
12	**61**	**0.0101**	16	**60**	**0.0202**	**61**	**0.0202**	**8**	55	0.1818
Summary	60.92	10/12	18.5	55.33	4/12	59.9	8/12	20	55.4	4/12

Table shows maximum and asymptotic (all feature) accuracy with significance levels of classification for transient (0–100 ms, 20–70 Hz) evoked response amplitude and ranked amplitude.
Significant classifications (P<0.05) are in BOLD font.