

Quadratic component analysis

Alain de Cheveigné^{a, b, c, *}

^a Laboratoire de Psychologie de la Perception, UMR 8581, CNRS and Université Paris Descartes, France

^b Département d'Etudes Cognitives, Ecole normale supérieure, France

^c Ear Institute, University College London, UK

ARTICLE INFO

Article history:

Received 31 May 2011

Revised 3 August 2011

Accepted 26 October 2011

Available online 4 November 2011

Keywords:

MEG

Magnetoencephalography

EEG

Electroencephalography

Noise reduction

Artifact removal

Principal component analysis

ABSTRACT

I present a method for analyzing multichannel recordings in response to repeated stimulus presentation. Quadratic Component Analysis (QCA) extracts responses that are *stimulus-induced* (triggered by the stimulus but not precisely locked in time), as opposed to stimulus-evoked (time-locked to the stimulus). Induced responses are often found in neural response data from magnetoencephalography (MEG), electroencephalography (EEG), or multichannel electrophysiological and optical recordings. The instantaneous power of a linear combination of channels can be expressed as a weighted sum of instantaneous cross-products between channel waveforms. Based on this fact, a technique known as Denoising Source Separation (DSS) is used to find the most reproducible “quadratic component” (linear combination of cross-products). The linear component with a square most similar to this quadratic component is taken to approximate the most reproducible evoked activity. Projecting out the component and repeating the analysis allows multiple induced components to be extracted by deflation. The method is illustrated with synthetic data, as well as real MEG data. At unfavorable signal-to-noise ratios, it can reveal stimulus-induced activity that is invisible to other approaches such as time-frequency analysis.

© 2011 Elsevier Inc. All rights reserved.

Introduction

Repeated stimulation can produce responses of two sorts: *evoked* activity that is time-locked to the stimulus, and *induced* activity that is triggered by the stimulus but not precisely time or phase locked. Evoked responses can be enhanced by averaging over trials, but induced responses cannot. Examples of induced response are an ongoing oscillatory process that is modulated by stimulation, or an event or burst of events that occurs with random latency. The aim of this paper is to facilitate the study of induced responses.

Data recorded by electroencephalography (EEG) or magnetoencephalography (MEG) often have very poor signal to noise ratio. Even with relatively low-noise recording techniques, such as multichannel electrode arrays used in animal electrophysiology or intracranial recordings in human subjects, a weak source of interest may be masked by competing brain activity. Techniques to reduce noise and separate sources are crucial to make sense of such data. There are three basic approaches to improving signal to noise ratio (SNR): averaging over trials, filtering and time-frequency analysis, and component analysis.

Induced responses are commonly approached using time-frequency analysis, which allows the experimenter to search for particular frequency regions in which the induced response appears most clearly. This can be understood as a technique to improve the SNR,

assuming that signal and noise occur in distinct spectral regions. After the phase is discarded, magnitudes may be averaged over trials to further enhance SNR. This approach works best if the activity of interest is locally narrow-band (for example oscillatory activity), but it is less effective for activity that is wide-band, for example consisting of impulsive events or spikes.

Component analysis leverages the degrees of freedom offered by multichannel data by forming linear combinations of observations that optimize SNR. It can take many forms, from simple channel selection or electrode re-referencing, to independent component analysis (ICA) (Delorme and Makeig, 2004; Makeig et al., 1996) and beamforming (Hillebrand et al., 2005; Sekihara et al., 2006). In each case, a weighted average of the sensor waveforms is formed to maximize sensitivity to the source of interest and minimize contributions from competing sources and noise. Random search for the best weights in the high-dimensional weight space would be prohibitively expensive, which is why methods that can find solutions that optimize plausibly useful criteria are precious. We use “denoising source separation” (DSS) (Valpola and Pajunen, 2000; Särelä, 2004; Särelä and Valpola, 2005), a general purpose component analysis method that can be used to optimize various criteria, for example reproducibility over trials (de Cheveigné and Simon, 2008b; de Cheveigné, 2010).

In brief, DSS involves the following steps: (a) orthogonalization of PCA of the data matrix, (b) normalization of the principal components so as to render the data “spherical”, (c) application of a contrast function to emphasize directions that satisfy the criterion, and (d) a final PCA to rotate the data so that the first component maximizes the

* Laboratoire de Psychologie de la Perception, UMR 8581, CNRS and Université Paris Descartes, France.

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

criterion, the second maximizes it within the subspace orthogonal to the first, and so on. The first two steps (PCA and normalization) are essential to neutralize the influence of variance, so that the solution is determined only by the contrast function. DSS is analogous to ICA, but differs in that it searches directly for linear components that optimize the criterion of interest, rather than indirectly via a measure of statistical independence (that may or may not isolate a similar component). It differs from PCA in that the criterion optimized is not power. Like beamforming it suppresses competing sources while preserving the source of interest, but differs in the criterion used to define that source. Whereas beamforming preserves components coming from a particular location, DSS preserves those with a particular property. DSS (and similar methods such as FOP and CSP) might thus be understood as a “geometry-free” version of beamforming. DSS is complementary to these well established tools. DSS closely resembles the Filtering by Optimal Projection (FOP) and Common Spatial Pattern (CSP) algorithms (Blankertz et al., 2008; Boudet et al., 2007; Koles et al., 1990) that have been widely used for classification in the Brain Computer Interface (BCI) literature, see also Cichocki (2004), Fukunaga (1990), Parra et al. (2005).

Recently we applied DSS to find linear combinations of channels that maximize reproducibility (de Cheveigné and Simon, 2008b; de Cheveigné, 2010). This approach is extremely effective to reveal evoked brain activity, even very weak, as it capitalizes on both trial-averaging and component analysis. Unfortunately it does not work for induced activity, both because averaging over trials provides no benefit, and because we lack a criterion to direct the component analysis.

The present paper extends the usefulness of DSS by applying it to cross-products between sensor waveforms. This allows the search to explore a larger space (of dimension $J(J+1)/2$, where J is the number of sensors), and to tap higher-order correlation structures, such as repeatable stimulus-induced patterns of power or correlation between sources. We use search within the space of cross-products as an auxiliary to search within the J -dimensional space of linear components, but the quadratic solutions may also be of interest in their own right, as indicators of significant brain activity that is reliably related to stimulation.

Methods

Signal model

Waveforms measured at the sensors $\mathbf{X}(t) = [x_1(t), \dots, x_J(t)]^T$ are assumed to be a linear superposition of interesting brain activity and uninteresting noise activity:

$$\mathbf{X}(t) = \mathbf{X}_B(t) + \mathbf{X}_N(t). \quad (1)$$

The first term results from the superposition within sensors of brain sources $\mathbf{X}_B(t) = \mathbf{A}_B \mathbf{S}_B(t)$, and the second from the superposition of noise sources $\mathbf{X}_N(t) = \mathbf{A}_N \mathbf{S}_N(t)$, where \mathbf{A}_B and \mathbf{A}_N are mixing matrices that describe how each source affects each sensor. Our aim is to enhance components $\mathbf{X}_B(t)$ and suppress components $\mathbf{X}_N(t)$ by combining observations linearly. Specifically, given a hypothetical source $s_b(t)$ we wish to find a linear combination of channels $\sum_j w_j x_j(t)$ that maximizes that source's SNR.

Here we assume that $s_b(t)$ consists of *induced* activity, triggered by repeated presentations of a stimulus but not accurately time-locked so that it cannot be enhanced merely by averaging over trials. Our operational definition of “induced” is that the short-term power is reproducible, and our objective is to find a linear combination of sensor waveforms with a maximally reproducible pattern of power (variance) over time relative to the stimulus.

Given a linear combination of sensor waveforms:

$$y(t) = \sum_j w_j x_j(t), \quad (2)$$

its instantaneous power can be expressed as:

$$y^2(t) = \left[\sum_j w_j x_j(t) \right]^2 = \sum_{jk} w_j w_k z_{jk}(t). \quad (3)$$

In other words, it is a linear combination of instantaneous cross-products $z_{jk}(t) = x_j(t)x_k(t)$. We will refer to a linear combination of cross-products (with arbitrary weights) as a *quadratic component*. Because of symmetry there are $J(J+1)/2$ distinct cross-products, and so the space of quadratic components is of dimension $J(J+1)/2$, whereas the space of linear components is of dimension J : not every quadratic component equals the square of a linear component as in Eq. (3).

The $J(J+1)/2$ weights of a quadratic component can be arranged in a symmetric matrix $[U_{jk}]$ by sharing weights equally between $z_{jk}(t)$ and $z_{kj}(t)$. The rank of such a matrix is usually J , whereas the rank of the matrix $[W_{jk}] = [w_j w_k]$ is one, which is yet another way of saying that not every quadratic component maps to the square of a linear component. Using singular value decomposition, $[U_{jk}]$ can be expressed as a sum of matrices of rank one. Discarding all terms except the largest (corresponding to the largest eigenvalue), $[U_{jk}]$ can be approximated by this rank-one matrix, that corresponds to a linear component which, squared, can be used as an approximation of the quadratic component. Thus, given an arbitrary quadratic component, we have a means to approximate it as the square of a linear component.

The method in this paper finds the maximally reproducible quadratic component of a set of data using the DSS algorithm, and then approximates it by the nearest squared linear component in the sense just defined.

Algorithm

Given a data matrix of size $[J, T, N]$ (channels \times time \times trials), the following steps are performed:

1. The mean of each sensor waveform is subtracted and the waveforms are normalized to give equal weight to all sensors. PCA is performed and only the first K components are retained to reduce computational cost.
2. Cross-products are formed and low-pass filtered (smoothed) by convolution with a square window of size P , and the initial P samples are removed to avoid the onset artifact. There are $K(K+1)$ distinct cross-products, which we augment with a constant component that serves to absorb DC power in subsequent steps. This results in a matrix of size $[K(K+1)+1, T-P, N]$.
3. DSS is applied to the matrix of cross-products. Specifically: a first PCA is performed and the PC waveforms are normalized. They are then averaged over repetitions and a second PCA is performed on the average to find the directions of maximal repeated power. The product of first PCA matrix, normalization matrix, and second PCA matrix defines a matrix that, applied to the matrix of cross-product waveforms, yields $K(K+1)+1$ quadratic component waveforms.
4. The first component (constant) is discarded, and the second component is selected as the most reproducible quadratic component.

This optimal quadratic component is of interest as a statistic that shows that the neural response is reliably related to the stimulus. However it does not fulfill our goal of finding a linear combination

of sensor waveforms with a square that is maximally reproducible. That goal can be approached by the following additional step.

- The $K(K+1)/2$ weights of the best quadratic component are rearranged within a symmetric $K \times K$ square matrix. The coefficients of the eigenvector with largest eigenvalue (in norm) define a linear component, the square of which is taken to approximate the quadratic component found in step 4.

A final step can be used to extract multiple components by deflation:

- The linear component found in step 5 is projected out of the data, and the analysis is repeated from step 1.

In the event that multiple components are found, we cannot guarantee that each maps to a neural source. All we can say is that the components collectively define a subspace that contains the induced response. Components are orthogonal by construction, whereas neural sources participating in the same computation are likely to be correlated, so a one-to-one correspondence is unlikely (except perhaps for the first component).

The purpose of smoothing in step 2 is to average out the fine temporal structure and better estimate the short-term power. The value of P is not critical, and smoothing can often be omitted. Insufficient smoothing seems to be detrimental if the target is pulse-like with little serial correlation (as in the first simulation below), whereas excessive smoothing may encourage multiple evoked sources to merge into a common quadratic component.

Implementation

The algorithm was implemented in Matlab. The PCA in step 1 serves to concentrate most of the variance in a small number of components, reducing computational cost (possibly at the expense of components of interest if they are weak). A value of $K=50$ seems a reasonable compromise in practice. To save space, rather than explicitly representing all cross-products for the entire data, covariance matrices may be calculated on smaller chunks and then added.

Results

The algorithm is evaluated first with synthetic test data, and then with MEG data.

Simulated data

Two channels, target is modulated wideband gaussian noise

The target consisted of white gaussian noise modulated by half a period of a sinusoid, repeated 100 times, each time with a different noise carrier (Fig. 1 left top), whereas interference consisted of unmodulated gaussian noise (Fig. 1 left bottom). Interference was identical in both channels, whereas the signal was added to both channels with opposite phase, with an RMS amplitude 0.0001 times that of the interference ($\text{SNR} = 10^{-8}$). It is thus very hard to guess the presence of the target in the mixtures (Fig. 1 center). The algorithm nevertheless finds a clearly repeatable quadratic component (Fig. 1 right top), and a corresponding linear component that resembles the target (Fig. 1 right bottom). This example demonstrates that the algorithm can extract a weak induced target in the absence of any spectral disparity between target and interferer.

Ten channels, target is modulated sinusoid with random phase

The target consisted of a random-phase sinusoidal carrier modulated by half the period of a sinusoid (Fig. 2 top left). The target was repeated 100 times, each time with a different carrier phase, and added equally to all channels with an RMS amplitude 0.0001 times that of the

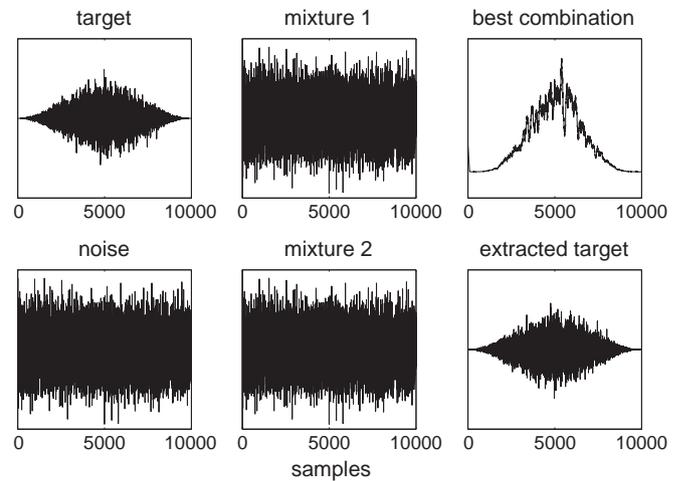


Fig. 1. Simulated data. The target consists of 100 repetitions of gaussian noise amplitude-modulated by a sinusoid (left top: one trial). This target is mixed with gaussian noise interference (left bottom) to form two “sensor signals” (center) with $\text{SNR} = 10^{-8}$. The algorithm extracts the most reproducible linear combination of cross-products (right top), and uses that to derive a linear combination of “sensors” similar to the target (right bottom).

interference ($\text{SNR} = 10^{-8}$). Interference consisted of 9 independent sources of white gaussian noise distributed over channels via a 9×10 random mixing matrix. Again, the presence of the target is impossible to guess visually within the interference (Fig. 2 top center), but the algorithm nevertheless recovers it (Fig. 2 top right). This example shows that the algorithm can extract a very weak oscillatory target from within multichannel data.

274 channels, target is pulse with random latency

The target consisted of a 60 ms pulse with a latency randomly chosen from within a 0.5 s interval (Fig. 2 bottom left), mixed across channels via a 1×274 random matrix. The noise consisted of background activity recorded from an MEG system with a subject present. The RMS amplitude of the signal was set to 0.0001 times that of the noise ($\text{SNR} = 10^{-8}$). The algorithm successfully extracted the target signal (Fig. 2 bottom right). This example shows that the algorithm can deal

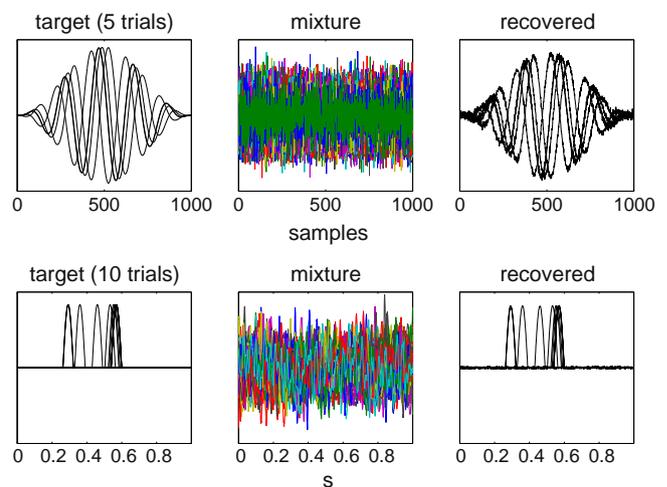


Fig. 2. Simulated data. In the top row, the target consists of a random-phase sinusoid modulated by one period of a sinusoid-shaped envelope, added equally to all channels. The interference consists of 9 white gaussian noise sources distributed over the sensor channels via a 9×10 random matrix, $\text{SNR} = 10^{-8}$. In the bottom row, the target consists of a 60 ms pulse with a random latency added equally to all 274 channels. The interference consists of real MEG background noise, recorded from an MEG system with a subject present, $\text{SNR} = 10^{-8}$. In both cases the algorithm successfully recovers the target (3rd column).

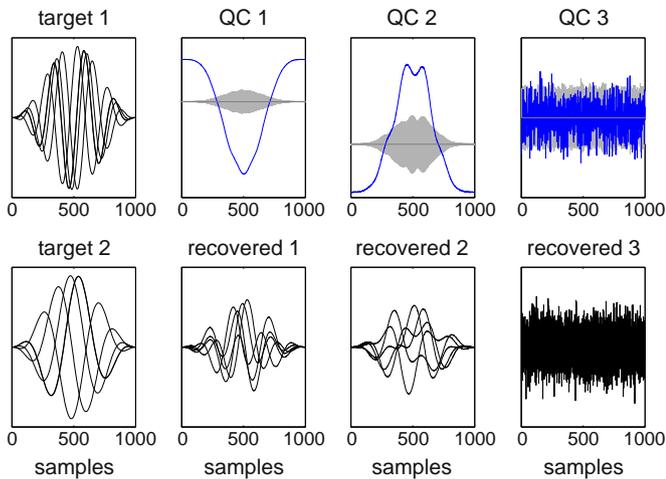


Fig. 3. Simulated data. Left: two targets, each consisting of a random-phase sinusoid modulated by one period of a sinusoid-shaped envelope (left column). Each line is a different trial. The targets were embedded in interference consisting of 8 white gaussian noise sources distributed over the sensor channels via a 8×10 random matrix, at $\text{SNR} = 10^{-8}$. A first application of the algorithm produced a quadratic component (column 2 top). The blue line represents the mean over repetitions, and the gray band ± 2 standard deviations of a bootstrap resampling of the mean. The corresponding linear component is shown in column 2, bottom. After this linear component was projected out of the data, a second application of QCA produced a second quadratic component and corresponding linear component (column 3). Subsequent applications of QCA revealed no further interesting components (column 4), suggesting that the first two iterations exhausted all induced activity. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with noise typical of an MEG system (with spectral, temporal and spatial correlations).

Ten channels, two targets

Two targets consisted each of a random phase sinusoidal carrier modulated by half the period of a sinusoid (Fig. 3, left column). The targets differed by the frequency of their carrier. Each was added equally to all channels via a 1×10 random matrix, with an RMS amplitude of 0.0001 times the noise ($\text{SNR} = 10^{-8}$). The interference consisted of 8 independent sources of white gaussian noise distributed over channels via a 8×10 random mixing matrix. The algorithm was applied three times, and each time the solution was projected out before the next iteration. The first two iterations (Fig. 3, columns 2 and 3) yielded two linear components that share some characteristics of both targets, but are not identical to either. This illustrates an important property of the method: the extracted components do not necessarily map to individual sources; rather they span a subspace that contains the sources. Subsequent iterations revealed no further components (Fig. 3, column 4), implying that the first two iterations exhausted the underlying targets. In this simulation both targets shared the same envelope, but simulations with different envelopes gave similar results (not shown). This example shows that the algorithm is not confused by the presence of multiple reproducible sources.

Real MEG data

MEG data were taken from a study that reported stimulus-specific gamma activity in a visual task (Duncan et al., 2009). Data were recorded from a whole-head MEG system equipped with 275 gradiometers, at a 600 Hz sampling rate after low-pass filtering in hardware at 150 Hz. The reader is referred to the original report for details of the experiment and interpretation of the data.

For the purpose of the present study, data were preprocessed by (a) projecting out components dominated by power at 50 Hz and harmonics (projection to the subspace orthogonal to those components, de Cheveigné and Simon, 2007), (b) projecting out components

dominated by the evoked response to the stimulus (de Cheveigné and Simon, 2008b), and (c) suppressing sensor-specific components due to sensor noise (de Cheveigné and Simon, 2008a). Data were then submitted to a PCA, and the first 50 principal components retained. The data set then consisted of a 3D matrix of size 1500 samples, 50 principal components, and 86 trials (responses to stimuli in the left visual field). QCA was applied as described in Methods, in three stages that differed in details of processing parameters.

In a first stage, the algorithm was applied to wideband data (no filter was applied other than the initial hardware filter). No smoothing was applied to the quadratic components. The first row of Fig. 4 shows the best quadratic component (left column, upper plot) and the linear component with closest square (left column, lower plot) averaged over trials. The gray band represents ± 2 standard deviations of a bootstrap resampling of the mean (Efron and Tibshirani, 1993), performed by drawing a random sample of trials (with replacement) repeated 200 times. The time course of the linear component for individual trials is shown as a raster plot in column 2. The activity consists in relatively slow deflections, with a density that is greatest during the second half of the epoch. The topography (column 3) is restricted to a single frontal channel, suggesting muscular activity in the eyes (despite the instruction given to subjects to fixate a fixed point on the screen). This linear component was then regressed out, and the processing repeated to produce a second component (row 2 of Fig. 4) that had a more widespread frontal distribution (column 3). Its spectrum (column 4) is low-pass but less steep than the first. In addition to these, the analysis revealed additional components with spectra and spatial distributions similar to the second (not shown). It should be stressed that individual components do not necessarily map to individual neural sources. In particular the spatial distribution may not reflect that of a neural source. Instead, components should be understood as spanning a subspace of the data that contains the neural activity of interest (see below).

In a second stage, data were high-pass filtered with 4 Hz cutoff (4th order Butterworth) to attenuate the slow components that dominated the results of the first stage, and thus improve the SNR of mid-frequency activity before attempting QCA. The best component, illustrated on row 3 of Fig. 4, has a relatively narrow spectrum peaking at 10 Hz (column 4), an occipital or parietal distribution (column 3), and an amplitude that dips after trial onset and after the stimulus transition at 2.5 s (columns 2 and 1). A second component (obtained after projecting out the first) has a similar time course but a spectrum that peaks at both 10 and 15 Hz. In addition to these, a total of about 10 similar components were found, mostly with occipital/parietal distributions and spectra peaking at 10 or 15 Hz (or both). Together they span a subspace of relatively high dimensionality, implying quite a complex set of sources. The time-course, similar for all these components, fits the description of an often-observed “event-related desynchronization” (ERD), although the rather rapid transition to the dip is perhaps suggestive of a drop in amplitude rather than desynchronization of ongoing oscillations.

In a third stage, data were high-pass filtered with a 30 Hz cutoff (4th order Butterworth) to attenuate the strong low and mid-frequency components and improve the SNR of gamma-band components (which were the focus of the original study). The best quadratic component, illustrated on row 5 of Fig. 4, has a relatively narrow spectrum peaking at about 50 Hz (column 4), a relatively focal occipital distribution (column 3), and an amplitude step at the stimulus transition at 2.5 s (columns 2 and 1).

A spectrogram (Fig. 5 top) shows a time course quite similar to that reported in the original study (Duncan et al., 2009, their Fig. 1C), with an initial downward chirp from about 56 Hz to 49 Hz. That study applied 20–70 Hz bandpass filtering and a spatial filter based on a beamformer algorithm. The similarity corroborates our analysis and suggests that QCA and beamforming may have found approximately the same component (i.e. spatial filter). The time course is much clearer than if time-frequency analysis were applied directly to selected sensor

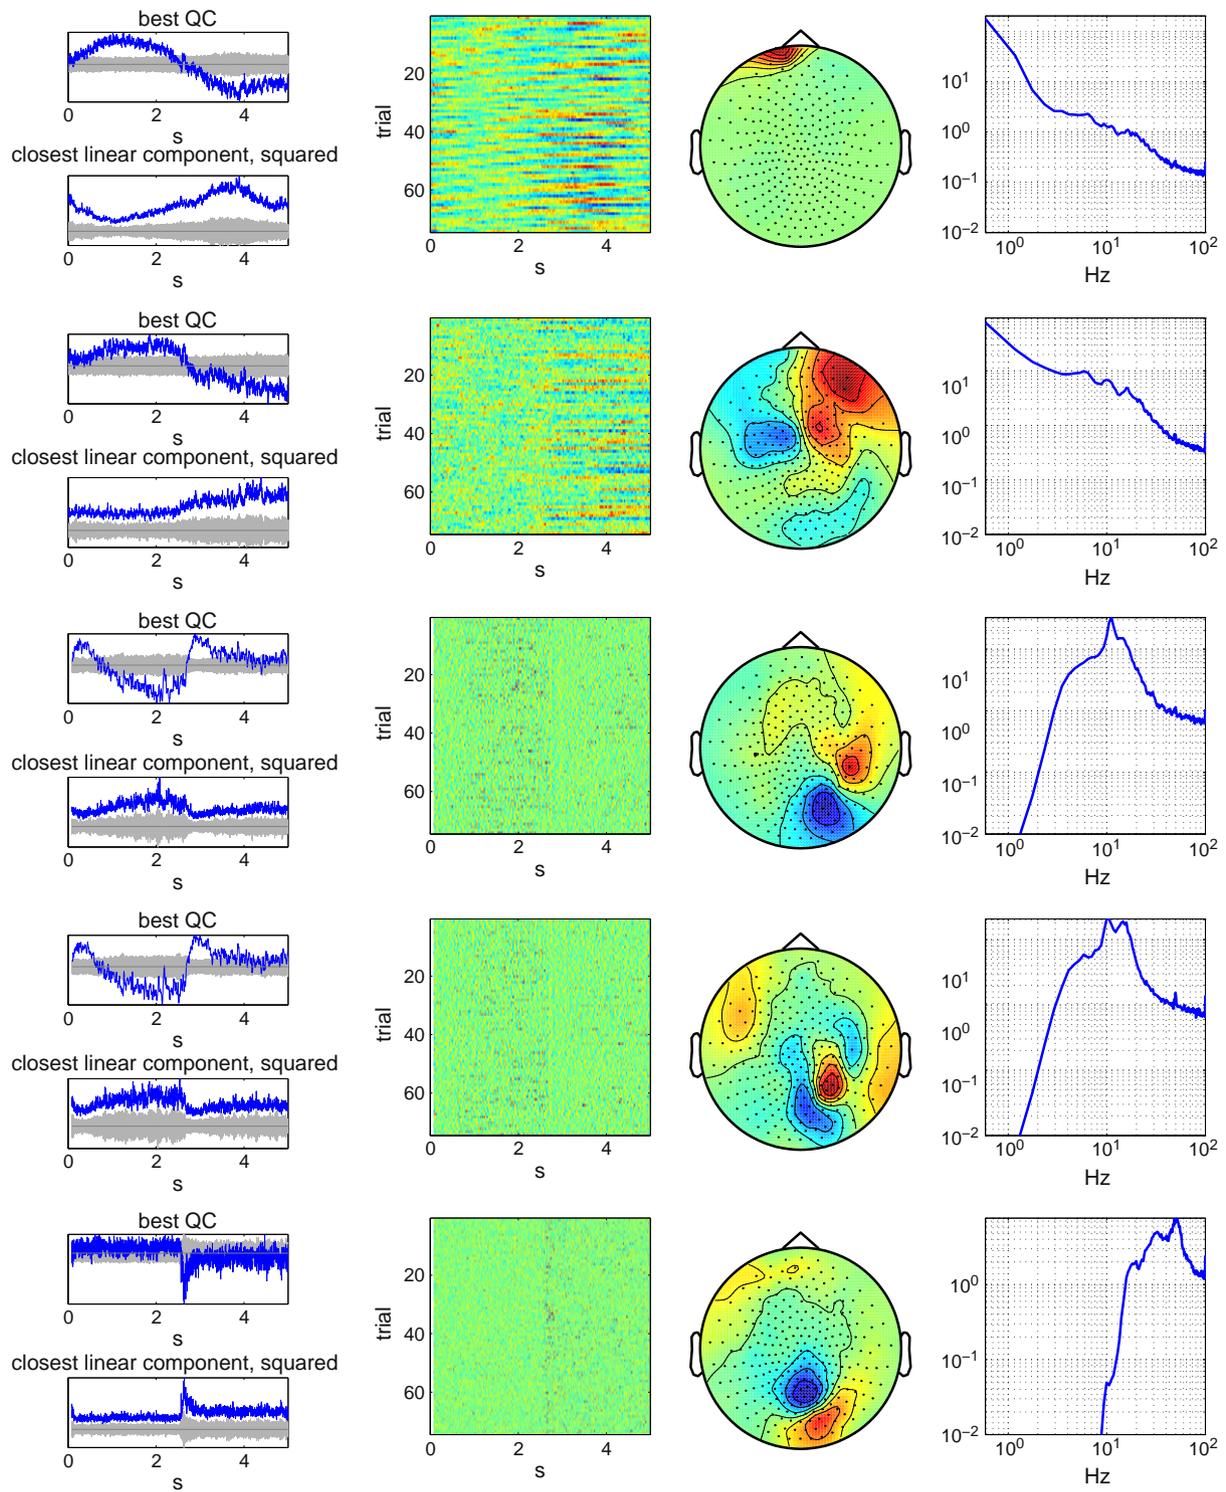


Fig. 4. Real data. Stimulus-induced components revealed by QCA analysis applied to MEG data recorded in response to visual stimulation (Duncan et al., 2009). Each trial consisted of two intervals of 2.5 s, the first containing a uniform field and the second an obliquely-oriented grating (see reference for further details). Each row describes one induced component of the response revealed by QCA. The first column shows the average over trials of the quadratic component (top), together with the instantaneous power of the linear component with closest square (bottom). Gray bands represent ± 2 standard deviations of a bootstrap resampling of the mean. The second column shows a raster plot of individual repetitions of the closest squared component, the third column shows its spatial distribution, and the last column its power spectrum (normalized by dividing by the mean over spectral bins). In general: cortical activity can be understood as belonging to a space spanned by these components, but there is not necessarily a one-to-one map between components and neural sources. See text for a discussion of each row.

channels (Fig. 5 bottom). Projecting out this component and repeating the analysis revealed no further gamma components. The reader is referred to Duncan et al. (2009) for a discussion of the significance of these responses for brain function, the purpose here being merely to illustrate the ability of our algorithm to extract them from MEG data.

Discussion

The simulations showed that the method can retrieve non-repeating activity (that cannot be enhanced by averaging), for wideband targets (time-frequency analysis would also have failed), and at very unfavorable

SNRs (as low as 10^{-8}). With real MEG data from a study using visual stimulation, the method retrieved a gamma-band response very similar to that obtained in that study by other means, together with multiple other induced response components. It thus seems to attain the goal defined in the Introduction, of extracting weak induced activity from multichannel data. The method produces linear components (weighted sums of sensor signals), in analogous fashion to other component analysis techniques such as beamforming or ICA.

Caveats and cautions

The method requires forming $J(J+1)/2$ cross-products and calculating covariance matrices of size $(J(J+1)/2)^2$, which may be prohibitive in processing and memory costs. The initial PCA alleviates this problem, but introduces the risk that a component that is of interest but weak in power might be discarded. If the data epoch is short and/or the sampling rate is small, the matrix of cross-products may not be of full rank. Smoothing (step 2) might exacerbate the effect by introducing serial correlation that reduces the “effective” number of samples. However this is not a serious issue as DSS is robust to rank-deficient data.

There are many free parameters and thus a risk of overfitting, that may result in seemingly repeatable components by chance. The risk is greater for larger values of the smoothing parameter P , as the serial correlation induced by smoothing reduces the degrees of freedom in the data. Any conclusions based on this method should be challenged by cross-validation or bootstrap, or corroborated in other ways. This risk is illustrated in Fig. 6. The top left hand panel represents the best component produced by QCA for the 30 Hz high-pass data of Duncan et al. (2009) for $P=100$ (instead of $P=1$ as in Fig. 4 bottom). The large amplitude of the mean (blue) relative to the bootstrap standard deviation (gray) suggests a clearly repeatable response. However the same analysis applied to a random matrix of same size also seems to produce a repeatable component (Fig. 6 top right), in this case the result of overfitting encouraged by the large value of the smoothing parameter. Bootstrap resampling (Fig. 6 bottom) resolves the issue: for the genuine response the component is reproducible across resamplings, for the noisy data it is not.

The method is sensitive to any source of repeatable power, not all of which are of interest. For example artifacts such as eyeblinks or muscular activity may tend to occur more at certain times within the epoch (e.g. Fig. 4, first row). More insidiously, repeatable power

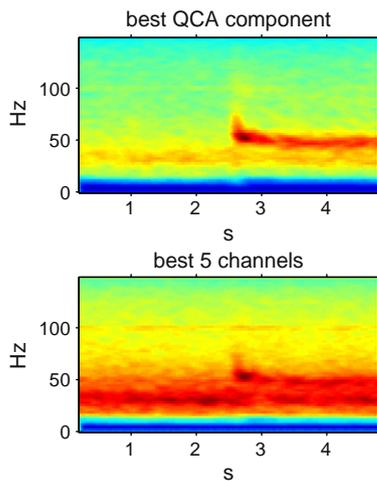


Fig. 5. Time-frequency analysis of best QCA component (top) or best 5 sensors (bottom) in the gamma band. Based on short-term Fourier analysis with a 256-point (0.43 s) Hanning analysis window and a 16-point (0.027 s) frame period. Color represents power raised to the exponent 0.2.

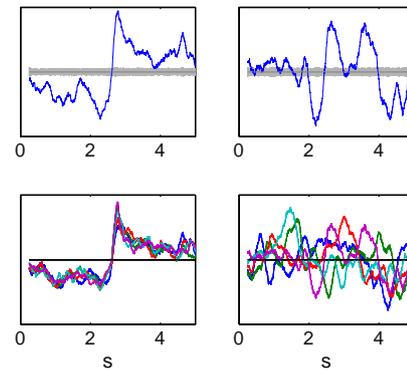


Fig. 6. Checking for overfitting. Top left: best quadratic component for data of Duncan et al. (2009) highpass filtered at 30 Hz, with the QCA smoothing parameter set to $P=100$. The blue line is the mean over trials, the gray band indicates ± 2 standard deviations of a bootstrap resampling of the mean. Top right: same for random data of same size. The seemingly reproducible pattern is entirely artifactual. Bottom: outcomes of QCA analysis for five resamplings of the data (trials drawn randomly with replacement). A genuine response is stable over resampling, an artifactual response is not. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

may arise from seemingly harmless preprocessing operations. For example if a filter is applied to each epoch individually, the onset artifact will appear as repeatable power. This may be avoided by applying the filter to the raw data before cutting into trials, or else by using FIR filters of order P and removing the first $P-1$ samples of each trial. As another example, in the presence of a slow drift in the sensor signals, removal of the mean from each trial will cause the power of the drift component to be smaller at the epoch center relative to its extremities, again resulting in a repeatable pattern of power. This can be avoided by applying high-pass filtering to the raw data before cutting them into trials.

DSS is guaranteed to find the most repeatable linear component (if applied to sensor waveforms) or quadratic component (if applied to cross-products). However this does not guarantee optimal repeatability of the power of the linear component found in step 5, although in practice the solutions seem to be quite good. In this sense the algorithm can be seen as a tool to direct the search within weight space to find induced activity.

As pointed out in the Methods section, multiple components found by repeating the analysis (as in Fig. 4) do not necessarily each correspond to an individual source within the brain. All that can be said is that observable neural sources reside within the subspace spanned by the components. Components produced by DSS or QCA are mutually orthogonal, which is unlikely to be the case of neural sources that participate in sensory processing. Further processing is required to find components within this subspace that map to individual neural sources. This is outside the scope of this paper.

If a weak component of interest is accompanied by stronger induced components, those may need to be “stripped away” before the component of interest can be extracted. There is a risk that the weaker component is lost in this process, particularly if it is partly colinear with others. For example the gamma response of Fig. 4 bottom was much weaker than the ocular, evoked, and induced alpha components. In such situations it may be useful to attenuate those competing components by filtering, as we did, prior to applying QCA.

Relation with time-frequency analysis

Time-frequency analysis is a common tool for the study of induced responses. It can be understood as performing two functions: demodulation (as in the squaring operation of Eq. (2)) and SNR enhancement (on the assumption that signal and noise have different

spectral properties). It also provides a descriptive tool to characterize properties of the source. Time-frequency analysis has three important drawbacks: (a) it is useful only if the target or interference have distinct spectro-temporal properties, (b) convolution with the filter or analysis kernel smears the response over time, possibly leading to incorrect conclusions about the latency and time course of neural processes, (c) it introduces a bias towards interpreting neural responses as oscillatory. QCA can reveal sources with little or no spectral structure embedded in wideband noise (Figs. 1, 2 bottom). It can also reveal genuine oscillatory activity (e.g. Fig. 4, rows 3–5) with good temporal resolution. Importantly, QCA can also be combined with filtering (as in Fig. 4, rows 3–5) or time-frequency analysis (as in Fig. 5), effectively enhancing the effectiveness of these tools.

Beyond the induced response

The approach might usefully be extended in at least three directions. First, induced activity, defined as a linear component with repeatable power, lives within an J -dimensional manifold of the $(J+1)/2$ space of quadratic components explored by QCA. Mathematically it is clear that there are other objects within that space, and these might be of interest if they relate reliably to processes within the brain. A hypothetical example might be if two neural sources, neither of which is repeatable (evoked) or with repeatable power (induced), synchronize at specific times relative to stimulation. In other words, QCA has the potential to reveal stimulus-related brain activity that is neither evoked nor induced. This possibility remains to be investigated.

Second, for simplicity and definiteness this paper focused on activity related to repeated stimulation, but the DSS operation is applicable to a wider range of situations, for example to contrast between two experimental conditions, or between different time intervals of a response, etc. Furthermore, the idea of applying linear analysis to cross-products might also be usefully used with other linear analysis techniques such as ICA (Lindgren and Hyvärinen, 2007).

Finally, QCA might usefully be applied to time-shifted data (similar to what was proposed for DSS in de Cheveigné, 2010). Time shifts allow the solution to implement finite impulse response filters (FIR) to be applied to the sensor waveforms, and to compensate for between-channel delays. They may also allow the analysis to discover higher-order spatiotemporal (or spatio-spectral) patterns within the data. These avenues remain to be explored.

Conclusion

Quadratic component analysis (QCA) is a new method for the discovery of patterns of neural activity within multichannel data. It finds the optimal linear combination of cross-products of sensor channels to maximize repeatability, and derives from it a linear combination that optimizes the signal-to-noise ratio of the neural response. It is applicable to *induced* responses (not stimulus-locked) that are hard to reveal by other methods. It can also, in principle, reveal neural activity that is neither evoked nor induced, as the space spanned by quadratic components is larger than that spanned by linear components. In simulations the method accurately retrieved sources embedded in noise (white noise or MEG noise) at extremely unfavorable signal-to-noise ratios. Applied to real MEG data the method revealed numerous stimulus-induced responses.

QCA is complementary with other established methods such as time-frequency analysis or beamforming, that are routinely used to characterize induced activity. As such it is potentially an important new tool to measure brain activity.

Acknowledgments

I wish to thank Gareth Barnes and Markus Bauer for MEG data used to test this method. Jonathan Le Roux and Maneesh Sahani provided useful advice for determining the linear component with square most similar to a quadratic component. Israel Nelken and Maria Chait provided useful comments on a previous version, as did two anonymous reviewers. This work was supported in part by a Royal Society International Joint Project grant, a grant from the Agence Nationale de la Recherche, and a grant from the CNRS Neuroinformatique programme.

References

- Blankertz, B., Tomioka, R., Lemm, S., Kawanabe, M., Müller, K.-R., 2008. Optimizing spatial filters for robust EEG single-trial analysis. *IEEE Signal. Process. Mag.* 25, 41–56.
- Boudet, S., Peyrodie, L., Gallois, P., Vasseur, C., 2007. Filtering by optimal projection and application to automatic artifact removal from EEG. *Signal. Process.* 87, 1978–1992.
- Cichocki, A., 2004. Blind signal processing methods for analyzing multichannel brain signals. *Int. J. Bioelectromagnetism* 6, 1–18.
- de Cheveigné, A., 2010. Time-shift denoising source separation. *J. Neurosci. Methods* 189, 113–120.
- de Cheveigné, A., Simon, J.Z., 2007. Denoising based on time-shift PCA. *J. Neurosci. Methods* 165, 297–305.
- de Cheveigné, A., Simon, J.Z., 2008a. Sensor noise suppression. *J. Neurosci. Methods* 168, 195–202.
- de Cheveigné, A., Simon, J.Z., 2008b. Denoising based on spatial filtering. *J. Neurosci. Methods* 171, 331–339.
- Delorme, A., Makeig, S., 2004. EEGLAB: an open toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* 134, 9–21.
- Duncan, K.K., Hadjipapas, A., Li, S., Kourtzi, Z., Bagshaw, A., Barnes, G., 2009. Identifying spatially overlapping local cortical networks with MEG. *Hum. Brain Mapp.* 7, 1003–1016.
- Efron, B., Tibshirani, R.J., 1993. Introduction to the bootstrap. *Monographs on Statistics and Applied Probability*. Chapman and Hall/CRC.
- Fukunaga, K., 1990. Introduction to Statistical Pattern Recognition, second ed. Academic Press.
- Hillebrand, A., Singh, K.D., Holliday, I.E., Furlong, P.L., Barnes, G.R., 2005. A New approach to neuroimaging with magnetoencephalography. *Hum. Brain Mapp.* 25, 199–211.
- Koles, Z.J., Lazar, M.S., Zhou, S.Z., 1990. Spatial patterns underlying population differences in the background EEG. *Brain Topogr.* 2, 275–284.
- Lindgren, J.T., Hyvärinen, A., 2007. Emergence of conjunctive visual features by quadratic independent component analysis. *Advances in Neural Information Processing Systems* 19, 897–90. MIT Press.
- Makeig, S., Bell, A.J., Jung, T.-P., Sejnowski, T.J., 1996. Independent component analysis of electroencephalographic data. *Adv. Neural Inf. Process. Syst.* 8, 145–151.
- Parra, L.C., Spence, C.D., Gerson, A.D., Sajda, P., 2005. Recipes for the linear analysis of EEG. *Neuroimage* 28, 326–341.
- Särelä, J. Exploratory source separation in biomedical systems. Technical University of Helsinki unpublished thesis, 2004.
- Särelä, J., Valpola, H., 2005. Denoising source separation. *J. Mach. Learn. Res.* 6, 233–272.
- Sekihara, K., Hild, K.E., Nagarajan, S.S., 2006. A novel adaptive beamformer for MEG source reconstruction effective when large background brain activities exist. *IEEE Transactions on Bio-medical Engineering* 53, 1755–1764.
- Valpola, H., Pajunen, P., 2000. Fast algorithms for Bayesian independent component analysis. *Proceedings of the Second International Workshop on Independent Component Analysis and Blind Signal Separation (ICA 2000)*, pp. 233–237.