UC Irvine

UC Irvine Previously Published Works

Title

A characterization of a single-trial adaptive filter and its implementation in the frequency domain.

Permalink

<https://escholarship.org/uc/item/96b946g3>

Journal

Electroencephalography and clinical neurophysiology, 73(4)

ISSN

0013-4694

Authors

Arpaia, JP
Isenhardt, R
Sandman, CA

Publication Date

1989-10-01

DOI

10.1016/0013-4694(89)90114-4

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

A characterization of a single-trial adaptive filter and its implementation in the frequency domain

Joseph P. Arpaia, Robert Isenhardt and Curt A. Sandman

State Developmental Research Institute, Fairview, Costa Mesa, CA 92626 (U.S.A.), and University of California Irvine Medical Center, Department of Psychiatry and Human Behavior, Orange, CA 92668 (U.S.A.)

(Accepted for publication: 28 April 1989)

Summary A single-trial adaptive filter (SAF) was implemented in the frequency domain (FDAF) by using the Fast Fourier Transform. The FDAF is significantly more efficient than the SAF. In the data presented the FDAF ran approximately 2 times faster than the SAF. For time series containing larger numbers of data points (n) the efficiency of the calculation will increase on the order of $N/\ln(N)$. The FDAF was tested under a variety of conditions to determine the limits of its usefulness. Pre-filtering the data was found to be necessary to prevent the FDAF from lining up on high frequency activity not related to the signal. The importance of minimizing the amount of low frequency noise was emphasized since it adversely affected the performance of the FDAF and was difficult to filter. The single-trial latencies predicted by the FDAF were much more sensitive to increasing noise than the final wave form. In the absence of excessive low frequency noise a negative exponential relationship was found between the mean error in latency prediction and the SNR estimate. Since the SAF technique is also used to determine signal latency in single sweep data the SNR estimate can be a useful test to determine if the FDAF is locating the signal correctly or merely amplifying chance regularities in noisy data.

Key words: Woody filter; Fast Fourier Transform; Single-trial latencies

One difficulty encountered by investigators studying event-related potentials (ERPs) is the need for accurate extraction of signal from background noise. Although averaging numerous epochs improves the signal-to-noise ratio, middle- and long-latency responses (> 25 msec) are unstable and not strictly time-locked to the stimulus (Sutton and Ruchkin 1984). This variable latency of the signal in single sweeps can greatly and artificially attenuate amplitude of the averaged signal. A single-trial adaptive filter technique (SAF) has been proposed (Woody 1967) to estimate the latency of the signal in each single sweep and enhance the signal-to-noise ratio (Pfefferbaum et al. 1980, 1984; Michalewski et al. 1986).

In the SAF technique each sweep is pre-filtered with a low-pass digital filter and shifted along an initial template. Pre-filtering is necessary to keep the SAF from amplifying high frequency activity unrelated to the signal. The cross-covariance between the sweep and the template is calculated and the temporal shift which gives the maximum value of the cross-covariance is chosen as the latency 'correction' for that sweep. Each sweep is then shifted by its corresponding latency. The average of all the shifted sweeps becomes the new template and the entire process is repeated (Woody 1967). One disadvantage of this procedure is that the number of calculations involved is approximately the square of the number of data points in each time series. Therefore, as either the sampling period or the sampling rate increases, the number of calculations increases.

However, the same procedure can be done in the frequency domain using the Fast Fourier

Correspondence to: Curt A. Sandman, Ph.D., State Developmental Research Institute, Fairview, 2501 Harbor Boulevard, Costa Mesa, CA 92626 (U.S.A.).

Transform (FFT). The pre-filtering can be done in the frequency domain by multiplying the Fourier coefficient for each frequency by a scaling factor between 0 and 1. The latency giving the maximum value of the cross-covariance function also can be found from the FFT of the two time series. It is well known that the complex conjugate product of the Fourier Transforms of two time series is the Fourier Transform of their cross-covariance function (Otnes and Enochson 1972; Elliott and Rao 1982). That is, given the two time series $f(t)$ and $g(t)$ with Fourier Transforms $F[f(t)]$ and $F[g(t)]$ respectively,

$$F^{-1}\{F[f(t)] \times F[g(t)]^*\} = z(\tau)$$

where $z(\tau)$ is the cross-covariance function between the two time series, and $F[g(t)]^*$ is the complex conjugate of $F[g(t)]$. If $f(t)$ is the template and $g(t)$ is the sweep then this efficient formula can be used to find the sweep latency correspond-

ing to the maximum value of the cross-covariance function. The trial is shifted by the latency which gives the maximum value for the cross-covariance function and the shifted trials are averaged to form a new template. The new template is used and another iteration is performed. A flow chart of the frequency domain adaptive filter (FDAF) is presented in Fig. 1. Note that the FDAF uses pre-filtered data only to calculate the cross-covariance function. Since unfiltered data are used to generate the new template none of the data is lost in the filtering procedure.

The number of calculations involved in finding the FFT of a time series with N data points increases as $N \ln(N)$ while the number of calculations needed to find the cross-correlation function increases as N^2 . Therefore, for time series containing large numbers of data points the FDAF will take less time than the SAF (Borgioli 1968).

In the present report the FDAF was tested on synthetic data while parametrically altering the signal-to-noise ratio to define the limits of its usefulness. Pre-filtering was performed by setting all Fourier coefficients above the pre-filter value to zero, while the coefficients corresponding to frequencies at or below the pre-filter value were left unchanged. Finally the FDAF was used on single-trial EEG data.

Fig. 1. Flow chart of the FDAF.

Methods and results

Synthetic sweeps were modeled to resemble the EEG by adding together 3 sine waves, gaussian random noise, and a signal. Each sine wave was given a frequency in one of the following bands: 0.5–1.5 Hz, 8.0–10.0 Hz and 13.0–17.0 Hz. A random frequency value in each band was picked and a sine wave of that frequency and a random phase was added to each sweep. The gaussian random noise was generated by averaging 8 numbers from a pseudo-random number generator for each noise point. The signal (Fig. 2) was given by the equation

$$s(t) = 60 \exp(-3.5(t - 0.4)) \times \sin(2(1.5 + 8.2(t - 0.4))(t - 0.4))$$

on the interval $0 \leq t \leq 0.4$ (Coppola et al. 1978).

Fig. 2. Signal used in the synthetic sweeps. Signal is a swept frequency from 4.8 to 1.5 Hz. The amplitude of the signal was held constant across sweeps and across all tests of the FDAF.

The midpoint of each sweep was defined as $t = 0$ and the signal was allowed to appear within ± 400 msec of the 0 point. The equation for the synthetic sweep is given by:

$$d(t) = A_0 n(t) + A_1 \sin(w_1 t + p_1) \\ + A_2 \sin(w_2 t + p_2) + A_3 \sin(w_3 t + p_3) \\ + A_4 s(t + p_4)$$

where $n(t)$ and $s(t)$ are the noise and the signal respectively, A_0 – A_4 are the amplitudes for each term,

$$0 \leq p_1, p_2, p_3 \leq 2 \quad \text{and} \quad -0.4 \leq p_4 \leq 0.4$$

are the phase shifts, and

$$0.5 \leq w_1 \leq 1.5, 8.0 \leq w_2 \leq 10.0, 13.0 \leq w_3 \leq 17.0$$

are the frequencies of the 3 sine waves, respectively. One hundred synthetic sweeps were used in each test of the filter. Signal jitter was modeled by randomly assigning a different latency value to the signal component of each sweep. The variance ratio (VR) over the 100 sweeps was defined as the variance of the signal divided by the mean of the variance of each sweep. This is equivalent to the mean signal-to-noise ratio over the period of the sweep. The signal amplitude, A_4 , was held constant for all tests of the FDAF and the other amplitude coefficients A_0 – A_3 were varied to test the behavior of the FDAF at different variance ratios.

The FDAF was first checked against the SAF to compare their relative speeds. Each trial consisted of 128 data points (2 sec at 64 Hz resolu-

tion) and the single trials were allowed to shift over ± 400 msec giving 50 points at which to calculate the cross-covariance function. The FDAF ran approximately twice as fast as the SAF. If twice the number of data points had been used (data points at 128 Hz resolution) the FDAF would have run more than 3 times as fast as the SAF.

The FDAF was next implemented on the synthetic data to determine which pre-filter value would give the best results. Two criteria were used to check the operation of the FDAF. The first was the shape of the wave form generated by the FDAF compared to the shape of the signal. The second was the difference between the signal latency for each trial. A summary statistic for the second criterion was defined as

$$\delta = \left[\left(\sum_{i=1}^{100} (fl_i - sl_i)^2 \right) / 100 \right]^{1/2}$$

where fl_i is the latency predicted by the FDAF for trial i and sl_i is the actual latency for trial i . The FDAF was tested separately on 3 different types of noise: 8–10 Hz activity, 0.5–1.5 Hz activity, and gaussian random noise.

First the FDAF was used on synthetic data generated with increasing activity in the middle frequency band. The wave form after 5 iterations was examined for different pre-filter values. A pre-filter value corresponding to a low pass filter with a cut-off frequency of 8 Hz or less is needed to eliminate the 8–10 Hz sine wave activity from the final wave form (Fig. 3). A pre-filter of 6 Hz gave the smallest value of δ in this case, but a pre-filter of 7 Hz gave a smaller value of δ when large amounts of gaussian noise were present (data not shown). Therefore a pre-filter value of 7 Hz was used in all subsequent tests of the FDAF.

The FDAF was next tested on synthetic data generated with increasing amplitude of gaussian noise. Since the amplitude of the signal was constant, the variance ratio (VR) decreased as the amplitude of the noise increased. The wave form and δ are presented for the initial template and after each iteration of the FDAF in Fig. 4. The variance ratios for the graphs are 0.19, 0.12, 0.08, 0.07, 0.06, 0.054, 0.048 for graphs 4a–4g respec-

Fig. 3. Wave forms generated after 5 iterations of the FDAF operating on synthetic data with high alpha band activity. The same synthetic data was used in each case, only the pre-filter value (see Fig. 1) was changed.

tively. The wave form labeled as iteration 0 in each set of graphs is the time-locked average of the sweeps and was used as the initial template for the FDAF. Surprisingly the shape of the final wave form was remarkably stable, with only the amplitude of the N2-P3 difference being artificially increased at lower VR, and the shape of N1 being obscured at the lowest VR of 0.048. The value of δ was much more sensitive to the increasing amounts of noise, climbing rapidly as the VR dropped below 0.08. This indicates that the ability of the FDAF to determine the signal latency in individual sweeps decreased as the amount of noise increased.

The FDAF finally was tested on synthetic data with increasing amounts of low frequency noise

added to the sweeps. As above, the wave form and δ are shown for the initial template and after each iteration of the FDAF (Fig. 5). Here the variance ratios are 0.19, 0.18, 0.16, 0.15 and 0.12 (Fig. 5a-e respectively). Again δ was much more sensitive

Fig. 4. a-g: wave forms generated by the FDAF. Each set of wave forms is from a different set of synthetic data containing different amounts of gaussian noise. Part a contains the least amount of noise and part g contains the greatest amount of noise.

Fig. 5. a-e: wave forms generated by the FDAF. Each set of wave forms is from a different set of synthetic data containing different amounts of low frequency noise. Part a contains the least amount of noise and part e contains the greatest amount of noise.

than the shape of the wave form to increasing low frequency noise and increased rapidly as the amount of power in the low frequency noise increased above the amount of power in the signal.

The complete set of tests described above was repeated using the positive half of a 2 Hz sine wave (Pfefferbaum et al. 1984) centered at 390 msec after the 0 points as the initial template. All values of δ as well as all the resultant wave forms were virtually the same as when the time-locked average was used as the initial template.

Since the VR is not available for real EEG data it is important to have an estimate that indicates if the Woody filter is finding a signal that is there, or is merely amplifying noise. A signal-to-noise estimate (SNR) was attended using the formula

$$s = \exp(2.66 - 1.56 * \exp(-1.16 * z + 1.56))$$

where z is the mean z -transform of the cross-correlation coefficient between each sweep and the template taken over the 400 msec signal period (Coppola et al. 1978). The 400 msec signal period was the data between 0 and 400 msec (which would completely contain an unshifted signal). A semi-log plot of $\ln(\delta)$ vs. the SNR after 5 iterations is shown in Fig. 6. The data can be divided into two sets based on the power present in the low frequency band. The triangular points correspond to data sets in which the mean power in the low frequency band was greater than 0.9 times the power in the signal. Two separate regressions were performed on the two data sets. In the absence of excessive low frequency power the relationship was

$$\delta(\text{msec}) = 399 e^{-3.4(\text{SNR})}$$

and if the low frequency power was greater than 0.9 times the signal power the relationship was

$$\delta(\text{msec}) = 659 e^{-1.8(\text{SNR})}$$

The sensitivity of the FDAF to low frequency noise occurred because 90% of the signal power appeared between 0 and 6 Hz with a peak between 4 and 5 Hz. It is not possible to effectively pre-filter the low frequency noise without overly attenuating the signal. It is also difficult to test for this condition since any test involves a priori assumptions about what is signal and what is noise. However, since the ERP paradigms are searching for a signal with similar spectral characteristics to the test signal described above, one test is to examine the power spectrum of the final wave form in the 0-2 Hz bandwidth. A plot of δ vs. the fraction of power in the 0-2 Hz band, F , is shown in Fig. 7. The steep negative relationship between δ and F emphasizes the need to control low frequency drift from the instrumentation used in collecting the EEG.

Single-trial EEG was collected from C3, C4, Fz and Pz using a Grass polygraph, Model 79, equipped with 7P511 amplifiers. Subjects were seated with eyes closed and pure tone stimuli of 450 and 550 Hz at 92 dB SPL were presented at 2.56 sec intervals. The 550 Hz tone was designated as the target stimulus and subjects depressed a key each time they detected the target. The EEG responses to the target stimuli were collected by

Fig. 6. Plot of $\ln(\delta)$ vs. SNR. Note that increased amounts of low frequency noise, the triangular points, cause a much greater increase in δ than increased amounts of gaussian random noise.

Fig. 7. Plot of δ vs. power ratio. The x-axis represents the percentage of power in the sweeps that were present in the 0–2 Hz frequency band. The steep negative relationship shows the sensitivity of the FDAF to noise at frequencies that overlap the frequency band of the signal.

Fig. 8. FDAF used on real EEG data. Part a from subject 1 shows a modest increase in N2-P3 amplitude after 5 iterations. Part b from subject 2 has a signal-to-noise ratio that is too low for the FDAF to work reliably. Part c from subject 3 illustrates the performance of the FDAF with a pre-filter of 10 Hz and is an example of how the FDAF may amplify extraneous high frequency noise in the single sweeps if the data are not pre-filtered appropriately. Part d shows the same data from subject 3 using a pre-filter of 8 Hz, illustrating how the pre-filter allowed the FDAF to amplify the signal without amplifying the extraneous high frequency noise.

sampling at 200 Hz for 1280 msec pre stimulus to 1280 msec post stimulus. A low pass filter was implemented by adding together 4 sequential samples to produce each data point for analysis, giving a final data resolution of 50 Hz. Forty-four sweeps were obtained for each subject and the FDAF was tested on these data from the Pz electrode. The FDAF showed improved SNR up to the third iteration and then no subsequent improvement. The useful parameters derived from the synthetic data were tested in 4 subjects as displayed in Fig. 8. Subject no. 1 (Fig. 8a) shows an increased P300 after 3 iterations of the filter. Subject no. 3 (Fig. 8c, d) shows an increased P300 with a pre-filter of 8 Hz and 10 Hz, but the pre-filter of 10 Hz (Fig. 8d) allowed the FDAF to line up on the subject's prominent alpha activity and thus obscured the P300 wave form. In both

subjects 1 and 3 the estimated SNR increased approximately 40% after 3 iterations of the FDAF. This required 1.5 min/subject. In subject 2 (Fig. 8b) the SNR increased after 3 iterations; however, there is no identifiable P300 in the final wave form. This is in agreement with the synthetic trials where the final SNR estimate was indicative of the reliability of the FDAF. Because the final SNR estimate was 0.08, the FDAF would not be expected to succeed, and the observed wave form was due to the FDAF lining up chance regularities in a noisy signal.

This study was supported in part by a grant from the National Institute on Aging (No. AGO3975-04) to C.A.S.

The authors would like to thank Dr. Henry J. Michalewski for his helpful comments and discussion.

References

- Borgioli, R.C. Fast Fourier Transform correlation versus direct discrete time correlation. *Proc. IEEE*, 1968, Sept.: 1602-1604.
- Coppola, R., Tabor, R. and Buchsbaum, M.S. Signal to noise ratio and response variability measurements in single trial evoked potentials. *Electroenceph. clin. Neurophysiol.*, 1978, 44: 214-222.
- Elliott, D.F. and Rao, K.R. *Fast Transforms: Algorithms, Analysis and Applications*. Academic Press, New York, 1982.
- Michalewski, H.J., Prasher, D.K. and Starr, A. Latency variability and temporal interrelationships of the auditory event-related potential (N1, P2, N2, and P3) in normal subjects. *Electroenceph. clin. Neurophysiol.*, 1986, 65: 59-71.
- Otnes, R.K. and Enochson, L. *Digital Time Series Analysis*. Wiley, New York, 1972.
- Pfefferbaum, A., Ford, J.M., Roth, W.T. and Kopell, B.S. Age-related changes in auditory event-related potentials. *Electroenceph. clin. Neurophysiol.*, 1980, 49: 266-276.
- Pfefferbaum, A., Ford, J., Johnson, R., Wenegrat, B.G. and Kopell, B.S. Manipulation of P3 latency: speed vs. accuracy instructions. *Electroenceph. clin. Neurophysiol.*, 1983, 55: 188-197.
- Pfefferbaum, A., Ford, J.M., Wenegrat, B.G., Roth, W.T. and Kopell, B.S. Clinical applications of the P3 component of event-related potentials. I. Normal aging. *Electroenceph. clin. Neurophysiol.*, 1984, 59: 85-103.
- Sutton, S. and Ruchkin, D.S. The late potential complex. *Ann. NY Acad. Sci.*, 1984, 425: 391-397.
- Woody, C.D. Characterization of an adaptive filter for the analysis of variable latency neuroelectric signals. *Med. Biol. Eng.*, 1967, BME-5: 539-553.