34th Annual International Conference of the IEEE EMBS
San Diego, California USA, 28 August - 1 September, 2012

Nonlocal Two Dimensional Denoising of Frequency Specific Chirp Evoked ABR Single Trials

J. Kristof Schubert[1], Tanja Teuber[2], Gabriele Steidl[2], Daniel J. Strauss[1,3,4] and Farah I. Corona–Strauss[1]

***Abstract*— Recently, we have shown that denoising evoked potential (EP) images is possible using twodimensional diffusion filtering methods. This restoration allows for an integration of regularities over multiple stimulations into the denoising process. In the present work we propose the nonlocal means (NLM) method for EP image denoising. The EP images were constructed using auditory brainstem responses (ABR) collected in young healthy subjects using frequency specific and broadband chirp stimulations. It is concluded that the NLM method is more efficient than conventional approaches in EP imaging denoising, specially in the case of ABRs, where the relevant information can be easily masked by the ongoing EEG activity, i.e., signals suffer from rather low signal-to-noise ratio SNR. The proposed approach is for the a posteriori denoising of single trials after the experiment and not for real time applications.**

I. INTRODUCTION

Providing a stimulus (auditory, visual, somatosensory etc.) to a subject while recording an electroencephalogram (EEG), one can define equally sized segments of the EEG being time-locked to the stimulus. These epochs $\mathbf{s}_i(t)$, $i = 1, ..., N$, called *single–sweeps*, *single–trials* or in a nutshell *sweeps* may show voltage changes specifically related to the brain's response to the stimulus. These voltage changes are called *evoked potentials* (EPs). Auditory evoked potentials appearing in the first ten milliseconds are referred to *auditory brainstem responses* (ABRs). Jewett and Williston [5] suggest a chronological roman numbering of the elicited dominant peaks from *I* to *VII* in the ABR. Wave *V* appears to be the most prominent wave with the highest amplitude. Based on this observation, the detection of wave *V* has grown to be a suitable indicating method for providing information regarding auditory function and hearing sensitivity. One prominent application is the newborn hearing screening (NHS) (see e.g. [3]), where hearing thresholds are objectively determined by means of detecting the wave *V*. Further, by measuring the absolute latencies of the first five waves and the inter–peak intervals $I - III$, $III - V$ and $I - V$, retrocochlear pathologies and frequency-specific estimation of auditory sensitivity can be discovered, too. Being neural correlates of sensory and cognitive processing, ABRs and their analysis will be applicable in clinical neuroelectrodiagnostics. The direct analysis of the single–sweeps is difficult as background activity (EEG components being not related to the event or stimulus), background noise or artifacts mask the ABR component resulting in a bad signal–to–noise ratio (SNR). To obtain a meaningful ABR waveform, the sweeps can be pointwise averaged, i.e., $\mathbf{s}(t) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{s}_i(t)$ (see Fig. 1 (top)). Based on the so–called *signal–plus–noise model* which states that a signal only consists of a stimulus-locked evoked component plus independent stationary background noise, this averaging method shows an improvement of the SNR with $\sqrt{N}$. The averaging technique is the most commonly used approach to detect small stimulus-locked and exogenous evoked potentials embedded in noise. But considering morphological changes of ABR components *during* measurements as a result of effects of non–stationary endogenous conditions such as attention or vigilance, obviously these changes will not be visible in the final average. Even changing a solution strategy during a task affects the morphology of an ABR and is therefore not visible, too. Furthermore, the analysis of individual sweeps per se is not possible due to a poor SNR. In [6] we introduced a two–dimensional image processing technique for the denoising of single-sweeps by means of the so-called *ERP image* (single–sweeps in matrix representation, see Fig. 1 (bottom) for an example). The amplitudes of the sweeps are encoded in a color–scale map (yellow to white colors represent high values and dark red to black colors represent small values). We keep the idea of two-dimensional single-sweep denoising. In this paper we introduce a nonlocal approach in order to denoise ABR images (ABRIs).

II. MATERIALS AND METHODS

A. Denoising of ABRIs by Nonlocal Means

Let $\mathcal{A} = \{\mathbf{s}_n \in \mathbb{R}^M : n = 1, 2, ..., N\}$ be a set of N sampled ABR single–sweeps within the time interval $[0, M/f_s]$ (f_s is the sampling frequency) of a particular experiment. From $\mathcal{A}$ we can construct the ABRI $\mathbf{S} \in \mathbb{R}^{N\times M}$ such that $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_N)^T$. As mentioned in the previous section the evoked individual sweeps $\mathbf{s}_n$ of $\mathbf{S}$ are time-locked to the stimulation for fixed stimulus settings. Hence, those responses result in morphologically stable transient potential fluctuations that form the typical ABR component (Fig. 1 (top)). Such reproducible fluctuations yield correlated intervals in the rows of $\mathbf{S}$ and induces self–similarity in the ABRI.

The nonlocal mean (NL–mean) method suggested independently in [9] and [10] exploit the self–similarity in the ABRI by so–called image patch methods. Here, we adjust the NL–means algorithm for our ABR image denoising problem and

[1]Systems Neuroscience & Neurotechnology Unit, Saarland University, Faculty of Medicine, Neurocenter, Building 90.5, D–66421 Homburg/Saar, Germany strauss at snn-unit.de

[2]Mathematical Image Processing and Data Analysis, Technical University Kaiserslautern, Felix Klein Zentrum, D–67663 Kaiserslautern, Germany

[3]Leibniz–Institut fuer Neue Materialien gGmbH, Campus D2 2, D–66123 Saarbruecken, Germany

[4]Key Numerics, Nelkenstrasse 31, D-66119 Saarbruecken, Germany

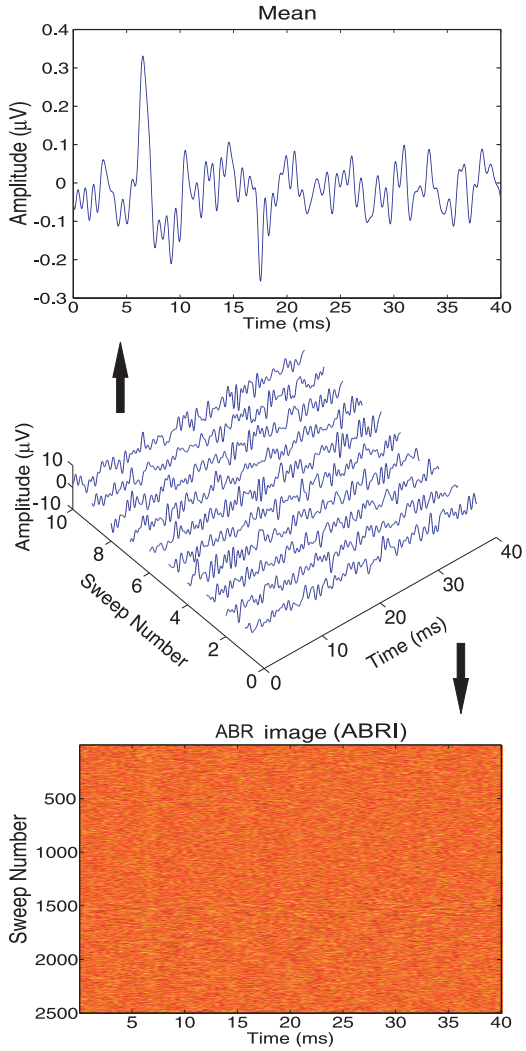

Fig. 1. Top: The commonly used averaged signal showing the wave *V* at approximately 7ms; Middle: Single sweeps (individual responses); Bottom: Single sweeps in matrix representation, where the amplitude is encoded in a color-scale map (yellow to white colors represent high values and dark red to black colors represent small values). A very fine vertical trace at around 7ms is discernable. Possible amplitude fluctuations or latency shifts would be visible in the ABRI, but not in the mean response. All time axes start at stimulus offset.

refer to [9], [10] for more general discussions and for general performance comparisons.

In contrast to conventional one–dimensional ABR denoising procedures $\mathcal{D}_1 : \mathbf{s}_n \mapsto \tilde{\mathbf{s}}_n$, where $\tilde{\mathbf{s}}_n \in \mathbb{R}^M$ is a denoised version of $\mathbf{s}_n$, we are interested in a denoising operator $\mathcal{D}_2 : \mathbf{S} \mapsto \mathbf{Q}$, where $\mathbf{Q} \in \mathbb{R}^{N\times M}$ is the denoised version of the two–dimensional ABRI $\mathbf{S}$. From $\mathbf{Q} = (\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n)^T$ the individual denoised sweeps $\mathbf{q}_n \in \mathbb{R}^M$ $(n = 1, 2, \ldots, N)$ can be extracted after the two–dimensional denoising process. Physiologically meaningful amplitude and latency changes can be extracted from $\mathbf{q}_n$, see Fig. 3 for an example. Thus the $\mathcal{D}_2$–denoising approach allows for an integration of regularities over multiple stimulations into the denoising process. Note that the inclusion of such regularities, although significant in ABR experiments, is completely missing in $\mathcal{D}_1$–denoising schemes. To apply the NL–means algorithm for the aforementioned $\mathcal{D}_2$ denoising of ABR images, each sample (i.e., entry or pixel in $\mathbf{S}$) s_i, $i = 1, \ldots, J$ with $J = NM$, is replaced by a denoised version q_i in the following way: Each pixel s_i is compared together with its neighborhood to other ABRI patches (i.e., the very same neighborhoods of other pixels s_j). For each comparison a weight coefficient $\xi_{i,j} \in \mathbb{R}$ $(i, j = 1, \ldots, J$ is assigned to the center pixel s_i, depending on the similarity of the image patches, i.e., the similarity of the neighborhood of s_i to the neighborhood of s_j. The restored/denoised entry q_i is now the weighted average of all the surrounding entries in $\mathbf{S}$ such that

$$q_i = \frac{1}{\gamma_i} \sum_{j=1}^{J} \xi_{i,j} s_j, \tag{1}$$

with $\gamma_i = \sum_{j=1}^{J} \xi_{i,j}$. Let $\mathcal{I}$ be an appropriate index set such that the two–dimensional patches of the ABRI with the centers s_i and s_j are given by $s_{i+\mathcal{I}}$ and $s_{j+\mathcal{I}}$. We further introduce the vector $\boldsymbol{\varphi}_\sigma = (\varphi_{\sigma,k})_{k\in\mathcal{I}}$ which denotes a sampled version of a two–dimensional Gaussian kernel with standard deviation $\boldsymbol{\sigma}$. The weights which quantify the similarity of $s_{i+\mathcal{I}}$ to $s_{j+\mathcal{I}}$ are now given by

$$\xi_{i,j} = \exp\left(-\frac{1}{\lambda} \sum_{k\in\mathcal{I}} \varphi_{\sigma,k} |s_{i+\mathcal{I}} - s_{j+\mathcal{I}}|^2 \right). \tag{2}$$

The parameter $\boldsymbol{\sigma}$ is steering the influence of neighboring pixels on the weight and $\lambda > 0$ is controlling the amount of denoising. The application of (1) to all pixels in the ABRI finally yields its denoised version $\mathbf{Q}$. Each row in $\mathbf{Q}$ is now the denoised single sweep $\mathbf{q}_n$ of $\mathbf{s}_n$ $(n = 1, \ldots, N)$, but in contrast to conventional 1D procedures, stemming from a D_2–denoising process. In order to reduce the computational load, we use the "search window" implementation proposed in [9] in our computations. We used a 15×15 patch and $\boldsymbol{\sigma} = (0.5, 5.0)^T$ for the numerical experiments in Sec. III.

B. Stimulus Selection

It is well known that the cochlear is tonotopically organized [1]. That means, different areas of the cochlea are only sensitive to certain frequencies. The *base* is more sensitive to higher frequencies and the *apex* is more sensitive to lower frequencies. Therefore, a signal containing low frequencies takes a longer time to reach the sensation locus. Compared to clicks, the high frequencies of a chirp are delayed from the low frequencies to compensate the temporal dispersion of the basilar membrane and to stimulate specific areas along the cochlear. Additionally, the chirps are designed to start and end exactly with zero to avoid stimulation of undesired areas due to abrupt onset or offset stimulation. It has been shown that chirp stimulations show better performance for evoking ABRs than click stimulations, especially at low intensity levels [8]. Further, Gorga et al. [4] showed, that the latency of the wave *V* is related to the frequency and to the intensity of the stimulus. We decided to use five different chirps for our ABR measurements (see Table I and [2] for detailed

information). One is a broadband chirp containing frequencies from 100Hz to 10kHz and the remaining four chirps are sub-bands containing only certain frequency intervals.

TABLE I
FREQUENCY SPECIFIC CHIRPS

Chirp	Center Frequency (Hz)	Interval (Hz)	Duration (ms)
Ch1	302	[108,490]	6.19
Ch2	813	[495,1135]	2.02
Ch3	1915	[1230,2600]	0.88
Ch4	6725	[2950,10500]	0.51
B-b	5050	[100,10000]	10.12

C. ABR Data

The chirp-evoked ABRs were collected from ten student volunteers (mean age: 25.1 ± 3 years; four female, six male) with no history of hearing problems and normal hearing thresholds (below 15dB (HL)) checked by an audiogram. Ag/AgCl electrodes were attached (according to the 10–20 system) ipsilateral to the stimulus at the right mastoid (A1), at the vertex (Cz) as a common reference and at the upper forehead (Fpz) serving as a ground. Impedances were kept below 5kΩ during the measurements. A total of 3000 sweeps (artifacts were removed by an amplitude threshold of $15\mu V$) were recorded for each chirp. The ABRs were acquired by a high-end 24 bit biosignal amplifier (gUSBamp, gTec, Austria) using a sampling frequency of 19.2 kHz and a bandpass filter with low and high cutoff frequencies of 0.1Hz and 1.5kHz, respectively. Additionally, a 50Hz notch filter is applied.
We used the EEG simulator of [12], [13] to generate a set of synthetic data with a SNR of 1 with the classical signal–plus–noise approach. The synthetic ABR data were trials of the mean of our ABR measurements.

D. Objective Evaluation Criteria

To estimate the residual noise, we denote the averaged ABR (mean) by $\bar{\mathbf{s}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{s}_i$. The residual noise estimation is computed as suggested in [11] by

$$g_m = \sqrt{\frac{1}{N(N-1)}\sum_{n=1}^{N}(s_{n,m} - \bar{s}_m)^2}, \quad (m = 1, 2, ..., M) \tag{3}$$

which we further reduce to a residual noise quantifying scalar value by

$$\alpha = \|g\|_2^2 = \sum_{m=1}^{M} g_m^2 \tag{4}$$

The *ABR–reproducibility* allows for another quantification of the ABR quality. For this, we decompose the ABRI $\mathbf{S}$ into two sub–matrices $\mathbf{S}^e$ and $\mathbf{S}^o$ which carry the responses for even (upper index e) and odd (upper index o) numbered stimulations n ($n = 1, 2, \ldots, N$, N even) or rows of $\mathbf{S}$, respectively. Let us denote the averaged even and odd ABR data by $\bar{\mathbf{s}}^{e/o} = \frac{1}{N}\sum_{n=1}^{N/2}\mathbf{s}_n^{e/o}$ ($\bar{\mathbf{s}}^{e/o} \in \mathbb{R}^M$) and the additional average over the time by $\tilde{s}^{e/o} = \frac{1}{M}\sum_{m=1}^{M}\bar{s}_m^{e/o}$. The ABR–reproducibility is now just the Pearson's correlation coefficient $\beta \in [-1, 1]$ between the average of the even and odd numbered trials

$$\beta = \frac{\sum_{m=1}^{M}(\bar{s}_m^e - \tilde{s}^e)(\bar{s}_m^o - \tilde{s}^o)}{\sqrt{\sum_{m=1}^{M}(\bar{s}_m^e - \tilde{s}^e)^2}\sqrt{\sum_{m=1}^{M}(\bar{s}_m^o - \tilde{s}^o)^2}} \tag{5}$$

III. RESULTS

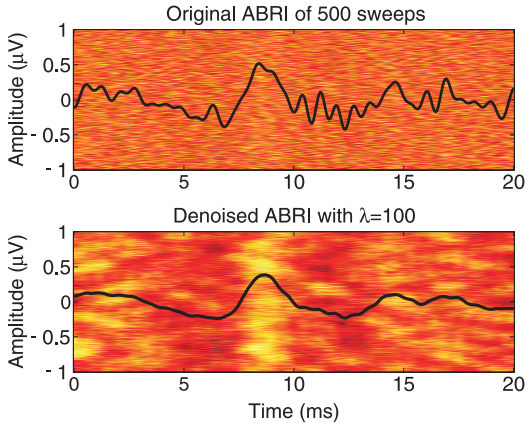

Fig. 2. Top: ABRI of 500 sweeps (where the amplitude is encoded in a color-scale map (yellow to white colors represent high values and dark red to black colors represent small values)) of one subject with the respective mean plotted on top. Stimulus: frequency-specific chirp "Ch2" presented at 50 dB (SPL). Bottom: NL–mean processed ABRI. The trace of the wave *V* is clearly visible after applying the NL–mean method. Possible amplitude fluctuations or latency shifts could be made visible in the ABRI. Note how the NL–mean method affects the average. Time axes start at stimulus onset.

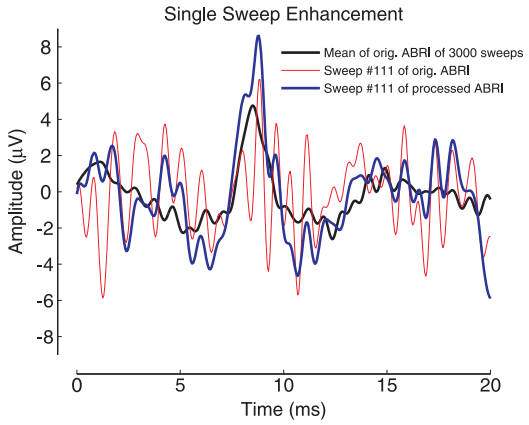

Fig. 3. Single–sweep enhancement. One sweep (number 111) of the unprocessed (red) and processed (blue) ABRI is compared to each other and to the average (black) of a whole measurement (3000 sweeps). Note the smoothing effect of the NL–mean method (blue) to the average. The processed sweep number 111 clearly shows the prominent wave *V* peak whereas the unprocessed sweep does not.

Fig. 2 shows exemplarily the result of applying the NL–mean method to an ABRI. The wave *V* trace is clearly visible. This implies an individual analysis of the trials which can be seen in Fig. 3. Here, one single–trial is compared to its denoised version and to the overall mean of 3000 sweeps. The denoised sweep recovers its morphology.
In section II-D we estimate the residual noise scalar value α by calculating the variances over the sweeps. Larger variances result in larger values of α.

Analyzing how well the structure of the ABRI is preserved is an impossible task since the noise model and thus the "real sweep" is unknown. However, in the case of synthetic data, the noise is superimposed to the signal according to the model described in II-C. In order to evaluate the deterioration of the original sweep morphology, we computed the Pearson's correlation coefficient between the denoised sweeps from the signal–plus–noise model to the sweeps without added noise. The result for the synthetic data for the NL–mean method is shown in Fig. 4.

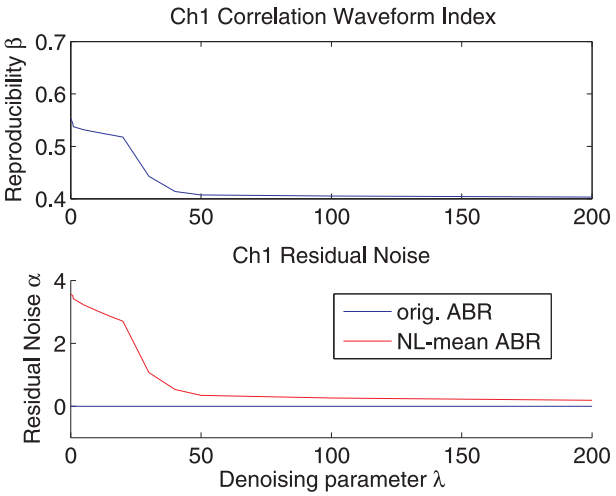

Fig. 4. Top: The Pearson's correlation coefficient β between the NL–mean processed ABR and the ABR without noise. The morphology is preserved by increasing λ. Bottom: The residual noise of the NL–mean processed ABR decreases by increasing λ. The residual noise of the the original ABRI is due to its synthetic properties very low compared to the real data.

IV. DISCUSSION

Time–locked responses result in typical traces in the ABRI. Depending on the stimulus settings, these traces vary. Obviously, those changes will not be visible in an average. The main advantage of using the NL–mean processed ABRIs is the possibility to a posteriori analyze *every* individual sweep extracted from the denoised image with respect to the clinical relevant and experimentally established parameters such as latency or amplitude.

Fig. 2 shows the improvement of the detection of the wave *V* after denoising. Due to the better visibility of the wave *V* trace, the ABRIs are suitable for the analysis of the common parameters. Additionally, possible fluctuations in amplitude or latency shifts during the measurement could be made visible. The importance of the visualization of these variations during the experiment is obviously, if we think of the correlation of tinnitus-phenomena with prolonged latencies [7]. In contrast to conventional $\mathcal{D}_1$ approaches which only consider pre- and post–stimulus interval, regularities in the data over multiple stimulations can also be integrated into the $\mathcal{D}_2$–denoising processes.

The paper has not the intention to estimate or to tune the parameters used in the method. It is important to know, that this two–dimensional denoising approach can be adapted to individual applications, improving the performance. Note that the method as presented in this paper is not suitable for real time processing. However, the results show a significant improvement in the wave *V* detection and allow the analysis of individual sweeps while preserving their morphology (see for Fig. 4 (top)).

V. CONCLUSIONS

In this paper we have shown, how the self–similarity of an ABRI can be exploited for the denoising by nonlocal means. It is shown, that the resulting ABRIs contain less residual noise by preserving their morphology and allow for a clear identification of the wave *V*. It is concluded that the NL–mean method might be a promising new approach in the field of denoising ABR data and *a posteriori* single sweep analysis.

REFERENCES

[1] E. de Boer, Auditory physics. Physical principles in hearing theory I, Phys. Rep.,1980: 62:87-174.

[2] F. I. Corona-Strauss, D. J. Strauss, B. Schick, W. Delb, A series of notched-noise embedded chirps for objective frequency specific hearing examinations, Conf. Proc. IEEE Eng. Med. Biol. Soc., 2009: 2074–7.

[3] F. I. Corona-Strauss, W. Delb, M. Bloching, D. J. Strauss, Ultra-fast quantification of hearing loss by neural synchronization stabilities of auditory evoked brainstem activity, Conf. Proc. IEEE Eng. Med. Biol. Soc.,2007: 2476–9.

[4] M. P. Gorga, J. R. Kaminski, K. A. Beauchaine, Auditory Brainstem Responses to Tone Bursts in Normally Hearing Subjects. Journal of Speech and Hearing Research, 1988: Vol.31 87–97.

[5] D. L. Jewett, J. S. Williston, Auditory-Evoked Far Fields Averaged From The Scalp Of Humans. Brain, 1971: 94(4):681-696.

[6] I. Mustaffa, C. Trenado, K. Schwerdtfeger, D. J. Strauss, Denoising of single-trial matrix representations using 2D nonlinear diffusion filtering, Journal of Neuroscience Methods, 2010: 185, pp.284–292.

[7] U. Rosenhall, A. Axelsson, Auditory brainstem response latencies in patients with tinnitus. Scand Audiol., 1995: 24(2):97-100.

[8] O. Wegner, T. Dau, Frequency specificity of chirp-evoked auditory brainstem responses. Acoustical Society of America, 2002.

[9] A. Buades, B. Coll, and J.-M. Morel. A non–local algorithm for image denoising. In *IEEE CVPR.*, volume 2, pages 60–65, San Diego, CA., 2005. IEEE Computer Society Press.

[10] S. P. Awate and R. T. Whitaker. Unsupervised, information–theoretic, adaptive image filtering for image restoration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(3):364–376, 2006.

[11] M. Granzow H. Riedel and B. Kollmeier. Single–sweep-based methods to improve the quality of auditory brain stem responses, part ii: Averaging methods. *Zeitschrift fuer Audiologie/Audiological Acoustics*, 40:62–85, 2001.

[12] N. Yeung, R. Bogacz, C. B. Holroyd, S. Nieuwenhuis, and J. D. Cohen. Theta phase resetting and the error–related negativity. Psychophysiology, 44:39–49, 2007.

[13] N. Yeung, R. Bogacz, C. B. Holroyd, and J. D. Cohen. Detection of synchronized oscillations in the electroencephalogram: An evaluation of methods. Psychophysiology, 41:822–832, 2004.