Ensemble-based Time Alignment of Biomedical Signals

A Garde[1,2,3], P Laguna[3,4], B F Giraldo[1,2,3], R Jané[1,2,3], L Sörnmo[5]

[1]Department of ESAII, Universitat Politècnica de Catalunya, Spain;
[2]Institut de Bioenginyeria de Catalunya (IBEC), Barcelona, Spain;
[3]CIBER de Bioingeniería, Biomateriales y Nanomedicina, Spain;
[4]Communications Technology Group, I3A, University of Zaragoza, Spain;
[5]Department of Electrical and Information Technology, Lund University, Sweden.

Abstract

In this paper, the problem of time alignment is revisited by adopting an ensemble-based approach with all signals jointly aligned. It is shown that the maximization of an eigenvalue ratio is synonymous to maximizing the signal-to-jitter-and-noise ratio. Since optimization of this criterion is extremely time consuming, a relaxed optimization procedure is introduced which converges much more quickly. Using simulations based on respiratory flow signals, the results suggest that the time delay error variance of the new method is much lower than that obtained with the well-known Woody's method.

Keywords *Time alignment, Signal ensemble, Subsample precision, Eigenvalue decomposition.*

1 Introduction

Time alignment is a classical problem in signal processing which has been extensively treated in many biomedical applications, including event-related brain potentials, conduction estimation in electromyography, and cardiac late potentials. The goal is often to align an ensemble of observed signals with similar shape so that noise reduction can be achieved through ensemble averaging [1].

Despite the long-standing interest in methods for time alignment, it is surprising that the methods which are intrinsically ensemble-oriented have received so limited attention. The guiding design principle has rather been to extend an alignment method which compares each signal of the ensemble to a reference signal, thus relying on pairwise alignment in the time [2], frequency [3, 4], or scale domain [5]. Such an approach is, however, empirical in nature and does not ensure that the resulting time delay estimates of the ensemble are optimal.

Woody's method is probably the most well-known method for ensemble time alignment, with matched filtering as its core operation [6]. An initial estimate of the filter's impulse response is obtained by averaging the unaligned ensemble. An iterative procedure is then applied by which the time delay of each signal of the ensemble is estimated, a new ensemble average is computed, and so on, until the time delays no longer change. Although no general proof of convergence has been presented, it is well-known that convergence is generally achieved within 5–10 iterations, provided that the waveforms are initially reasonably well-aligned and that the signal-to-noise ratio is reasonably good. The Woody method has later been extended so that the colored noise situation can also be handled [7], not only white noise which was implicit to the original work.

It was not until 2008 the time alignment problem was recast into a formulation by Cabasson and Meste which accounts for all the waveforms of the ensemble [8]. They proposed a statistical model framework in which each observed signal was assumed to be composed of an unknown, fixed-amplitude signal with unknown delay and additive, Gaussian white noise. The joint maximum likelihood (ML) estimator of the unknown signal waveform and the time delays was derived, and found to differ slightly from the procedure that defines Woody's method. The main difference is that the ensemble average is no longer needed when estimating the time delays. Since an iterative version of the estimator was implemented, optimality in the ML sense could not be guaranteed. Simulation results showed that the modified method produced time delay estimates with lower error variance for small ensemble sizes (i.e., containing less than 20–25 waveforms), and offered slightly faster convergence.

Maximum a posteriori estimation (MAP) has been investigated for a generalized statistical model in which each signal waveform is multi-component and characterized by random amplitude and time delay of each component [9]. Since the noise was assumed to be Gaussian, the MAP approach led to the minimization of a quadratic cost function whose solution entails matched filtering and a constrained time delay search intervals. Although the authors outlined the optimal solution in their paper, they implemented and studied nevertheless a suboptimal estimator which was computationally more efficient than the optimal one. Since the time delay estimates were determined iteratively using matched filtering, it was concluded that the Woody's method can be viewed as a suboptimal implementation of the MAP estimator. When a uniform probability density function (PDF) of the time delays is assumed, and the multi-component, variable-amplitude signal replaced by a single-component, fixed-amplitude signal, the alignment method becomes closely related to that in [8].

This paper introduces a novel approach to time alignment where the eigenvalue properties of the correlation matrix for the entire ensemble of data are explored. No assumptions on time delay PDF have to be made, nor is an estimate of the underlying signal waveform required as is the case in Woody's method. These differences are due to the simple fact that an ensemble of any similar-shaped waveforms produces an observation matrix whose ratio of the largest eigenvalue and the sum of remaining eigenvalues is larger when the waveforms are aligned. It can be shown that this eigenvalue ratio has the attractive interpretation of a signal-to-jitter ratio.

2 Methods

2.1 Signal model and cost function

Time alignment of a signal ensemble takes its starting point in the following simple statistical model where the observed signals $x_i(n), i = 1, \ldots, M$, are all assumed to be characterized by

$$x_i(n) = s(n - \theta_i) + v_i(n), \quad n = 0, \ldots, N-1, \quad (1)$$

where $s(n)$ is a signal with energy E_s which is constant across the ensemble, θ_i the unknown, integer-valued time delay to be estimated, and $v_i(n)$ zero-mean, white noise with variance σ_v^2. The amplitude of $s(n)$ is assumed to be fixed and equal to one. The entire ensemble is compactly represented by the column matrix

$$\mathbf{X} = \begin{bmatrix}\mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_M\end{bmatrix}, \quad (2)$$

where the i-th column contains the samples $x_i(n)$,

$$\mathbf{x}_i = \begin{bmatrix} x_i(\Delta) \\ x_i(\Delta+1) \\ \vdots \\ x_i(N+\Delta-1) \end{bmatrix}. \quad (3)$$

where Δ is the start sample index. When all signals $\mathbf{x}_i$ are perfectly aligned, the corresponding correlation matrix is given by [1]

$$\mathbf{R}_x \triangleq E\left[\mathbf{x}\mathbf{x}^T\right] = \mathbf{s}\mathbf{s}^T + \sigma_v^2\mathbf{I}, \quad (4)$$

where the vector $\mathbf{s}$ contains the signal samples $s(n)$. The eigenvalues of $\mathbf{R}_x$ equal

$$\lambda_i = \begin{cases} E_s + \sigma_v^2, & i = 1; \\ \sigma_v^2, & i = 2, \ldots, N. \end{cases} \quad (5)$$

The eigenvalue ratio

$$\Lambda \triangleq \frac{\lambda_1}{\sum_{i=2}^{N}\lambda_i} \quad (6)$$

is here proposed as a performance measure of time alignment because it reaches its maximum when all signals are perfectly alignment, i.e., the assumption which leads to (5) and a rank-one correlation matrix $\mathbf{R}_x$.

Interestingly, Λ can be interpreted as a signal-to-jitter-and-noise ratio. This fact is shown by considering the continuous-time counterpart to the signal model in (1), given by

$$x_i(t) = s(t - \theta_i) + v_i(t). \quad (7)$$

The same assumptions apply for the discrete-time case, except that θ_i is a real-valued instead of integer-valued random variable. Assuming that θ_i is small and zero-mean with variance σ_θ^2, the observed signal can be approximated by

$$x_i(t) \approx s(t) - \theta_i s'(t) + v_i(t), \quad (8)$$

where $s'(t)$ denotes the first derivative of $s(t)$. The signals $s(t)$ and $s'(t)$ are orthogonal, i.e.,

$$\begin{aligned}\int_{-\infty}^{\infty} s(t)s'(t)dt &= \frac{1}{2\pi}\int_{-\infty}^{\infty} S(\Omega)(-\jmath\Omega)S^*(\Omega)d\Omega \\ &= -\frac{1}{2\pi}\int_{-\infty}^{\infty} \jmath\Omega|S(\Omega)|^2 d\Omega = 0, \quad (9)\end{aligned}$$

where $S(\Omega)$ denotes the Fourier transform of $s(t)$. It is here tacitly assumed that $s(t)$, which is a finite duration signal, can be extended from $-\infty$ to ∞.

Assuming that the observed signal $x_i(t)$, expressed by (8), has been sampled at the Nyquist rate, the correlation matrix of the sampled counterpart is given by

$$\mathbf{R}_x = \mathbf{s}\mathbf{s}^T + \sigma_\theta^2\mathbf{s}'\mathbf{s}'^T + \sigma_v^2\mathbf{I}, \quad (10)$$

where use has been made of the sampled counterpart to (9), stating that the vectors $\mathbf{s}$ and $\mathbf{s}'$ are orthogonal. It can be shown that $\mathbf{s}$ and $\mathbf{s}'$ are eigenvectors, with eigenvalues $E_s + \sigma_v^2$ and $\sigma_\theta^2 E_{s'} + \sigma_v^2$, respectively, whereas the remaining eigenvectors can be arbitrarily chosen as long as they are orthogonal to $\mathbf{s}$ and $\mathbf{s}'$. Thus, the eigenvalues of $\mathbf{R}_x$ equal

$$\lambda_i = \begin{cases} E_s + \sigma_v^2, & i = 1; \\ \sigma_\theta^2 E_{s'} + \sigma_v^2, & i = 2; \\ \sigma_v^2, & i = 3, \ldots, N, \end{cases} \quad (11)$$

where $E_s = \mathbf{s}^T\mathbf{s}$ and $E_{s'} = \mathbf{s}'^T\mathbf{s}'$ denote the energy of the signal and its derivative, respectively. Inserting the eigenvalues in (6), it is easy to show that Λ can be interpreted as a signal-to-jitter-and-noise ratio,

$$\Lambda(\boldsymbol{\theta}) = \frac{\lambda_1(\boldsymbol{\theta})}{\sum_{i=2}^{N}\lambda_i(\boldsymbol{\theta})} \approx \frac{E_s}{\sigma_\theta^2 E_{s'} + (N-1)\sigma_v^2} \quad (12)$$

where the numerator $E_s + \sigma_v^2$ has been approximated by E_s since $\lambda_1 \gg \lambda_N$ in many biomedical applications of interest. The dependence of the eigenvalues on the time delay vector $\boldsymbol{\theta}$, defined by $\theta_1, \ldots, \theta_M$, is indicated in (12).

2.2 Brute force integer optimization

The sample correlation matrix

$$\hat{\mathbf{R}}_x = \mathbf{X}\mathbf{X}^T \quad (13)$$

is computed, its eigenvalues determined, and the optimal time delays are determined from

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \Lambda(\boldsymbol{\theta}), \tag{14}$$

i.e., by finding those column shifts $\hat{\theta}_1, \ldots, \hat{\theta}_M$ which maximize the cost function $\Lambda(\boldsymbol{\theta})$. Each column is shifted symmetrically around its initial position, i.e., it is shifted Δ locations upwards as well as downwards.

The maximization in (14) is implemented with a brute force technique which tests all possible shift combinations, amounting to about $(2\Delta+1)^M$ shifts. This implies that the singular value decomposition must be computed about 10^{84} times for the case when $\Delta = 3$ and $M = 100$, this being a case which is representative for many applications. Evidently, this amount of computation is far too demanding to be of any practical use, and therefore it is essential to find an alternative, more efficient approach to multidimensional optimization.

2.3 Relaxed optimization

A fruitful approach to dramatically speed up the optimization is to reformulate (14) so that the integer-valued time delays $\theta_1, \ldots, \theta_M$ become continuous-valued, i.e., to relax the optimization problem, thereby making it possible to employ an algorithm which does not require a gradient of the cost function. The first step in such an approach is to increase the original sampling rate of the augmented ensemble $\tilde{\mathbf{X}}$ through interpolation so that the signals become good approximations of the continuous-time signals. The interpolated signals are denoted $x_{i,c}(t)$. In practice, this conversion is accomplished by interpolating the signals by a factor of 50 (or similar), and then applying a zero-order hold operation to produce a signal which is continuous in time.

For every set of time delays $\theta_1, \ldots, \theta_M$, now treated as continuous-valued variables, the column vectors of $\mathbf{X}$ are determined by resampling the continuous-time signals $x_{i,c}(t)$, using the following expression

$$x_i(n) = x_{i,c}(nT_s+\theta_i), \quad n = 0, \ldots, N+2\Delta-1, \tag{15}$$

where T_s denotes the sampling interval of the original observed signal. Similar to brute force integer optimization, the eigenvalues of $\hat{\mathbf{R}}_x$ are then determined, and so on. Since it is difficult to derive an analytical expression of the gradient of $\Lambda(\boldsymbol{\theta})$, the well-known Nelder–Mead simplex algorithm (Matlab implementation) is here employed for optimization.

The assumption of an integer-valued time delay θ_i in (1) is obviously due to that the observed signal is sampled. However, θ_i is typically continuous-valued in practice and, consequently, relaxed optimization is particularly suitable for maximizing the cost function.

3 Simulations

The simulations involved a signal waveform which was extracted from a respiratory flow signal, acquired from a patient with coronary heart failure which was part of a database investigated in a recent study [10]. One representative breathing cycle of about 2.5 s was extracted from a patient with periodic breathing, corresponding to 25 samples at a sampling rate of 10 Hz. Twenty zero-valued samples were then inserted both before and after the extracted cycle in order to produce the transient waveform $s(n)$, displayed in Fig. 1(a) when being part of a small ensemble.

The integer-valued time delay θ_i was assigned a uniform PDF over the interval $[-\delta, \delta]$, where $\delta = 3$. The performance was evaluated using the Monte Carlo simulation technique with $R = 50$ different realizations for a small ensemble size of $M = 10$; this size was chosen so as to make the computations of the brute force method less painful. The resulting ensemble contained signals with a mixture of signal-to-noise ratios (SNRs), ranging from a variable lower limit (the SNR indicated below in the results) to a fixed upper limit equal to 30 dB; SNR is here defined as $10 \cdot \log(E_s/\sigma_v^2)$.

Both variants of the present eigenvalue-based method embrace only one single parameter, namely the maximum time shift Δ_{max} which defines the length of the shift interval $[-\Delta_{\text{max}}, \Delta_{\text{max}}]$ around the true time delay. This parameter is here set equal to δ, i.e., $\Delta_{\text{max}} = 3$. The default values of Matlab's implementation of the Nelder–Mead algorithm were employed.

In order to evaluate the performance, the present method was compared to the Woody method [6], see also [1]. The Woody alignment involved, just as the eigenvalue-based method, a symmetrical search interval of the same length around the true time delay to determine the time of the maximal output of the matched filter.

Alignment performance is here synonymous to the variance of the time delay errors $(\hat{\theta}_i - \theta_i), i = 1, \ldots, M$, and averaged over the R different Monte Carlo runs. All computations were performed in Matlab on a server (BLUE) with 8 parallel processors.

4 Results

Figure 1(b) and (c) illustrate the alignment that results from using Woody's method and brute force optimization, respectively. In this example, it is visually evident that the present eigenvalue-based method offers better alignment when the ensemble is heterogenous with respect to SNR.

Figure 2 presents the time delay error variance for different lower limits of the mixed SNR. It is noted that the error variance is 0 for the brute force method for an SNR equal to 25 dB. As the SNR drops, the performance of the Woody method deteriorates much faster than do the two variants of the present method. Indeed, they perform essentially the same when the SNR is equal to 15 dB. However, the brute force method requires about 6 weeks to finish whereas the relaxed optimization finishes in a few seconds!

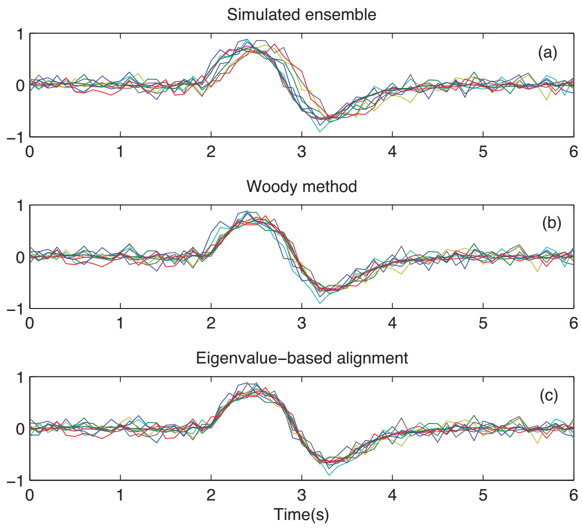

Figure 1: An example of (a) a simulated ensemble containing $M = 10$ noisy respiratory flow waveforms. The ensemble is aligned using either (b) the Woody method or (c) the new eigenvalue-based method (brute force).

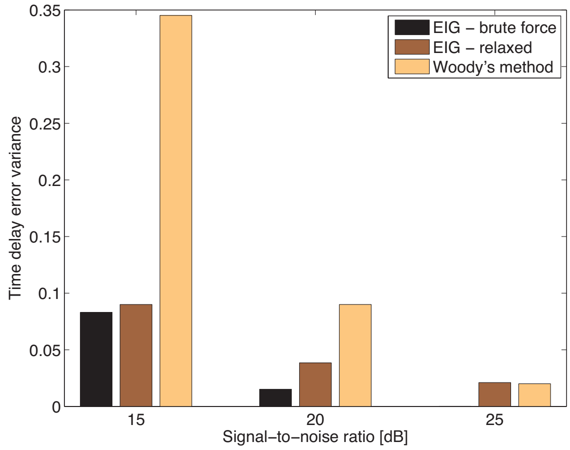

Figure 2: The time delay error variance of the new eigenvalue-based method (EIG), using either brute force or relaxed optimization, and the Woody method. Note that the error variance is equal to zero for brute force optimization when the lower SNR limit is equal to 25 dB.

5 Conclusions

A new method is presented for ensemble time alignment which explores the properties of the eigenvalues of the data matrix. The eigenvalue criterion to be maximized is given the interpretation of a signal-to-jitter-and-noise ratio. While the results are yet preliminary in nature, they suggest nonetheless that the new eigenvalue-based criterion can offer much better performance than does the widely used Woody's method for time alignment of waveforms with fixed morphology. No estimate of the underlying signal waveform is required.

Acknowledgment

This work was supported in part by Ministerio de Ciencia e Innovación from Spanish Government (grant TEC2010-21703-C03) and the Swedish Research Council (grant 2009-4584).

References

[1] L. Sörnmo and P. Laguna. *Bioelectrical Signal Processing in Cardiac and Neurological Applications*. Elsevier (Academic Press), Amsterdam, 2005.

[2] A. Zviagintsev, Y. Perelman, and R. Ginosar. Algorithms and architectures for low power spike detection and alignment. *J. Neural Eng.*, 3:35–42, 2006.

[3] D. T. Pham, J. Möcks, W. Köhler, and T. Gasser. Variable latencies of noisy signals: Estimation and testing in brain potential data. *Biometrika*, 74:525–533, 1987.

[4] L. Xu, P. Stoica, J. Li, S. L. Bressler, X. Shao, and M. Ding. ASEO: A method for the simultaneous estimation of single-trial event-related potentials and ongoing brain activities. *IEEE Trans. Biomed. Eng.*, 56:111–121, 2009.

[5] E. Laciar, R. Jané, and D. H. Brooks. Improved alignment method for noisy high-resolution ECG and Holter records using multiscale cross-correlation. *IEEE Trans. Biomed. Eng.*, 50:344–353, 2003.

[6] C. D. Woody. Characterization of an adaptive filter for the analysis of variable latency neuroelectric signals. *Med. Biol. Eng.*, 5:539–553, 1967.

[7] G. H. Steeger, O. Hermann, and M. Spreng. Some improvements in the measurement of variable latency acoustically evoked potentials in human EEG. *IEEE Trans. Biomed. Eng.*, 30:295–303, 1983.

[8] A. Cabasson and O. Meste. Time delay estimation: A new insight into the Woody's method. *IEEE Signal Proc. Letters*, 15:573–576, 2008.

[9] W. Truccolo, K. H. Knuth, A. Shah, S. L. Bressler, C. E. Schroeder, and M. Ding. Estimation of single-trial multicomponent ERPs: differentially variable component analysis (dVCA). *Biol. Cybern.*, 89:426–438, 2003.

[10] A. Garde, L. Sörnmo, R. Jané, and B. F. Giraldo. Correntropy-based spectral characterization of respiratory patterns in patients with chronic heart failure. *IEEE Trans. Biomed. Eng.*, 57:1964–1972, 2010.

Address for correspondence:

Dr. Leif Sörnmo
Deprtment of Electrical and Information Technology
Lund University, Sweden
leif.sornmo@eit.lth.se