FORTRAN subroutine for an EEG feedback digital filter

PHILLIP L. EMERSON

Cleveland State University, Cleveland, Ohio

A method of constructing digital filters is outlined and illustrated in an application for research uses of EEG feedback. The method is based on least squares estimation with geometric weighting. Old data in the time series are discounted according to age, as new data arrive. Recurrence relations are used, allowing the computations to be performed in real time. A FORTRAN digital filter is formulated that simultaneously estimates (1) the spectral power in a band centered at a target frequency, (2) the total spectral power, and (3) the dominant frequency.

Pope and Gersten (1977) pointed out some advantages for using a general-purpose digital computer in biofeedback research. Many chores, such as record keeping, statistical analysis, and equipment calibration, can be formulated in a coordinated set of programs. Another advantage is that transcription of data, between collection and analysis, can be avoided. The process of extracting the feedback information from the signals produced by the subject is a statistical estimation process. The estimates must be updated continually for adaptation to changes in the parameters being estimated. This estimation process, the heart of the biofeedback method, can be extremely simple. The biofeedback task considered here requires that the subject learn to produce EEG waves at or near a given target frequency. The task allows for a shaping technique to be used, according to the general principle of discrimination training.¹ The subject is given an easy goal to attain in early training, and the acceptable range of responses is gradually narrowed during training as the subject's skill increases.

EQUIPMENT REQUIREMENTS

A general-purpose laboratory computer system is needed with FORTRAN capabilities, an analog-to-digital converter, and a programmable real-time clock. The clock should have at least millisecond resolution and should be programmable to produce time intervals between 10 and 500 msec. An EEG amplifier should connect into the A/D converter. The EEG amplifier must have selectable band limiting to reject frequency activity above a selected upper limit and below a selected lower limit. The upper limit

should be set at approximately 2 to 3 times the target frequency. This band limiting is the one function that cannot be done effectively with digital filtering (Blackman & Tukey, 1958). The setting for the upper limit is suggested because the procedure to be outlined assumes the Nyquist frequency to be twice the target frequency. The lower limit can be set to about 1/10 the target frequency, because zero crossings are used for frequency estimation, and the zero-frequency component therefore must be removed.

The method outlined here is offered for research rather than clinical or hobbyist applications. [For example, Regan (1975) used evoked EEG feedback in a closed-loop system of tracking visual thresholds adaptively with changing experimental conditions.]

THEORY

Other than for the band limiting mentioned above, filtering can be done either with hardware analog circuits or by programmed digital (arithmetical) methods. There is a large body of engineering literature on digital filtering (Liu, 1975). Most of that literature requires familiarity with Fourier, Laplace, and Z transforms, as well as contour integration in the complex plane. However, one useful kind of digital filter can be constructed merely on the basis of least squares theory and the following three principles:

$$(1-r)\sum r^n = 1 \quad |r| < 1, \quad (1)$$

where r may be real or complex, and the summation extends from $n=0$ to ∞ , for integer n ; and

$$\cos\theta = [\exp(i\theta) + \exp(-i\theta)]/2, \quad (2)$$

$$\sin\theta = [\exp(i\theta) - \exp(-i\theta)]/(2i), \quad (3)$$

where $i = \sqrt{-1}$ and θ is real. Using these three relations,² one can evaluate the quantities P , Q , and R —defined

Part of this paper was written while the author was on professional leave from Cleveland State University during 1984. The ideas were developed from research begun while the author was supported by a Research Initiation Award from Cleveland State University during 1975. The author is indebted to anonymous reviewers and the editor for advice on earlier versions of this paper. The author's mailing address is: Department of Psychology, Cleveland State University, Cleveland, OH 44115.

as $P = \sum r^n \cos^2(n\theta)$, $Q = \sum r^n \sin^2(n\theta)$, $R = \sum r^n \cos(n\theta) \sin(n\theta)$ —which arise in the least squares problem to be formulated below. There are general closed-form expressions for these sums (Jolley, 1961), but they simplify under the special condition to be imposed below, $\theta = \pi/2$.

Let $[\dots, y_2, y_1, y_0]$ be a time series of observed data at equally spaced discrete time points, with the sense of the time index reversed. Thus, y_0 is the most recent observation, y_1 is the preceding one, and so forth. (This reversed time indexing is used here throughout this paper, as are the summation limits, $n = 0$ to ∞ .) Let $\hat{y}_n$ be an estimate of y_n , based on the pair of Fourier terms for a single target frequency, that is,

$$\hat{y}_n = a_0 \cos(n\theta) + b_0 \sin(n\theta). \quad (4)$$

Here, a_0 and b_0 are the two Fourier coefficients to be estimated immediately after y_0 has been observed. The estimates of a_0 and b_0 are to be determined so as to minimize W , where

$$W = \sum r^n (y_n - \hat{y}_n)^2. \quad (5)$$

The parameter r ($0 < r < 1$) is a geometric decay constant. This geometric weighting discounts old data according to age. The sum of the squares of the estimates of a_0 and b_0 provides an estimate of how much recent activity has been in the vicinity of the target frequency.³ This sum of squares is referred to as the in-band power estimate.

The estimates of a_0 and b_0 can be obtained recursively, which makes this procedure feasible for real-time calculations. Let $U_0 = \sum r^n y_n \cos(n\theta)$, and $V_0 = \sum r^n y_n \sin(n\theta)$. Setting the partial derivatives of the weighted sum of squared residuals (W) equal to zero yields the normal equations, $U_0 = a_0 P + b_0 R$ and $V_0 = a_0 R + b_0 Q$. P , Q , and R are fixed. Their values do not change from one time step to the next. However, U_0 and V_0 must be revised with each new datum. Because of the identities,

$$\cos[(n+1)\theta] = \cos(\theta)\cos(n\theta) - \sin(\theta)\sin(n\theta) \quad (6)$$

and

$$\sin[(n+1)\theta] = \cos(\theta)\sin(n\theta) + \sin(\theta)\cos(n\theta), \quad (7)$$

the following recurrence relation holds:

$$U_0 = y_0 + U_1 \cos\theta - V_1 \sin\theta, \quad (8)$$

$$V_0 = V_1 \cos\theta + U_1 \sin\theta. \quad (9)$$

One can proceed to find that these relations yield recurrence relations, generally for the Fourier coefficients to be estimated. However, it is convenient at this point to specialize, taking $\theta = \pi/2$, which is the case with the particular filter formulated here. This yields the welcome consequences that $P = 1/(1-r^2)$, $Q = r/(1-r^2)$, and $R = 0$. Also, under this condition, $U_0 = y_0 - rV_1$, and $V_0 = rU_1$. Only a few simple manipulations are needed then to show that

$$a_0 = y_0 - r^2(y_0 + a_2) \quad (10)$$

and

$$b_0 = a_1 \quad (11)$$

for the case of $\theta = \pi/2$.

IMPLEMENTATION

Figure 1 contains a subroutine that serves as a digital filter. This can be used for EEG biofeedback tasks in which the subject is to learn to produce brain waves at a given target frequency. This filter simultaneously estimates (1) the spectral power in a bell-shaped bandpass window centered at the target frequency, (2) the total spectral power being produced by the subject, and (3) the actual dominant frequency of waves being produced by

```

SUBROUTINE FILTER(DECAY, FTARG, DT, YO, POW, FEST, TOT)
DATA A1, A2, Y1/O., O., O./
TOT = (1.-DECAY)*YO**2 + DECAY*TOT
AO = YO - (YO+A2)*DECAY**2
POW = (AO**2+A1**2)/2.
G = 0
IF (YO*Y1) 100, 200, 200
100 G = 1./(DT+DT)
200 FEST = G + (FEST-G)*DECAY
Y1 = YO
A2 = A1
A1 = AO
RETURN
END

```

Figure 1. Digital filter for estimating the in-band spectral power (POW) at a target frequency (FTARG), the total spectral power (TOT), and the dominant frequency (FEST) of EEG activity produced by the subject in a biofeedback session. DECAY and DT are parameters determined as explained in the text, and YO is the most recent EEG voltage, as determined by an A/D conversion. These conversions, and the calls to FILTER, are assumed to occur at intervals of DT seconds.

the subject. These estimates are updated continually at each new step in the process of digitizing the EEG input at equally spaced time points. The simplest of biofeedback devices usually provides some form of estimate of the in-band power, as mentioned in (1) above. Some of these normalize with respect to a total power estimate as mentioned in (2), and some provide an additional display based on a frequency estimate, as mentioned in (3).

Thus, there is no fundamentally new estimated quantity with this digital filter. The new aspect is that a single bandwidth parameter, BW, adapts all estimates to different circumstances, such as in the characteristic pattern of discrimination training in operant conditioning paradigms. At the beginning of the training process, the value of BW can be set for a broad bandpass window, which has several effects. The first effect is that the production of any frequency within a fairly broad range will be reinforced. As training progresses, the parameter can be used to narrow the bandpass window to encourage the subject to produce waves more concentrated near the target frequency. The effect in itself usually is achievable with commonly available biofeedback devices. However, there are some secondary effects of the value of BW that give rise to a trade-off between frequency selectivity and settling time. With a broad bandpass window, the filter reacts rapidly to any change, providing almost immediate feedback. As the band is narrowed gradually, the information that is fed back becomes more precise. The time delay (settling time) necessarily increases because of the requirement of effectively more data for greater precision. This increasing precision applies to the frequency estimate as well as the power estimates. For the frequency estimate, a running estimate of the rate of zero crossings is maintained, step by step, by a geometrically weighted moving average.

The parameter BW controls the adaptiveness of all three estimates via geometric weighting. The sine and cosine coefficients for the target frequency are weighted least squares estimates with this geometric weighting. The sum of their squares determines the in-band power estimate. The estimate of total power is a geometrically weighted average of the squares of the raw time series observations. The geometric weighting in all three cases places most weight on the most recent quantities, where the weighting of old quantities decays, according to the terms of a geometric sequence, with the age of the data. Geometric (or exponential) decay is a very common way of discounting old data in an adaptive system as new data arrive. It has the added advantage, in the present application, that all estimates can be calculated recursively. Thus, no large data arrays need to be maintained. The recursiveness of the calculations with the geometric weighting makes this one of the few digital methods that is feasible for real-time adaptive estimation.

These design principles do not entirely determine the filter. Two choices were therefore made on the grounds of practicality and simplicity. The target frequency is represented in the subroutine of Figure 1 as FTARG, the time spacing between A/D conversions as DT, and the

geometric decay constant as DECAY. For the first choice, DT is determined as $1/(4\text{FTARG})$, in seconds. This is equivalent to the choice of $\theta = \pi/2$ in the theoretical section. This choice makes each time shift equal to a quarter of a cycle at the frequency FTARG. This choice results in extremely simple recurrence relations for the least square estimates of the Fourier coefficients. However, a consequence is that a fully programmable real-time clock is needed, to produce the time interval, DT, with approximately millisecond resolution. Secondly, for convenience, DECAY was determined as $\exp(-\text{BW}/\text{FTARG})$. Thus, BW and FTARG are both in units of cycles per second, and BW is a free parameter that can be used to control the estimation process in the way mentioned earlier. With these choices decided, the decay of the effects of old data follows the course of $\exp(-4t\text{BW})$, where t is in seconds. Therefore, settling time can be figured as $t = -\ln(p)/(4\text{BW})$, where p is the relative weighting of data that are t seconds old. For example, with $p = .5$ (for "half-life"), $t = .17/\text{BW}$. For $p = .05$, $t = .75/\text{BW}$, which is more like settling to "zero." The convenient formula, $t = 1/(2\text{BW})$, corresponds approximately to $p = .135$, which seems to be a reasonable criterion for a conventional designation of "settling time."

The arguments YO, POW, FEST, and TOT for the FILTER subroutine of Figure 1 are as follows. The first, YO, is the value of the most recently observed time series datum. The values of YO are acquired at equally spaced time intervals of DT seconds. POW is the returned estimate of the in-band power. FEST is the returned estimate of the dominant frequency of the activity being monitored, and TOT is the returned estimate of the total spectral power. Thus, the subroutine should be called after each new datum becomes available.

It is difficult to construct an example calling program that can be used generally in laboratory research. There is too much variability in the calling sequences for real-time clocks and A/D converters from one system to another. However, a main program is shown in Figure 2, together with a function, SUBJ, that simulates the subject, the real-time clock, and the A/D conversions. This function generates a cosine wave of frequency FTRUE, and adds a random component to yield YO, the simulated result of the A/D conversion. The parameter SNRAT is a specified signal-to-noise ratio. This main program can be used for test purposes to obtain rough estimates of how the filter behaves under various circumstances. The function SUBJ provides a rough simulation of real data. For one thing, the noise is not band limited to reject frequencies above 2 to 3 times FTARG, as it should be. These programs are in PDP-11 RT-11 FORTRAN. The main peculiarities of this version of FORTRAN are the input and output which are from and to the console (Devices 5 and 7), and the random number generator, RAN(K1,K2). This is a uniform pseudorandom number generator producing numbers between 0 and 1.

Rough as the simulation program is, it can be used to verify some of the following features of the filter. For

```

C EXAMPLE MAIN PROGRAM
      DATA T/O./
      WRITE(7,100)
100    FORMAT(1X,'TYPE FTRUE,SNRAT,FTARG,BW')
      READ(5,200)FTRUE,SNRAT,FTARG,BW
200    FORMAT(4F7.0)
      DT = 1./(4.*FTARG)
      DECAY = EXP(-BW/FTARG)
250    YO = SUBJ(FTRUE,SNRAT,DT)
      CALL FILTER(DECAY,FTARG,DT,YO,POW,FEST,TOT)
      POWT = POW/TOT
      WRITE(7,300)T,YO,POW,POWT,FEST
300    FORMAT(1X,5F9.4)
      T = T + DT
      GO TO 250
      END
      FUNCTION SUBJ(FTRUE,SNRAT,DT)
      DATA T/O./
      SUBJ = SNRAT*COS(2.*FTRUE*T*3.1415927)
      SUBJ = SUBJ + SQRT(12.)*(RAN(K1,K2)-.5)
      T = T + DT
      RETURN
      END

```

Figure 2. Test program for the digital filter of Figure 1. The SUBJ function produces simulated EEG data (YO) with a true sinusoidal frequency of FTRUE, and a random component whose average size depends on the parameter SNRAT. The main program gets parameter values from the console, and then proceeds with time steps, DT, to obtain data (YO) and incorporate them stepwise into the current estimates, by calling FILTER. Results at each step are printed on the console display.

BW less than about $FTARG/4$, the power bandpass window is essentially symmetrical about $FTARG$, and is bell shaped. The power is attenuated to about 30% of peak at $FTARG \pm BW$, and to about 10% of peak at $FTARG \pm 2BW$. The in-band power estimates (POW) undulate or "beat" at frequencies determined by the difference between FTRUE and $FTARG$. Also, they vary randomly from step to step, to an extent related positively to BW and negatively to SNRAT. For BW greater than about $FTARG/4$, the power bandpass window begins to lose its symmetry. For $BW > FTARG/2$, the window peaks below $FTARG$ and has a long skewed tail in the positive direction. This is undesirable for estimation purposes, but it does not entirely rule out the use of the filter in such cases. Also, for large values of BW, the in-band power estimates sometimes exceed the estimates of total power. The frequency estimates, FEST, show no regular beat-frequency undulations, but they are quite variable with large BW and small SNRAT. They appear to be biased upward for large BW and small SNRAT, but this may be partly due to the artifact of high frequencies that are not properly rejected with the noise component of the SUBJ function. With decreasing BW, the power and frequency estimates decrease in variability, and the settling time increases. Also, the in-band power estimate for the target frequency becomes more orthogonal to adjacent frequencies,⁴ and the relative power estimates become more accurate. Figure 3 illustrates several of the characteristics of the POW estimates. These curves were generated with a different program, with pure sinusoidal inputs.

One of the advantages of using a computer and a digital filter, rather than a hardware analog filter, is that it is easy to change the design of the filter for improved per-

formance or for adaptation to different circumstances. The filter described here seems adequate for this purpose, and it is extremely simple. At the cost of some simplicity, it can be modified without adding parameters by changing

Figure 3. Power-gain windows for several target frequencies, for $BW = 3\text{Hz}$. The graphs show the near symmetry and near invariance of the window shapes for fixed BW, under variation of the target frequency, except when the target frequency is below about $4BW$. Arrows above the curves are at the target frequencies. Note the severe bias and gain overestimation when $FTARG = 2$, with $BW = 3$. These curves were calculated with pure sinusoidal inputs at frequencies varied above and below the target frequency. They are averages and do not show the beat-frequency oscillations whose extrema would form narrow envelopes enclosing these curves. The band limiting mentioned in the text would curtail the long skewed tails somewhat, for the cases of $FTARG = 2$ and 6 .

the choice of DT from $DT = 1/(4FTARG)$ to, for instance, $DT = 1/(8FTARG)$. The FILTER subroutine would become somewhat more complicated with new operations in the calculations of the least square estimates of the Fourier coefficients, though most of these calculations can still be implemented recursively. The main purpose of such a modification would be to raise the prescribed upper band limit to perhaps 4 to 6 times FTARG, rather than the factor of 2 to 3 prescribed here. This change would increase accordingly the possible range of frequency estimates.

The main disadvantage of digital filtering, when it is done in real time, is the time required for the stepwise computations. If the filter becomes very complicated, and a user is attempting to target on a high frequency, the computations may not be performable rapidly enough. When adaptive real-time estimation is not demanded, the fast Fourier transform is about the most efficient method of treating complete records of finite length. The present FORTRAN filter keeps up in real time with $FTARG = 10$, on a DEC LSI-11 with an extended arithmetic processor. A BASIC version was tried, and it failed to keep up. The computations could be speeded up somewhat by program optimization techniques. Also, if any of the three estimates is not needed for a given application, the corresponding computations can be deleted from the subroutine of Figure 1.

REFERENCES

- ATKINSON, R. C., BOWER, G. H., & CROTHERS, E. J. (1965). *An introduction to mathematical learning theory*. New York: Wiley.
- BLACKMAN, R. B., & TUKEY, J. W. (1958). *The measurement of power spectra*. New York: Dover.
- JOLLEY, L. B. W. (1961). *Summation of series*. New York: Dover.
- LIU, B. (Ed.). (1975). *Digital filters and the fast Fourier transform*. New York: Halstead.
- POPE, A. T., & GERSTEN, C. D. (1977). Computer automation of bio-feedback training. *Behavior Research Methods & Instrumentation*, 9, 164-168.
- REGAN, D. (1975). Colour coding of pattern responses in man investigated by evoked potential feedback and direct plot techniques. *Vision Research*, 15, 175-183.

NOTES

1. With nonhuman subjects, some reward contingency would be needed. With human subjects who are cooperative and attentive, the information feedback is usually sufficient.

2. For the first of these, and several of its ramifications, see Atkinson, Bower, and Crothers (1965).

3. Also, the ratio a_0/b_0 yields information about the current phase of the incoming sinusoidal signal. It is interesting that gradual stepwise phase changes correspond to frequency differences. This general principle is used in some electronic technologies to calibrate frequencies by a scheme known as "phase locked loop." I have attempted to use such a principle in digital filtering for frequency estimation, but the zero crossing rate seemed to provide a more stable and robust frequency estimate.

4. For the sake of simplicity, I dealt only with the case of estimating the Fourier coefficients for a single frequency. It is a straightforward matter to extend the treatment for the simultaneous estimation of all equally spaced frequencies from 0 up to the Nyquist limit, which is $1/(2\Delta t)$, where Δt is the time spacing. Adjacent frequency estimates are almost orthogonal under the condition that $r > \exp(-\Delta f/f_n)$, where Δf is the frequency spacing and f_n is the Nyquist frequency. Practically, if more than just a few of the frequencies are to be treated simultaneously, real-time computations become unfeasible. Off-line, it is more efficient to use the fast Fourier transform on finite fixed-length segments of the time series, or on the complete series as a single segment. Thus, this type of digital filtering is advantageous in a practical sense only when there is a demand for adaptive estimation in real time, and when not very many components need to be estimated simultaneously.

(Manuscript received May 1, 1985;
revision accepted for publication August 20, 1985.)