

Scanning for oscillations

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2015 J. Neural Eng. 12 066020

(<http://iopscience.iop.org/1741-2552/12/6/066020>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 213.174.99.137

This content was downloaded on 26/10/2015 at 15:51

Please note that [terms and conditions apply](#).

Scanning for oscillations

Alain de Cheveigné^{1,2,3} and Dorothée Arzounian^{1,2}

¹Laboratoire des Systèmes Perceptifs, UMR 8248, CNRS, France

²Département d'Études Cognitives, Ecole Normale Supérieure, PSL* Research University, France

³UCL Ear Institute, UK

E-mail: Alain.de.Cheveigne@ens.fr

Received 23 June 2015, revised 24 August 2015

Accepted for publication 16 September 2015

Published 27 October 2015



CrossMark

Abstract

Objective. Oscillations are an important aspect of brain activity, but they often have a low signal-to-noise ratio (SNR) due to source-to-electrode mixing with competing brain activity and noise. Filtering can improve the SNR of narrowband signals, but it introduces ringing effects that may masquerade as genuine oscillations, leading to uncertainty as to the true oscillatory nature of the phenomena. Likewise, time–frequency analysis kernels have a temporal extent that blurs the time course of narrowband activity, introducing uncertainty as to timing and causal relations between events and/or frequency bands. **Approach.** Here, we propose a methodology that reveals narrowband activity within multichannel data such as electroencephalography, magnetoencephalography, electrocorticography or local field potential. The method exploits the between-channel correlation structure of the data to suppress competing sources by joint diagonalization of the covariance matrices of narrowband filtered and unfiltered data. **Main results.** Applied to synthetic and real data, the method effectively extracts narrowband components at unfavorable SNR. **Significance.** Oscillatory components of brain activity, including weak sources that are hard or impossible to observe using standard methods, can be detected and their time course plotted accurately. The method avoids the temporal artifacts of standard filtering and time–frequency analysis methods with which it remains complementary.

Keywords: EEG, brain oscillations, ICA, DSS, JD, joint decorrelation, MEG

1. Introduction

Oscillations play important roles in cortical processing (Buzsáki 2006, da Silva 2013, Wang 2010). They are manifest visually as a regular series of oscillatory deflections in electroencephalography (EEG) or local field potential (LFP), or equivalently as a local concentration of power in their Fourier spectra. Oscillatory activity is usually superimposed on less structured activity, and it may be hard to see it in recorded data due to source-to-sensor mixing and competing sources. Indeed, oscillatory activity known to exist based on depth electrode recordings may be invisible from surface recordings in EEG (da Silva 2013). Scientists interested in measuring this activity face a never-ending battle against noise.

Spectral filtering (bandpass, etc) is an effective means to attenuate the background and improve signal-to-noise ratio (SNR). However, filters are prone to ringing triggered by noise or transients, and such ringing may masquerade as oscillations. Even for activity that is genuinely oscillatory, convolution of the signal with the impulse response of the filter changes its time course, both phase and temporal envelope. This may obscure temporal or causal relationships between a stimulus and its response, or between activity in different parts of the brain. Time–frequency analysis is prone to similar issues due to the temporal extent of its analysis kernels, which introduce smearing of response patterns and uncertainty as to their temporal relation to underlying brain events. An alternative means to improve the SNR of oscillatory activity is therefore welcome.

The correlation structure of multichannel data such as EEG, LFP, magnetoencephalography (MEG) or electrocorticography (ECoG) can be exploited to factor out noise sources. Many techniques have been proposed under the general name of *component analysis*. Here we use a



Content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

technique known as joint decorrelation (JD) (de Cheveigné and Parra 2014) to determine optimal linear combinations of observations. JD subsumes techniques such as denoising source separation, common spatial patterns and earlier techniques (Koles *et al* 1990, Fukunaga 1990, Särelä and Valpola 2005, de Cheveigné and Simon 2008, de Cheveigné 2010). Here we use it to scan the data for evidence of narrowband activity. The method is closely related to ideas proposed by Särelä 2004 and to the spatial spectral decomposition method of Nikulin *et al* (2011). Our contribution is to describe the method in an accessible way, within the JD framework, and to illustrate it on a number of common tasks.

The aim of this paper is to draw attention to a useful tool for the study of oscillatory activity, that can either replace filtering and time–frequency analysis, or else complement them. We describe the tool and illustrate its capabilities using synthetic and real EEG and MEG data.

2. Methods

2.1. Optimizing the SNR of narrowband activity

We use JD to maximize the SNR of narrowband activity. The J sensor or electrode signals, arranged as columns of a matrix $\mathbf{X} = [x_{ij}]$, where t is time, are combined linearly to produce K component signals y_{tk} :

$$y_{tk} = \sum_{j=1}^J x_{tj} w_{jk}, \quad (1)$$

where w_{jk} are optimal weights produced by the JD algorithm. In matrix notation, $\mathbf{Y} = \mathbf{X}\mathbf{W}$, where \mathbf{W} is the analysis matrix of dimensions $J \times K$, which converts J sensors into K components, $K \leq J$. The JD algorithm finds weights that optimize some criterion, implemented as a ‘bias filter’ applied to the data: $\tilde{\mathbf{X}} = \mathcal{F}(\mathbf{X})$. The nature of this filter determines what aspect of the signal is optimized by the linear combination of channels. In our case the filter \mathcal{F} is simply a narrowband filter centered on f , and thus JD finds components that maximize power near f . JD finds this solution by jointly diagonalizing the covariance matrices of the raw and filtered data.

The covariance matrix of the raw data is calculated as $\mathbf{C}_0 = \mathbf{X}^T \mathbf{X}$, and that of the filtered data as $\mathbf{C}_1 = \tilde{\mathbf{X}}^T \tilde{\mathbf{X}}$, and the desired transformation matrix is obtained by joint eigendecomposition of these matrices: $\mathbf{C}_0 \mathbf{W} = \mathbf{W} \mathbf{D}_0$, $\mathbf{C}_1 \mathbf{W} = \mathbf{W} \mathbf{D}_1$, where diagonal matrices \mathbf{D}_0 and \mathbf{D}_1 represent the eigenvalues, and the columns of \mathbf{W} the eigenvectors (Fukunaga 1990). The matrix \mathbf{W} defines the desired transform.

The first component y_{t1} (defined by the first column of \mathbf{W}) gives the largest possible ratio of filtered power to raw power: it is the linear combination with the best possible SNR for oscillatory activity with frequency f . The second component y_{t2} is the best linear combination within the signal subspace orthogonal to the first, and so-on. Depending on the purpose of the analysis, we may keep just one component (the optimal linear combination of channels), or several that

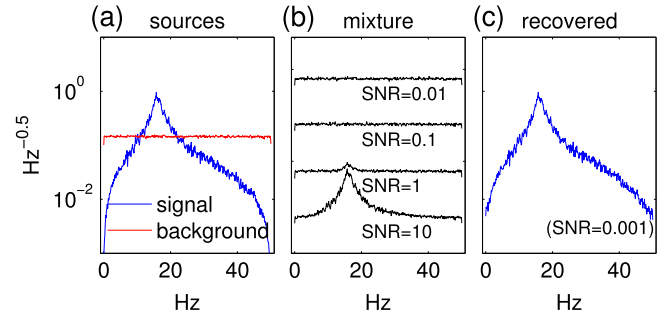


Figure 1. Illustration of the basic method with synthetic data. (a) Power spectra of oscillatory target (blue) and background activity (red). (b) Power spectrum of mixture at several values of SNR. (c) Oscillatory component recovered from mixture at SNR = 0.001.

together define an ‘oscillatory subspace’. Note that the y_{tk} are *not* filtered by the filter \mathcal{F} , which serves merely to find the weights \mathbf{W} . Once the weights are found, the components are obtained according to the formula $\mathbf{Y} = \mathbf{X}\mathbf{W}$ that does not involve \mathcal{F} , and thus the component signals are not distorted by spectral filtering.

The method thus far optimizes the power at f relative to the entire spectrum. It is possible to optimize it relative to the immediate spectral vicinity instead, by preprocessing the data with a bandpass filter wider than the bias filter and centered on f . This maximizes the local *spectral contrast* in the vicinity of the desired frequency (Nikulin *et al* 2011).

2.2. Scanning for oscillatory activity

Given a frequency f , we can find a linear combination of channels that maximizes the SNR of oscillatory activity at that frequency. If f is *unknown*, we can discover it by applying the procedure with a series of filter center frequencies f_i covering the range of interest. For each f_i the optimal component y_{t1} is examined for the presence of narrowband activity using standard tools such as the spectrum or spectrogram, or visual examination of the waveform for the presence of oscillations. The JD algorithm being computationally efficient, it is quite feasible to apply it repeatedly in this fashion.

3. Results

3.1. Simulated data

The method is first illustrated with synthetic data. A target source was created by filtering white Gaussian noise through a narrowband filter centered on 16 Hz (figure 1(a), blue), and mixed into 20 observation channels via a 1×20 mixing matrix with random coefficients. In addition to the target, 19 independent white Gaussian noise sources (figure 1(a), red) were added to the same channels via a 19×20 mixing matrix. The relative amplitudes of target and background were adjusted to attain a predefined SNR. Spectra of mixtures for various SNR are shown in figure 1(b). For SNR = 10, the

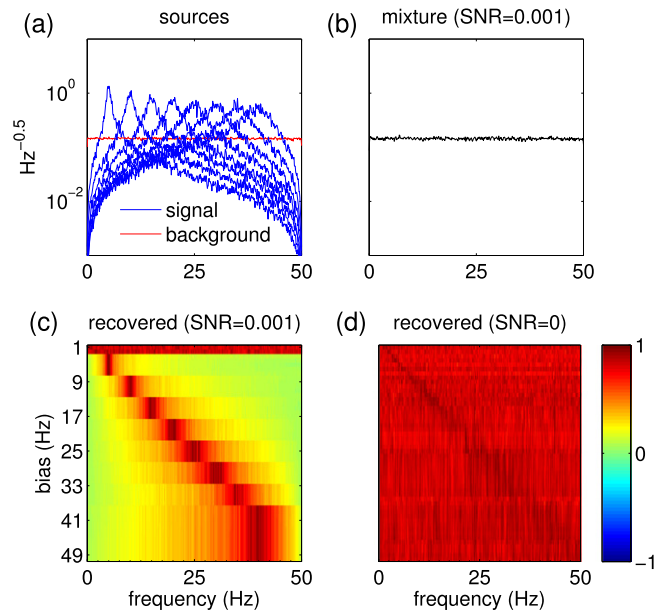


Figure 2. Exploratory scanning with synthetic data. (a) Power spectra of oscillatory targets (blue) and background (red). (b) Power spectrum of mixture at SNR = 0.001. (c) Raster plot of power spectra of recovered components for each value of the bias, for SNR = 0.001. Each line represents the color-coded spectrum of the first component recovered for a particular bias filter center frequency (ordinate). (d) Same for SNR = 0 (no sources). Each spectrum is normalized to [0, 1]. The color bar applies to all figures in this paper.

presence of the target is evident in the power spectrum of the mixed data, but for less favorable SNR it is quite difficult to discern. Applying JD with a bias filter tuned to 16 Hz reveals the target even for an unfavorable SNR, in this case SNR = 0.001 (figure 1(c)). This example demonstrates how a weak oscillatory source can be extracted from a much stronger noise background.

In a second example, eight narrowband target sources with different center frequencies (figure 2(a), blue) were mixed with 12 white Gaussian background sources at an overall SNR = 0.001. At this SNR the targets are completely concealed within the noise (figure 2(b)). The data were then scanned systematically at 1 Hz intervals using as bias filter a second-order resonator filter with quality factor $Q = 8$ (Smith 2007). The presence of the sources is evident in the raster plot of color-coded power spectra (figure 2(c)). Specifically: for each source there exists a transform, defined as a set of weights, that preserves that source and locks out the noise and the competing sources. The sources found are genuine: in the absence of a source (SNR = 0) the method finds nothing (figure 2(d)). The faint ridge along the diagonal reflects slight overfitting (see below). This example shows how multiple weak oscillatory sources can be discovered and isolated within the data.

3.2. Simulated 8 Hz target embedded in real EEG

In this example an 8 Hz oscillatory pulse (figure 3(a)) was added to 72-channel EEG at overall SNR = 0.1 (figure 3(b)). At this SNR the target is hard to discern within the EEG. The




Figure 3. Synthetic oscillatory pulse in EEG. (a) Target pulse (8 Hz). (b) Target mixed with EEG at SNR = 0.1. (c) Result of applying a causal narrowband filter centered at 8 Hz. (d) Same for a zero-phase filter. (e) Result of applying JD.

target SNR can be enhanced by applying a narrowband filter (resonator filter with $Q = 8$), but the response is temporally stretched and delayed (figure 3(c)). Applying instead a zero-phase filter (using the Matlab `filtfilt` function) eliminates the delay but causes the response to appear to start earlier than it should (figure 3(d)), potentially leading to erroneous inferences concerning the causal relation between events. In addition, spurious oscillatory patterns occur at other times triggered by the background noise. Time–frequency analysis is prone to similar issues (not shown). If instead we apply JD (using the same narrowband filter as bias) the SNR is improved without introducing any temporal distortion or spurious oscillations (figure 3(e)). In this case the noise rejection is not perfect, presumably because JD was calculated on a data segment too short (3s) to find an optimal filter. This example shows how JD can be useful as an alternative to the standard practice of applying narrowband filters or time frequency analysis.

3.3. Oscillatory activity in spontaneous EEG

Spontaneous EEG activity was recorded from a human subject in an exploratory study on brain state. The subject was placed within an electromagnetically shielded booth and instructed to remain with eyes closed. Data were recorded with a BioSemi EEG system with 32 electrodes with standard layout. Data were sampled at a 2048 kHz sampling rate, downsampled in software to 256 Hz, and high-pass filtered with a 2nd order Butterworth filter with cutoff 1 Hz. A sample of 850 seconds of data was scanned for narrowband activity at 1 Hz intervals. The bias filter at each frequency was a second-order resonator filter with quality factor $Q = 8$ (Smith 2007). Figure 4(a) shows the outcome as a raster plot. Each line represents the power spectrum of the first (optimal)

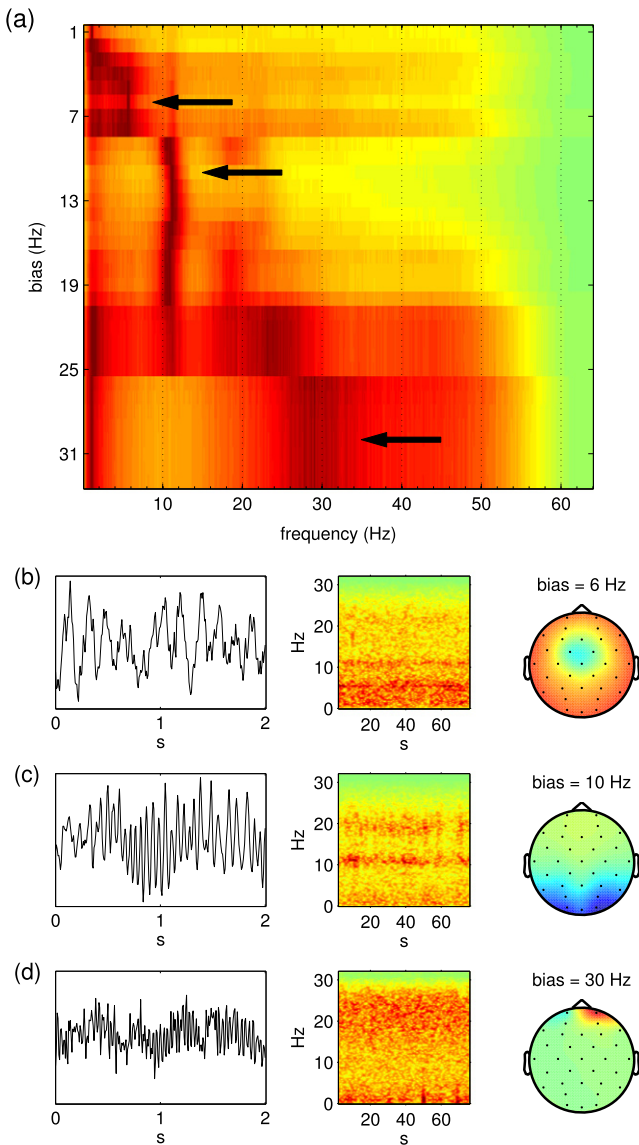


Figure 4. Exploratory scanning of spontaneous EEG data. (a) Raster plot of power spectra of components extracted by JD. Each line represents the color-coded spectrum of the optimal component for a particular bias filter center frequency (ordinate). For each bias (ordinate) the spectrum has power concentrated at certain signal frequencies (abscissa). (b)–(d) Waveform snippets (left), spectrograms (middle) and topographies (right) for three selected bias frequencies.

component for one bias frequency. The shape of the power spectrum (horizontal lines in the raster plot) changes according to the bias frequency. For certain bias frequencies (e.g. 6, 11 or 28 Hz) the spectrum has a clear concentration of power at a particular frequency, that is typically (but not always) close to the bias frequency.

Figures 4(b)–(d) shows waveforms, spectrograms, and topographies for three components discovered in this way. The first (6 Hz) has a focal frontal topography typical of theta band activity (Cavanagh and Frank 2014). The second (11 Hz) has a topography typical of occipital alpha activity originating from visual cortices, and the third (28 Hz) has a topography suggestive of ocular microsaccades (Yuval-

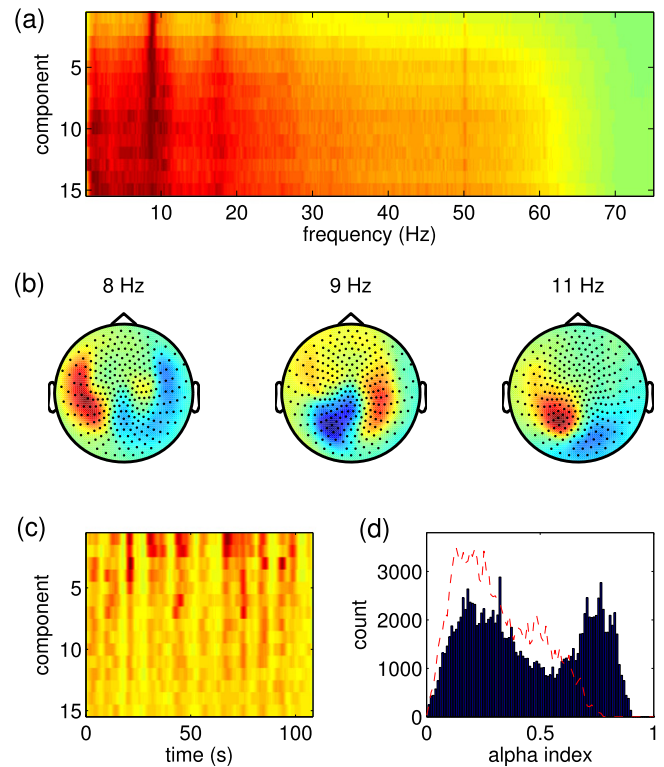


Figure 5. Alpha subspace in MEG data. (a) Raster plot of power spectra of the first 15 components of a JD analysis using a bias filter centered on 9 Hz. (b) Topographies associated with the first component for bias filters centered on 8, 9 and 11 Hz. (c) Raster plot of time courses of the rms temporal envelopes (1 s boxcar sliding window) for the first 15 JD components. (d) Histogram of instantaneous alpha-to-total power ratio for the first component (black bars) and the best-correlated raw sensor waveform (red dotted line).

Greenberg *et al* 2008). It is worth remarking that the signals shown in figures 4(b)–(d) are *not* affected by filtering (see methods). Thus there is no ambiguity as to whether the oscillatory pattern observed is genuine or the result of filter ringing (Yeung *et al* 2004). The time course of the oscillation can be followed accurately without the temporal smearing entailed by filtering. This example shows how multichannel data such as EEG may be scanned to reveal the presence of narrowband activity.

3.4. Exploring the alpha subspace in MEG

For each bias frequency, the JD algorithm gives a *series* of components ordered by decreasing score (ratio of filtered power to raw power), of which only the first is plotted in figure 4(a). For some bias frequencies (e.g. in the alpha range), several of these components may show a high proportion of narrowband power. Together these components define a *subspace* of the data containing narrowband activity at that frequency. This indicates that the data reflect *several brain sources* with similar spectral characteristics but distinct time courses and spatial signatures.

As an example, MEG data were recorded from a human subject during an exploratory study in which the subject was

performing a task that required responding to auditory stimuli. Here we ignore the sensory responses and focus instead on the background activity. Data were acquired from 274 radial gradiometers at a 600 Hz sampling rate and high-pass filtered at 1 Hz with a second-order Butterworth filter. Environmental noise was suppressed by (a) projecting out signals from reference sensors that pick up environmental noise using the TSPCA algorithm (de Cheveigné and Simon 2007), and (b) projecting out remaining 50 Hz-related power using the JD algorithm with a bias filter emphasizing 50 Hz and harmonics (de Cheveigné and Parra 2014). The cleaned data were submitted to PCA and 100 components were selected to reduce dimensionality. The data were then scanned for oscillatory components as described in the previous section. This analysis (not shown) revealed narrowband activity circa 9 Hz.

JD was then applied with a narrowband bias filter centered on 9 Hz ($Q = 4$), yielding a series of components ordered in terms of decreasing alpha-to-total power ratio. The power spectra of the first 15 components are shown as a raster plot in figure 5(a). Based on visual inspection, it would seem that roughly 13 components are dominated by alpha activity, although there is no clearcut transition between these and the remaining ‘non-alpha’ components. Collectively the selected components span an alpha-dominated subspace of the 274-dimension MEG signal space. There appears to be some diversity in their spectral properties (position and sharpness of the spectral peak, salience of a harmonic). Indeed, tuning the bias filter to neighboring frequencies emphasizes different aspects of this activity, and bias filters tuned to 8, 9 and 11 Hz yield different topographies (figure 5(b)). The different topographies imply different source geometries, and the spatial progression between these topographies is consistent with reports of lower frequency for frontal alpha sources (da Silva 2013).

Spectral filtering is not involved in the JD transform (the bias filter that served to determine the weights is *not* retained in the transform, see methods), and the method is thus useful to determine the time course of oscillatory activity without the confounding effects of ringing or temporal smearing (figures 3(c), (d), but see Caveats below). Figure 5(c) shows the root mean (rms) temporal envelopes of the first 15 components for the 9 Hz bias. To calculate the envelopes, the component waveforms were normalized (equal power for each component), filtered by a mild bandpass filter centered on 9 Hz ($Q = 1$), squared, and smoothed with a 1 s boxcar window. Their time course is deeply modulated, reflecting bursts of alpha power separated by periods of quiescence. The envelopes are correlated between components, suggesting that multiple sources tend to activate simultaneously. Between-component correlation is however not perfect, implying some temporal disparity between the activations of the various sources.

‘Alpha-dominated’ intervals in the data are characterized by a large ratio of power at 9 Hz relative to total power. Figure 5(d) right shows a histogram of this ratio for the first JD component (black). The distribution is bimodal, supporting a relatively clear distinction between ‘alpha’ and ‘non-alpha’ temporal intervals. For comparison, the same analysis




Figure 6. Stimulus-induced narrowband alpha activity in MEG data. (a) Raster plot of power spectra of the first 30 components of a JD analysis with bias filter centered on 15 Hz. (b) Time course of smoothed trial-averaged instantaneous power of the first 30 components. (d) Time course of instantaneous power of the first component for each trial. (c) Number of trials for which the power over interval 2.5–3.5 s is smaller than the power over interval 0.5–1.5 s (ERD), for each of the first 30 components. Chance and 5% confidence intervals are indicated as vertical full and dashed lines, respectively.

applied to a raw data channel (the one most correlated with this component) yields a unimodal distribution dominated by lower values of the ratio (red dashed line). This example shows how JD-based oscillatory analysis can help to characterize the multidimensional cortical activity that underlies the complex phenomenon known as ‘alpha’.

3.5. Stimulus-induced narrowband activity in MEG

This example uses MEG data borrowed from a published study that measured responses of human subjects to visual stimulation (Duncan *et al* 2009). During each 5 s trial, the subject fixated a cross during 2.5 s, followed by a grating within the lower right or left quadrant during 2.5 s. Stimuli were repeated for a total of 160 trials, of which a subset of 30 are used in the examples in this paper. Data were recorded with a 274-channel gradiometer MEG system (CTF) at a 600 Hz sampling rate. Further details can be found in the original study of Duncan *et al* (2009). The same data were also used for illustration in a later study on induced responses (de Cheveigné 2012).

Scanning for oscillatory activity, as explained above, revealed several narrow band components in the 14–17 Hz region. Here we focus on the time course of their power within each trial. The power spectra of the first 30 components of a JD analysis with a narrowband bias filter centered on 15 Hz are shown in figure 6(a) as a raster plot. The first ten or so are dominated by power near 15 Hz, whereas the




Figure 7. Stimulus-induced narrowband gamma activity in MEG. Left: spectrogram of selected component. Right: corresponding topography.

following ones have more power near 10 Hz. Interestingly, the time course of the trial-averaged power of these various components (figure 6(b)) shows a lull (event-related desynchronization, ERD) after the stimulus transition at 2.5 s (figure 6(b)). Figure 6(d) shows a raster plot of instantaneous power for the first component. The activity occurs as a series of bursts, with a reduced density after the stimulus transition. There is considerable trial-to-trial variability but a test based on the binomial distribution for no effect shows that the power reduction is highly significant, a conclusion that holds for several of these components (figure 6(c)). This result is remarkable given that the components were selected purely on the basis of their *spectral*, not temporal, properties.

The existence of such narrowband ERD effects is well known (Pfurtscheller and Lopes da Silva 1999). What narrowband JD analysis contributes is the ability to overcome the poor SNR of such activity relative to background, and to follow its time course accurately without the temporal smearing entailed by filtering or time–frequency analysis.

3.6. Stimulus-induced narrowband gamma in MEG

This last example involves extracting a narrow gamma band event-related synchronization response from the same MEG data as used in the previous example. The data were high-pass filtered at 30 Hz to de-emphasize the prominent low-frequency power, and scanned for narrowband activity between 30 and 60 Hz. This scan revealed a component near 53 Hz. Its amplitude increased after the stimulus transition at 2.5 s, with a slight downward sweep in frequency, and its topography was focal in the occipital region (figure 7). The same component had previously been found using other methods: beamforming (Duncan *et al* 2009), and quadratic component analysis (de Cheveigné 2012). The convergence of these widely different methods is reassuring.

These examples illustrate the range of situations where the method can be useful to observe oscillatory activity embedded within multichannel data at low SNR.

4. Discussion

Brain oscillations play a prominent role in the neurosciences of the brain, and tools to extract and enhance them are of

crucial importance. Classic tools such as spectral filtering or time frequency analysis are hobbled by issues related to ringing and temporal smearing. The tool that we review in this paper uses instead the between-channel correlation structure of the data to design a spatial filter that isolates the desired activity. In favorable conditions it can extract extremely weak oscillatory activity, without the temporal distortion entailed by spectral filtering.

4.1. Scanning for spectrally shaped activity

JD is effective to extract narrowband cortical activity at a frequency f . Applying it repeatedly with a range of values f_i allows the data to be scanned to discover unknown oscillatory activity. Advantages over alternative methods, discussed below, are (a) the analysis is sensitive to weak sources embedded in competing background activity, (b) the time course of oscillatory activity can be followed accurately without temporal smearing, (c) the method is easy to apply, requiring no selection of channels or knowledge of cortical or sensor geometry. It is a good exploratory tool to look for oscillatory patterns hidden within the data.

The method can in principle also be used to search for components with spectral characteristics other than narrowband, for example high-pass, lowpass, bandpass (e.g. beta or theta), or harmonic (multiple narrow bands at multiples of a common fundamental frequency).

4.2. Alternative methods

A straightforward approach to enhance the SNR of narrowband activity is to apply a bandpass filter (e.g. Basar *et al* 1999, Ben-Simon *et al* 2008, Mazaheri and Jensen 2009). A downside is that non-oscillatory events may excite the filter poles to produce ringing patterns that masquerade as oscillations, raising questions concerning the genuine oscillatory nature of the underlying cortical activity (Yeung *et al* 2004). Even if the oscillations are genuine, convolution with the impulse response of the filter alters their time course, smoothing the envelope and introducing latency shifts or non-causal effects that complicate the interpretation (figures 3(c), (d)). These are all the more marked as the filter is selective: there is a tradeoff between the benefit of filtering and the salience of ringing effects produced by the filter. In contrast JD improves SNR without distorting the waveform, and it can also operate at much lower SNRs (figure 1).

Time–frequency analysis is subject to the same issues as filtering. Each pixel of the representation is obtained by convolving the signal by an analysis kernel. The temporal span of this kernel is inversely related to spectral resolution (constant for short-term Fourier transform, frequency-dependent for wavelet methods), and the temporal alignment of the kernel relative to the analysis point is arbitrary. Typical choices are to align its *center of gravity* (to minimize latency at the expense of causality) or its *origin* (to preserve causality at the expense of latency). Unfortunately such details are often not fully reported, and this may lead to uncertainty

concerning the interpretation of temporal relations and causality (Zoefel and Heil 2012).

A different approach, in the case of multichannel data, is to enhance SNR by linear spatial filtering as in ICA, CCA, beamforming, PCA, and so-on. JD belongs to this family. The methods differ in how they find the weights to apply to the different channels (equation (1)). For ICA the weights are chosen to maximize a measure of statistical independence such as non-Gaussianity, hopefully yielding components that map to neural sources. These are then scanned for oscillatory activity. Unfortunately, narrowband signals are only weakly non-Gaussian and may be missed if SNR is low. For JD, in contrast, the weights are chosen to maximize narrowband activity directly, so oscillatory sources are less likely to be missed. Components are also ordered, obviating the need for a post-hoc selection phase. Nikulin *et al* (2011) describe a method similar to ours that enhances known narrowband components, by maximizing spectral contrast between a target frequency band and adjacent bands. They also briefly discuss a sliding analysis to discover unknown narrowband components, using as criterion the mean ratio between ten successive eigenvalues. Haufe *et al* (2014) apply a similar process for dimensionality reduction as an alternative to PCA.

4.3. Caveats

The method is sensitive, but it may nonetheless miss oscillatory activity that is spatially collinear with high-amplitude activity in another frequency band. The analysis selects components that are active at the target frequency *relative to other frequency bands*, and so weak narrowband activity might be missed if it is spatially collinear with strong activity in those other spectral regions. This can be alleviated by prefiltering to attenuate power in remote unwanted bands. A useful strategy is to preprocess the data with a bandpass filter centered on the same frequency as the bias filter, but wider, so as to enhance spectral contrast (see methods).

For a given bias frequency, JD may produce several narrowband components that together span a subspace of narrowband activity (e.g. figures 4(a) and 5(a)). There is usually not a one-to-one correspondence between each component and a neural source. Indeed, brain sources with similar frequencies are likely to be correlated, whereas component signals are necessarily uncorrelated. The amplitude of each component signal depends on the phase-dependent vector summation over sources. It is thus difficult to infer the time course of neural sources from the data.

The method is prone to overfitting, particularly if (a) the number of channels is large, (b) the amount of data is small, and/or (c) the available degrees of freedom have been reduced by prior filtering. Overfitting is manifest as an apparent peak in the spectrum at the same frequency as the bias filter, in the absence of any genuine oscillatory activity (e.g. the faint diagonal in figure 2(d)). Overfitting may be tested for by cross-validation techniques (de Cheveigné and Parra 2014), and the tendency to overfit may be reduced by initially applying PCA and discarding low-amplitude PCs to reduce the dimensionality of the data.

Artifactual ‘narrowband’ features may arise due to pre-filtering of the data (lowpass, high-pass, band reject, etc). For example the first row of the raster plot of figure 4(a) seems to indicate narrowband activity at 1 Hz. Instead it simply reflects the interaction between the overall low-pass spectrum of the data and the 1 Hz high-pass preprocessing filter. As explained in de Cheveigné and Parra (2014), activity that propagates across the sensor or electrode array may also incorrectly emerge as oscillatory activity. One must be attentive to such artifacts.

4.4. Oscillatory activity in the brain

The method is relatively sensitive, and can reveal the presence of multiple oscillatory sources within brain data. It is sobering to observe that these components represent only a fraction of the total activity. Most of the power in the data appears to be non-oscillatory, as far as we can tell using the method. It is of course possible that additional sources exist that are collinear with stronger non-oscillatory sources, or too weak to emerge even after processing. It is also conceivable that the apparently non-oscillatory activity is actually the superposition of multiple oscillatory sources, too numerous to be resolved within electrophysiological data. Nonetheless, we cannot exclude the alternative hypothesis that much of brain activity is inherently non-oscillatory.

Conclusions

Oscillatory patterns constitute an important aspect of brain activity, often hard to discern within data recorded from EEG or MEG because of the poor SNR relative to other activity. Spectral filtering, often employed to improve the SNR, is unsatisfactory because filter ringing triggered by non-oscillatory events may masquerade as oscillations, and the temporal extent of time–frequency analysis kernels may likewise blur temporal and causal relations. In this paper we introduced a methodology that relies instead on spatial filtering to reveal narrowband activity within multichannel data such as EEG, MEG, ECoG or LFP. The method exploits the between-channel correlation structure of the data to suppress competing sources and improve the SNR. Applied to real data, it can reveal weak sources that are hard to observe using standard methods, and allow their time course to be plotted accurately.

Acknowledgments

This work was supported by the EU H2020-ICT grant 644732 (COCOHA), and grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL* awarded to the Department of Cognitive Studies at ENS. The second author was supported by the Région Ile-de-France via the DIM Cerveau & Pensée and the Ecole Doctorale Frontières du Vivant (FdV) Programme Bettencourt. Thanks to Gareth Barnes for permission to use the data of Duncan *et al* (2009), and to two anonymous reviewers for comments on a previous version of this paper.

References

- Basar E, Basar-Eroglu C, Karakas S and Schurmann M 1999 Oscillatory brain theory: a new trend in neuroscience *IEEE Eng. Med. Biol. Mag.* **18** 56–66
- Ben-Simon E, Podlipsky I, Arieli A, Zhdanov A and Hendler T 2008 Never resting brain: simultaneous representation of two alpha related processes in humans *PLoS One* **3** e3984
- Buzsáki G 2006 *Rhythms of the Brain* (Oxford: Oxford University Press)
- Cavanagh J F and Frank M J 2014 Frontal theta as a mechanism for cognitive control *Trends Cogn. Sci.* **18** 414–21
- da Silva F L 2013 EEG and MEG: relevance to neuroscience *Neuron* **80** 1112–28
- de Cheveigné A 2010 Time-shift denoising source separation *J. Neurosci. Methods* **189** 113–20
- de Cheveigné A 2012 Quadratic component analysis *NeuroImage* **59** 3838–44
- de Cheveigné A and Parra L C 2014 Joint decorrelation, a versatile tool for multichannel data analysis *NeuroImage* **98** 487–505
- de Cheveigné A and Simon J Z 2007 Denoising based on time-shift PCA *J. Neurosci. Methods* **165** 297–305
- de Cheveigné A and Simon J Z 2008 Denoising based on spatial filtering *J. Neurosci. Methods* **171** 331–9
- Duncan K K, Hadjipapas A, Li S, Kourtzi Z, Bagshaw A and Barnes G 2009 Identifying spatially overlapping local cortical networks with MEG *Hum. Brain Mapp.* **31** 1003–16
- Fukunaga K 1990 *Introduction to Statistical Pattern Recognition* (New York: Academic)
- Haufe S, Dähne S and Nikulin V V 2014 Dimensionality reduction for the analysis of brain oscillations *NeuroImage* **101** 583–97
- Koles Z, Lazar M and Zhou S 1990 Spatial patterns underlying population differences in the background EEG *Brain Topography* **2** 275–84
- Mazaheri A and Jensen O 2009 Rhythmic pulsing: linking ongoing brain activity with evoked responses *Front. Hum. Neurosci.* **4** 177
- Nikulin V V, Nolte G and Curio G 2011 A novel method for reliable and fast extraction of neuronal EEG/MEG oscillations on the basis of spatio-spectral decomposition *NeuroImage* **55** 1528–35
- Pfurtscheller G and Lopes da Silva F H 1999 Event-related EEG/MEG synchronization and desynchronization: basic principles *Clin. Neurophysiol.* **110** 1842–57
- Särelä J 2004 Denoising source separation *PhD Thesis* Helsinki University of Technology
- Särelä J and Valpola H 2005 Denoising source separation *J. Mach. Learn. Res.* **6** 233–72
- Smith J O 2007 *Introduction to Digital Filters with Audio Applications* (W3K Publishing) (<http://w3k.org/books/filtersIp3/index.html>)
- Wang X-J 2010 Neurophysiological and computational principles of cortical rhythms in cognition *Physiol. Rev.* **90** 1195–268
- Yeung N, Bogacz R, Holroyd C B and Cohen J D 2004 Detection of synchronized oscillations in the electroencephalogram: an evaluation of methods *Psychophysiology* **41** 822–32
- Yuval-Greenberg S, Tomer O, Keren A S, Nelken I and Deouell L Y 2008 Transient induced gamma-band response in EEG as a manifestation of miniature saccades *Neuron* **58** 429–41
- Zoefel B and Heil P 2012 Detection of near-threshold sounds is independent of EEG phase in common frequency bands *Front. Psychol.* **4** 262