NeuroImage 66 (2013) 672–680

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

Automatic denoising of single-trial evoked potentials

Maryam Ahmadi [a], Rodrigo Quian Quiroga [a,b,*]

[a] Department of Engineering, University of Leicester, UK
[b] Leibniz Institute for Neurobiology, Magdeburg, Germany

ARTICLE INFO

Article history:
Accepted 18 October 2012
Available online 7 November 2012

Keywords:
Event related potentials
Single trial analysis
Wavelet transform
Denoising
Visual evoked potentials
Auditory evoked potentials

ABSTRACT

We present an automatic denoising method based on the wavelet transform to obtain single trial evoked potentials. The method is based on the inter- and intra-scale variability of the wavelet coefficients and their deviations from baseline values. The performance of the method is tested with simulated event related potentials (ERPs) and with real visual and auditory ERPs. For the simulated data the presented method gives a significant improvement in the observation of single trial ERPs as well as in the estimation of their amplitudes and latencies, in comparison with a standard denoising technique (Donoho's thresholding) and in comparison with the noisy single trials. For the real data, the proposed method largely filters the spontaneous EEG activity, thus helping the identification of single trial visual and auditory ERPs. The proposed method provides a simple, automatic and fast tool that allows the study of single trial responses and their correlations with behavior.

Introduction

Event related potentials (ERPs) are voltage fluctuations within the Electroencephalogram (EEG) due to external stimulation or internal processes. They are routinely used for clinical diagnosis, as they allow the identification of dysfunctions along the visual, auditory and somatosensory pathways (Luck, 2005; Niedermeyer and Lopes da Silva, 2004; Regan, 1989). ERPs are also widely used in neuroscience research, given that the amplitude, latency and localization of different peaks or oscillatory patterns have been correlated to a large variety of sensory and cognitive functions (Quian Quiroga, 2006). Compared to single-neuron studies, the gold standard in neuroscience, ERPs, and EEGs in general, give only an indirect and noisy measure of the neuronal activity. The large advantage of ERPs, however, is that, unlike single-cell recordings which are rarely performed in humans (Quian Quiroga et al., 2005, 2008), their recording involves a non-invasive procedure with a relatively simple setup, and therefore, they continue to be one of the preferred tools for studying sensory and cognitive processes in human subjects.

One of the main problems in the analysis of ERP data is that the single-trial responses have a small amplitude compared to the ongoing EEG in which they are embedded. By far the most popular technique to enhance the observation of ERPs is by averaging several repetitions of the stimulus (Dawson, 1954; Lopes da Silva, 2004). However, the drawback of ensemble averaging is that critical information about trial-by-trial changes of the evoked responses is lost. In particular, the conventional approach in the design of an ERP paradigm is to try to avoid these single-trial fluctuations in order to get better-defined average responses. But there are many interesting questions that are in fact related to systematic or unsystematic trial-by-trial variations, such as those related to the study of learning processes (Quian Quiroga et al., 2007). Thus the need to develop algorithms to filter the background EEG activity in order to observe the single trial evoked responses. For this, the use of Wiener filtering was suggested (Walter, 1968). Wiener filtering minimizes the mean square estimation error of average evoked potentials and could in principle be used to denoise single trials. However, it is a time-invariant method — i.e. it assumes stationary of the signal — and it does not give optimal results when dealing with time-varying transient signals such as ERPs (Quian Quiroga, 2000; Quian Quiroga and Garcia, 2003). For the same reason, other standard digital filters are not suitable for the analysis of single-trial ERPs, given that ERPs are a series of waves appearing at different times and with different frequency compositions. To deal with the non-stationary issue, De Weerd and co-workers proposed a time-varying Wiener filter, which, however, could not provide a good reconstruction of the signal (De Weerd, 1981; De Weerd and Kap, 1981).

Another set of algorithms to filter the single-trial ERPs use wavelets. The wavelet transform is a time–frequency decomposition ideally suited for non-stationary signals (Mallat, 1999), which has been used in the analysis of ERPs since the early 1990s (Bartnik et al., 1992; Hanrahan, 1990; Quian Quiroga et al., 2001; Thakor et al., 1993). In particular, Bartnik et al. (1992)) proposed to use an algorithm based on a wavelet decomposition to extract single trial auditory evoked potentials from the ongoing EEG. This algorithm was unsupervised, but it led to large errors in the estimation of the single-trial ERPs. Following this approach, an ad-hoc wavelet denoising technique showed optimal results for the identification of the single trial responses (Quian Quiroga, 2000). Given that ERPs have specific time and frequency localizations, after wavelet decomposition, the idea is to reconstruct the signal but using only those

* Corresponding author at: Department of Engineering, University of Leicester, University Road , Leicester, LE17RH, UK.
E-mail address: rqqg1@le.ac.uk (R. Quian Quiroga).

1053-8119/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2012.10.062

coefficients related to the evoked responses (deleting the ones related to the ongoing EEG). An analysis with synthetic ERP data showed that this ad-hoc implementation improved the signal to noise ratio of the single-trial responses as well as the estimation of their latencies and amplitudes (Quian Quiroga and Garcia, 2003). But the main caveat of this method is that it requires a manual selection of the stimulus-related coefficients, using prior knowledge of the time and frequency ranges of the ERPs. This makes this denoising procedure subjective, time consuming and not practical for the analysis of large number of channels (given that the selection of wavelet coefficients is not necessarily the same for the different channels). To overcome these problems, in this study we propose an automatic denoising implementation to visualize the single trial evoked responses. The method is based on the wavelet transform and it introduces an automatic selection of wavelet coefficients based on the inter- and intra-scale correlation of neighboring wavelet coefficients, and how their values deviate from baseline. We show its performance with synthetic ERP data, as well as with real visual and auditory ERPs.

Materials and methods

Real data

EEG recordings

Recordings were performed in an electrically shielded chamber in 10 voluntary healthy subjects (18–30 years old). Subjects were seated comfortably in a chair and were asked to remain still and relax while they did a visual and an auditory oddball paradigm (see below). The EEG data was recorded continuously using 64 electrodes placed according to the international 10-20 system, band pass filtered between 0.1 Hz and 250 Hz and sampled at 512 Hz, using an average reference. After the recording, the EEG signals were re-referenced to the average of the left and right mastoids and trials that were contaminated with eye blinks were removed manually from each data set. For each trial, one second pre- and one second post-stimulation were stored for further analysis.

Visual oddball Paradigm

As in previous studies (Quian Quiroga and Schürmann, 1999), pattern visual event-related potentials (VEPs) were obtained with a checkerboard pattern (side length of checks: 50' visual angle). A sequence with two different stimuli was presented pseudo randomly ($N=250$ stimuli): the frequent or non-target stimuli were a colour reversal of checks (80% of the stimuli), while the less frequent or target stimuli were colour reversals with a half check displacement (both horizontal and vertical) of the pattern (20% of the stimuli). Subjects were asked to ignore the non-target stimuli and press a key whenever they saw the target ones. Each pattern reversal was shown for 1 sec and the inter stimulus interval varied pseudo-randomly between 2 and 2.2 sec. No two target stimuli appeared in succession. Subjects were asked to fixate on a small red circle in the centre of the screen during the recording (Quian Quiroga and Schürmann, 1999; Schürmann et al., 1995).

Auditory oddball paradigm

Auditory event-related potentials (AEP) were obtained with an oddball paradigm, using a sequence with two different tones: non-target stimuli (80%) had a frequency 2000 Hz and target stimuli (20%) a frequency 1000 Hz (Goodin et al., 1978). Subjects were instructed to press a key whenever they heard the target tone and ignore the non-target ones. Each stimulus was presented for 100 ms and the inter stimulus interval varied pseudo randomly between 1.5 and 1.7 sec. As with the VEP, subjects were asked to gaze on a small red circle in the centre of the screen during the recording to avoid eye movements.

Synthetic data

To evaluate the performance of the proposed algorithm, as in previous works (Quian Quiroga and Garcia, 2003), the typical ERP components obtained with a visual oddball, the P1, N2 and P3, were simulated using three Gaussian functions added to background EEG activity (Fig. 1). Random fluctuations in the latency of the simulated components were introduced in order to resemble the latency variability across single trials (ranges, P1: 90–125 ms, N2: 120–155 ms and P3: 400–700 ms). The background EEG activity was taken from the recording of one subject with eyes open fixating on a red circle in the centre of the screen. Thirty single trials of the noisy ERPs, 2 sec each, were generated with different signal to noise ratios (SNR). The SNR was defined as the ratio between the standard deviations of the simulated ERPs and the one of the background EEG activity.

The performance of the algorithm was quantified by the root mean square error (RMS) of the denoised single-trial ERPs — i.e. the difference between the denoised signal and the original "clean" signal without background EEG — and was compared to the performance obtained with the simulated ERPs (without denoising) and the one obtained using a standard denoising implementation (Donoho, 1993) (see next section). Moreover, given that the most important information to be extracted from a single trial ERP analysis is typically the amplitude and latency of the single trial responses, we also quantified the error in the estimation of the single trial amplitudes and latencies of the ERPs. For this, for each peak in each single trial, a time window around each component was chosen (P1: 55–155 ms, N2: 95–170 ms, P3: 300–700 ms) and the maximum (or minimum) peak in the corresponding window was identified. The error for each single trial amplitude (latency) was defined as: $e = \langle|\hat{x}_i - x_i|\rangle_i$, where x_i is the actual and $\hat{x}_i$ is the estimated amplitude (latency) of the simulated ERP component. For each SNR, the statistical difference between the different methods was assessed with pairwise t-tests.

Wavelet denoising

Wavelet transform

The wavelet transform is the inner product of a signal with dilated and translated versions of a wavelet function (Mallat, 1999). Given a signal $x(t)$ and a wavelet function $\psi_{a,b}(t)$ the continuous wavelet transform (CWT) is defined as:

$$W_\psi X(a,b) = \left\langle x, \psi_{a,b} \right\rangle, \tag{1}$$

$$\psi_{a,b} = |a|^{-1/2} \psi\left(\frac{t-b}{a}\right), \tag{2}$$

where $a,b \in \Re$ are the scale and translation parameters, respectively. The translation parameter changes the location of the wavelet function, while the scaling parameter dilates or compresses it (Grossmann and Morlet, 1984). The correlation of the signal $x(t)$ with the dilated (contracted) versions of the wavelet $\psi_{a,b}(t)$ gives the low (high) frequency components. The CWT is very redundant and, without any loss of information, it is practical to define the wavelet transform only at discrete scales $a_j = 2^j$ and times $b_{j,k} = 2^j k$, which is called the dyadic wavelet transform (DWT). The dyadic wavelet transform can be computed using a hierarchical and very efficient algorithm — faster than the Fast Fourier Transform — called multiresolution decomposition (Mallat, 1999). This algorithm successively divides the signal into coarse approximations and details at different scales. The end result is the decomposition of the original signal into a series of detail scales and a final approximation, corresponding to the time-localized activity in different frequency bands.

In this study we used a 5-scale decomposition of the ERPs, obtaining the detail levels D1 to D5 and a final approximation A5. The lower scales