Documenta Ophthalmologica **99:** 69–82, 1999.

Do’s and don’ts in Fourier analysis of steady-state potentials

MICHAEL BACH and THOMAS MEIGEN
Elektrophysiologisches Labor, Universitäts Augenklinik, Freiburg, Germany

Accepted 19 January 2000

Abstract. Fourier analysis is a powerful tool in signal analysis that can be very fruitfully applied to steady-state evoked potentials (flicker ERG, pattern ERG, VEP, etc.). However, there are some inherent assumptions in the underlying discrete Fourier transform (DFT) that are not necessarily fulfilled in typical electrophysiological recording and analysis conditions. Furthermore, engineering software-packages may be ill-suited and/or may not fully exploit the information of steady-state recordings. Specifically:

- In the case of steady-state stimulation we know more about the stimulus than in standard textbook situations (exact frequency, phase stability), so ‘windowing’ and calculation of the ‘periodogram’ are not necessary.
- It is mandatory to choose an integer relationship between sampling rate and frame rate when employing a raster-based CRT stimulator.
- The analysis interval must comprise an exact integer number (e.g., 10) of stimulus periods.
- The choice of the number of stimulus periods per analysis interval needs a wise compromise: A high number increases the frequency resolution, but makes artifact removal difficult; a low number ‘spills’ noise into the response frequency.
- There is no need to feel tied to a power-of-two number of data points as required by standard FFT, ‘resampling’ is an easy and efficient alternative.
- Proper estimates of noise-corrected Fourier magnitude and statistical significance can be calculated that take into account the non-linear superposition of signal and noise.

These aspects are developed in an intuitive approach with examples using both simulations and recordings. Proper use of Fourier analysis of our electrophysiological records will reduce recording time and/or increase the reliability of physiologic or pathologic interpretations.

Key words: evoked responses, ERG, Fourier analysis, pattern ERG, steady-state, VEP

Purpose

Fourier analysis is a powerful tool in signal analysis that can be fruitfully applied to steady-state evoked potentials (flicker ERG, pattern ERG, VEP, etc.). However, occasionally this is done ineffectively or even inadequately. A number of inherent theoretical assumptions (e.g., the cyclic nature of the analysis interval) in the discrete Fourier transform (DFT) are not necessarily fulfilled in electrophysiological data, which may lead to erroneous results. Furthermore, software packages developed for general engineering are often ill-suited. They may automatically apply undesirable operations (e.g., windowing) and may not fully exploit the extra information that we possess in the case of electrophysiological steady-state recordings.

The analytical power of Fourier analysis is not fully used in electrophysiological software packages. The ISCEV standard for recording and analyzing the full-field flash-ERG does not mention the Fourier transform as a tool to isolate the 30-Hz flicker response [1]. The pattern-ERG guidelines [2] do offer some hints concerning the use of Fourier analysis in the interpretation of steady-state responses. However, the possible artifacts and problems that may arise when these hints are not taken into account could not be illustrated in that context. The objective of the present tutorial is to point out how Fourier analysis may help to analyze and interpret electrophysiological steady-state responses. We will show in which way our situation differs from the general case and how this can be exploited. Furthermore, we will show a number of pitfalls that lurk for naive users of standard Fourier packages. We attempt to develop these ideas in an intuitive way without going into any mathematical detail. We will present graphic examples from simple simulations as well from actual electrophysiologic recordings. While Fourier analysis is also useful for transient records (e.g. for phase-free filtering), we will concentrate here on the steady-state situation. This applies to 8 rev/s and faster PERG and VEP stimulation, and also to flicker ERGs.

‘Fourier analysis’ – what does it mean?

While we want to avoid mathematics here, we need to lay the ground with a few concepts. Our experiments typically are of the following nature: We apply stimuli that can be described by their frequency, and we record electrophysiological responses that are microvolt samples every millisecond or so, stringed together to form the ‘analysis interval’. Thus our analysis interval is a discrete series of numbers over time, or the response in the ‘time domain’. The Fourier theorem states that any function of time over a finite interval can also be expressed in the ‘frequency domain’, that is as a sum of sinusoidal functions with the correct amplitude and phase. The information in the time domain and in the frequency domain is identical, so why should we transform from one representation to the other? Because some information is more accessible in the frequency domain, especially the response to steady-state stimulation. This is explained in the legend of Figure 1.

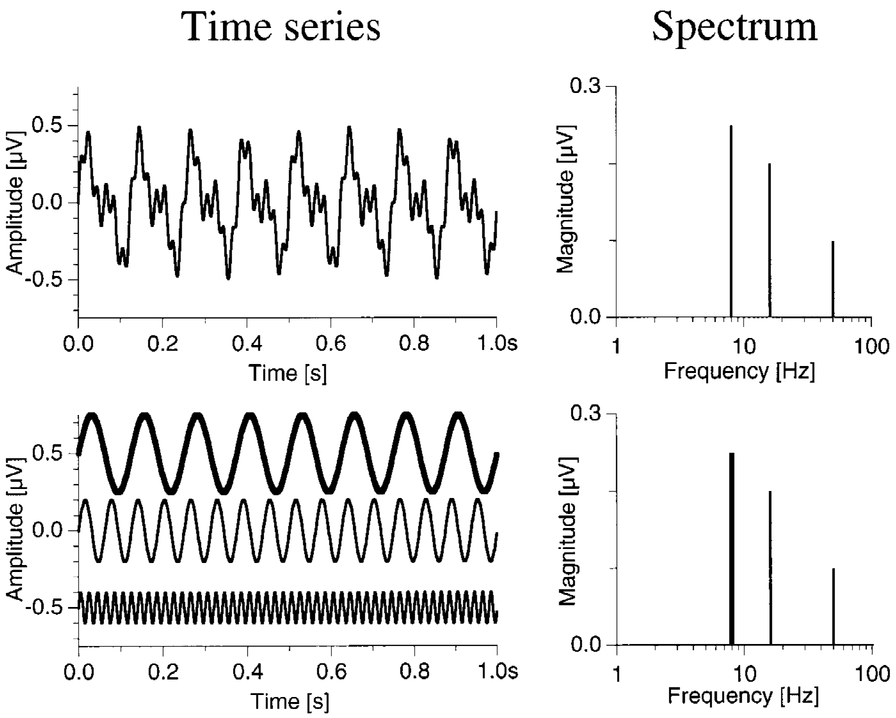

Figure 1. Fourier analysis and synthesis. The top part shows a somewhat irregular waveform with both slow and fast oscillations. The bottom part shows the three sinusoidal waveforms which, when added together, produce the top trace. The lowest frequency (thick trace) contains exactly 8 periods in the recording interval (=analysis interval) of 1 s length. Thus the corresponding spectral line (right) is located at 8 Hz. The spectrum further reveals the second frequency of 16 Hz and a third 50 Hz component which could well stem from a non-physiological source like mains interference (60 Hz in the US). This figure represents Fourier analysis when it is viewed from top to bottom (from the complex waveform via its spectrum to the underlying sinusoids). The figure represents Fourier synthesis when viewed in the opposite sequence: 3 simple sinusoidal waveforms (bottom left) lead to a rather complex waveform (top left) due to a superposition of the different peaks and troughs.

What is special in electrophysiological steady-state situations?

When we record responses to rapid stimulation, there are a number of parameters completely under our control – if limited by the available machinery. We know exactly the stimulation frequency, and we know that phase remains stable over time [3]. We are free to choose – within limits – the recording time and the analysis interval. This means that we can apply Fourier techniques to the entire record or we can pre-average, yielding a shorter analysis interval (the nature of compromise involved here is discussed below under ‘Number of stimuli per sweep’). The practical consequence is that we can choose an analysis interval that comprises an exact integer number of stimulus epochs (e.g. for a 8 rev/s stimulus we could choose a 1.125 s analysis interval which would contain 9 responses). Very often a computer monitor is used to present the stimuli. Then it is only possible to use stimulation frequencies that are integer dividers of the monitor frame rate [4]. Analog-to-digital conversion typically has a temporal resolution of $\sim$1 ms, and it is very desirable that there is an integer relation between the sampling rate and the monitor frame rate (we do this by triggering analog sampling from a counter that is clocked by the monitor's line frequency). Finally, knowing exactly the stimulation frequency allows us to apply narrow-band filtering yielding a high noise-rejection ratio.

So what is special in visual electrophysiology? We know exactly the stimulation frequency and we can choose stimulus and recording parameters to avoid all pitfalls described below.

The basic assumption of Fourier analysis is that the analysis interval is only one cycle of a periodically repeating time series. Thus only frequencies with an integer number of cycles can be extracted. Frequencies with a period longer than the analysis interval will be aliased and may appear as 'trend' (see below). Figure 2 illustrates that a 2 Hz sinusoid (with 4 peaks in a 2 s interval) will result in a clean spectrum with only one line at 2 Hz. However, if the frequency is lower and only 1.75 periods appear in the same time interval, it looks 'clean' to our eye, but Fourier analysis will detect many additional frequencies, peaking close to the actual frequency (Figure 2B, right). This effect is called 'leakage' or 'overspill'.

The emergence of the high harmonics can be understood when one reconsiders that the Fourier analysis assumes the analysis interval to be periodic. By concatenating two traces, a vertex appears in the center (Figure 2C). Many high and some low frequencies need to be 'mixed in' by the Fourier analysis to approximate this sharp vertex.

As we know exactly the stimulus frequency in our electrophysiological recordings, we can avoid this artifact and choose the length of the analysis interval such that it exactly comprises an integer number of cycles. (When employing a raster-based CRT, the stimulus frequency also has to be an integer divider of the frame frequency [4].)

What about analyzing old data, where Fourier analysis was not envisaged at the time of the recording? Figure 3 shows the recording of a response to a 3 Hz stimulus with a sampling rate of 50 Hz. The analysis interval of 0.4 s contains 20 data points. Thus two requirements for the application of Fourier analysis are violated: The stimulus period of 333 ms is not a multiple of the sampling rate and the analysis interval does not contain an integer number of stimulus periods. The spectrum consequently shows overspill artifacts (top

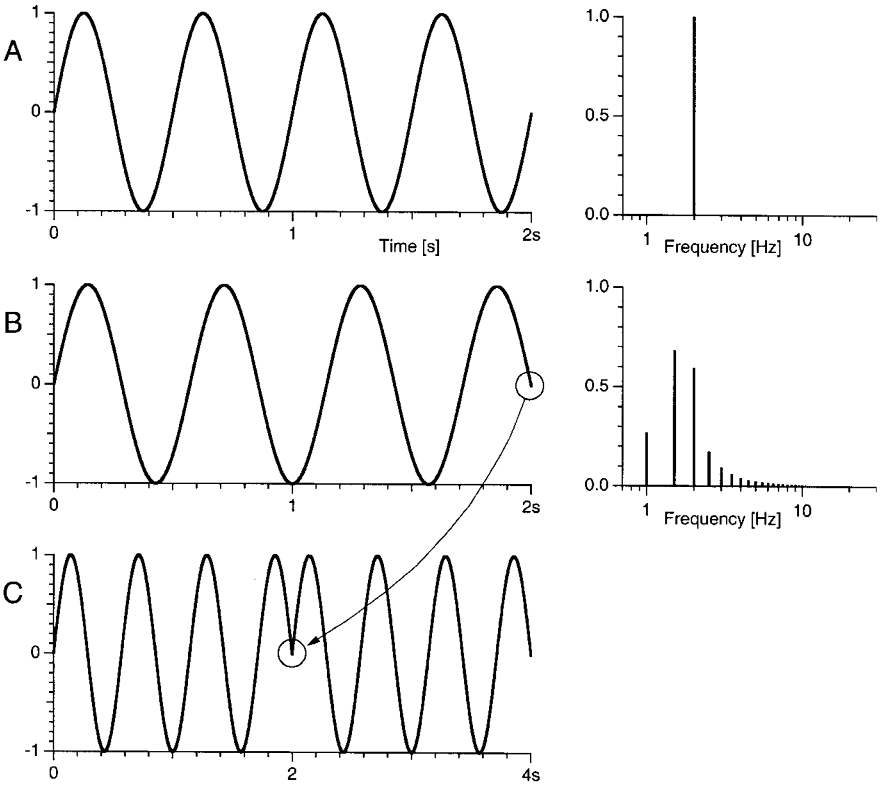

Figure 2. 'Overspill'. The top part shows a 2 Hz sinusoid over a 2 s analysis interval (left) and its 'clean' one-line spectrum on the right. Below is a 1.75 Hz sinusoid in a 2 s analysis interval. Since this is a non-integer number of cycles, this frequency cannot be represented by Fourier analysis, leading to 'overspill' in the spectrum at right. The emergence of the high harmonics can be intuitively understood when one considers that the Fourier analysis considers the analysis interval to be periodic. By cyclically extending the trace, a sharp vertex (high frequencies) occurs (bottom, 4 s analysis interval).

right). After proper interpolation (demonstrated in Figure 6) a 'clean' on-line spectrum results (bottom).

Trend artifacts

A phenomenon closely related to overspill occurs when the record contains many 'trend artifacts', as we call it. A record like that of Figure 4A is often obtained with the pattern ERG: The response at 16 Hz is superimposed on a rising ramp.

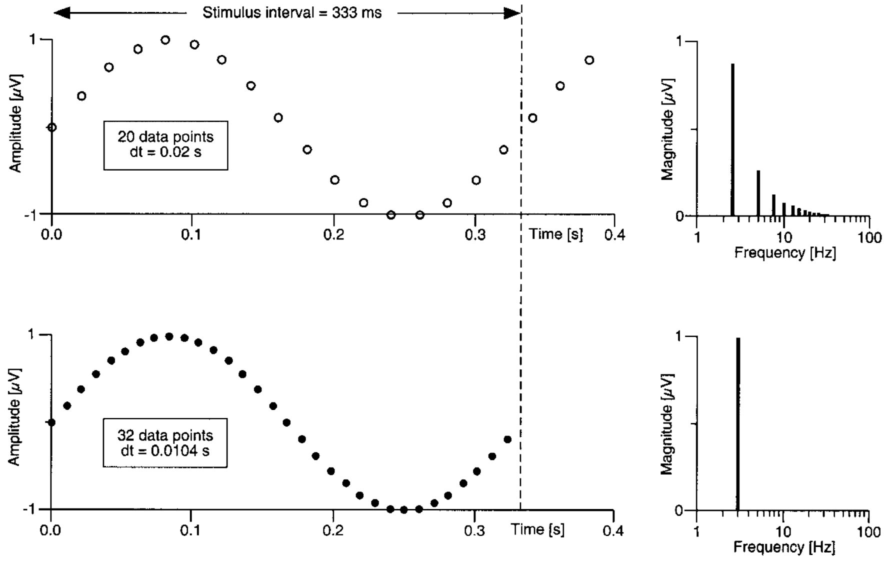

Figure 3. Post-hoc repair of an inappropriate sweep length. The top left trace violates two requirements for discrete Fourier analysis: The stimulus period of 333 ms is not a multiple of the sampling rate and the analysis interval does not contain an integer number of stimulus periods. Consequently, the Fourier spectrum shows overspill artifacts (right). The bottom part of Figure 3 shows that even under these difficult conditions a Fourier analysis can be performed which yields a spectrum without overspill. The trick is to interpolate the original waveform by ‘resampling’, an operation that yields a second waveform that fulfills all requirements mentioned. Suppose we want the second waveform to contain exactly one stimulus period of 333 ms and 32 data points. Then the sampling rate for the second waveform would be 333 ms/32=10.4 ms (Figure 3 bottom). When each of 32 data points receives the interpolated value of the original waveform, the corresponding Fourier spectrum contains only one component at 3 Hz (Figure 3 bottom right). Some requirements for interpolation are discussed with Figure 6.

This can result from blink-artifact removal: The momentary eye closure causes a large excursion (here negative), exceeding the artifact window of the averager, which consequently rejects this sweep. This negative excursion slowly decays, following the exponential time course set by the high pass filter (or time constant) of the amplifier. At a specific point in time, the signal is admitted again by the artifact rejection window while the exponential decay is still following its course. This decay will add to the response, and a short part of the exponential decay will look like the linear ramp in Figure 4A.

Both components of the trace have their counterpart in the spectrum: While the 16 Hz response (plus its 32 Hz harmonic) still sticks out, there are also many lower frequencies, which roughly represent a typical ‘sawtooth spectrum’. Again, we need to recall that Fourier analysis views the analysis interval as one cycle of a periodically repeating waveform. When many such

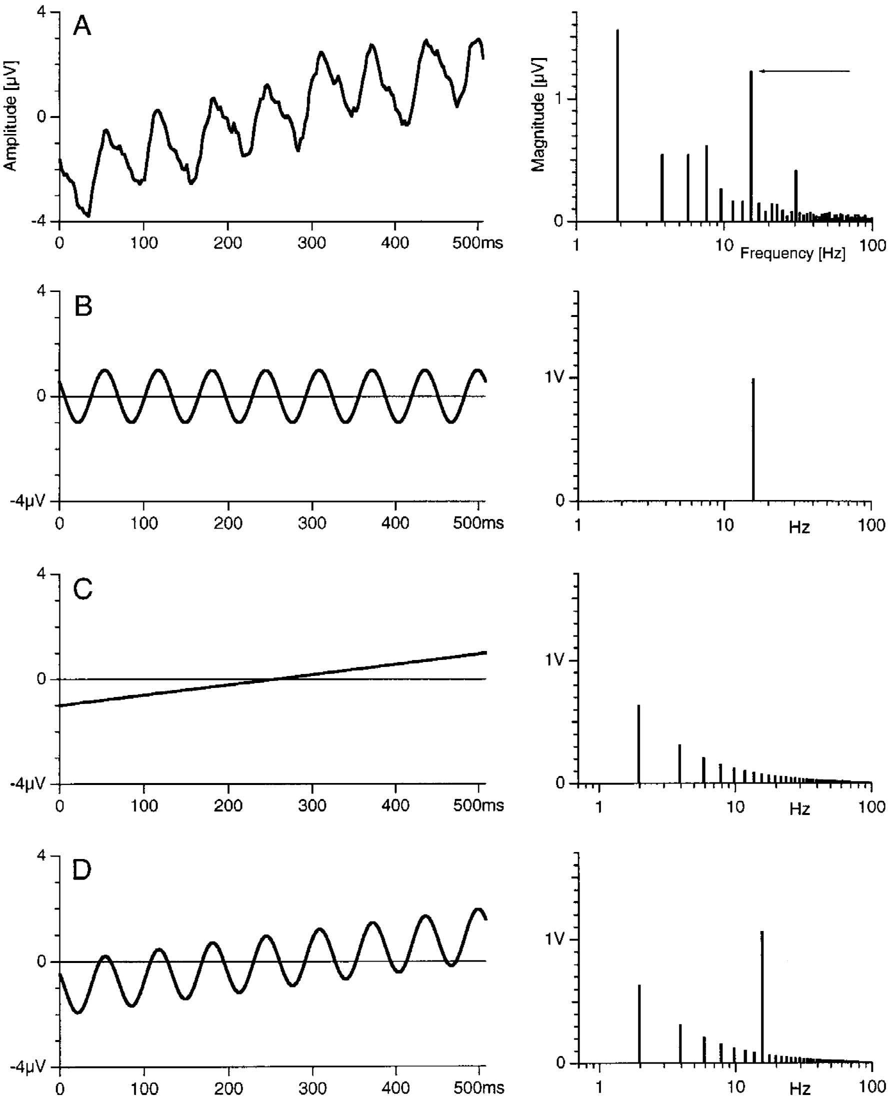

Figure 4. ‘Trend artifact’. Trace A shows a typical PERG recording, where the response (8 peaks) is superimposed on a rising ramp. Traces B–D show a simulation with fabricated data. Although the ‘noise’ (the ramp) and the response (the 16 Hz sinusoid) are of equal size, in the spectrum on the right this strong artifact results only in a 10% error of the response magnitude.

ramps (including the step at the end) are concatenated, they look like the teeth in a saw, hence the term ‘sawtooth’. The sawtooth spectrum is well known, it contains all harmonics with systematically decreasing amplitude (Figure 4C). If a pure sinusoidal response (Figure 4B) is added to the ramp of Figure 4C, the result (Figure 4D) closely resembles the actual recording (Figure 4A), illustrating the term ‘trend artifact’. The spectrum of Figure 4D also looks like it was the sum of the spectra of Figure 4A and B. Since the spectra, as depicted here, only represent the magnitude of the response at a given frequency and not the phase, the composite magnitude spectrum is unlikely to be an *exact* sum of the participating components.

Since eye movement and blink artifacts cannot be completely avoided in routine clinical situations, we suggest that the recording be set up such that a sizable number of stimulus cycles (e.g., ≥ 4) are contained in it. This separates the sawtooth harmonics sufficiently wide from the desired response so that the resulting error becomes acceptable. This point will be taken up in the section ‘Number of stimuli per sweep’ below.

Another possibility is to subtract the trend from the recorded trace before performing the Fourier analysis. How can we estimate the trend? When we assume that the analysis interval contains an integer number of stimulus cycles then the difference between the first and the last data point gives us an approximation of the trend. By subtracting this trend the sawtooth artifact is strongly reduced. A problem with this approach is that other noise sources (e.g., 50 Hz mains interference) may also introduce trend artifacts. Thus it would be the best way to filter out these high frequency components in a first Fourier analysis before trend subtraction and a second Fourier analysis are performed. While this procedure might be too complicated for the routine use, it may help to analyze ‘pathological data’ where both trend artifacts and large noise intrusion are present.

Windowing? – Not necessary

Textbooks suggest to ‘window’ the data [5, 6]. This means that the signal (Figure 5A) is multiplied over the analysis interval by a function starting with zero, reaching at (or before the middle of the interval) a value of 1.0, regressing to zero at the end of the interval. The result appears in Figure 5B. The choice of the window function is more art than science [5], here a Hanning window was used (a ‘raised cosine’).

Why and when can windowing be useful?

- Windowing does make sense in single transient recordings, e.g. when oscillations of car crash dummies’ limb movements are to be analyzed. In

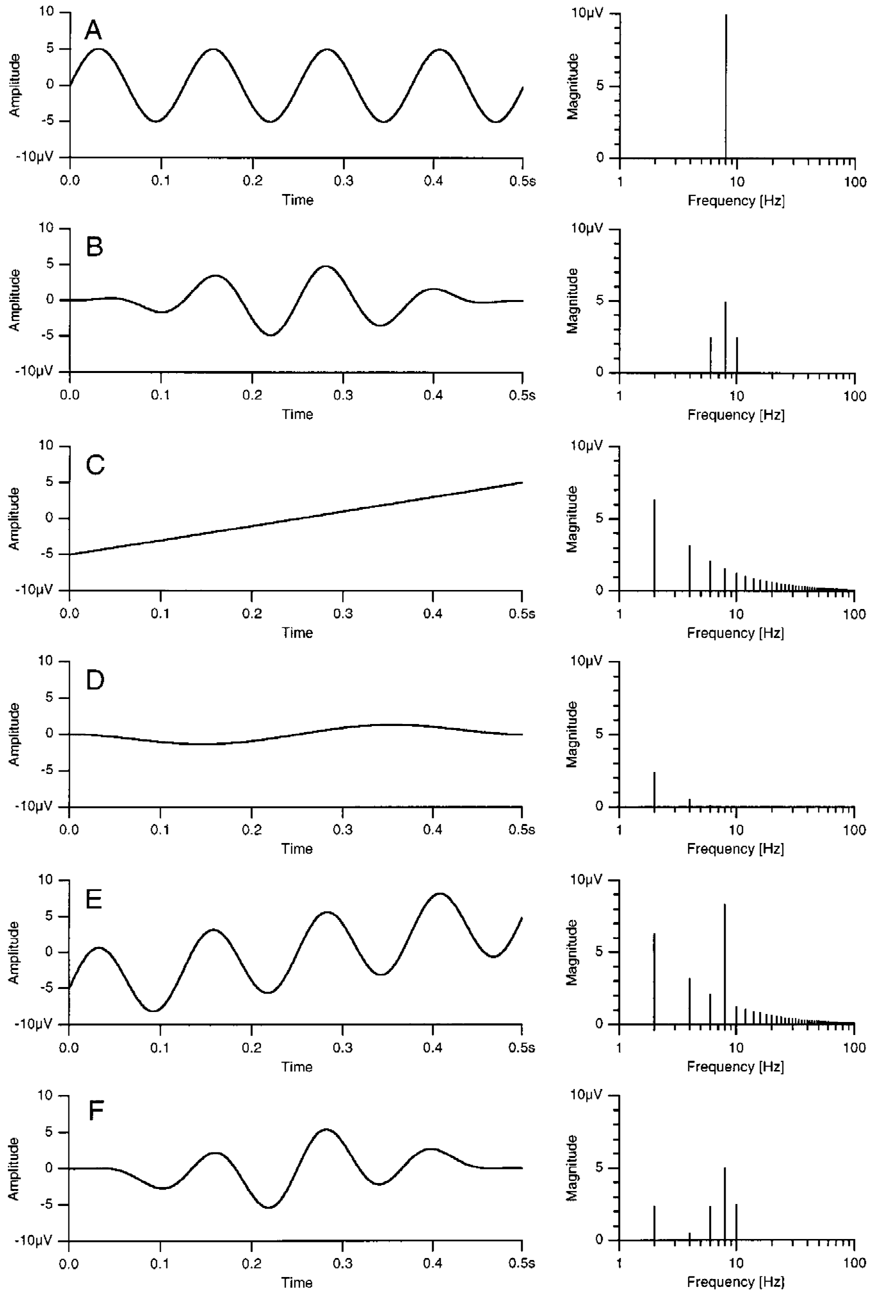

Figure 5. ‘Windowing’. A: When steady-state data is analyzed with no window (which corresponds to a rectangular window), a ‘clean’ single-line spectrum obtains. B: When the raw data is multiplied by a Hanning window, the amplitude decreases to zero at the start and end of the sweep. Such a ‘modulated’ signal has a broadened spectrum (on the right of trace B), and ‘sidebands’ have developed. Windowing can be useful when strong trend artifacts are present (C–F), although with a known signal frequency it is still not the method of choice.

the field of visual electrophysiology it might be fruitful for the analysis of oscillatory potentials in the ERG.

- When the record contains unknown frequencies, so the ‘integral number of cycles’ (see Overspill, above) cannot be applied.
- Reduction of trend artifacts (see above). This makes sense, in principle, because it reduces the influence trend artifacts (Figure 5C, D and E, F), though at the cost of broadening the signal spectrum.

However, windowing is not useful in analyzing steady-state responses: When a sinusoid proper is multiplied with a window function, ‘overspill’ in the spectrum necessarily occurs, as shown in Figure 5A and B. The lines in the spectrum adjacent to the 8 Hz response (also called ‘sidebands’) arise from the fact that after windowing the sinewave is no longer of constant amplitude, rather it is modulated. In clinical electrophysiology the stimulus frequency is normally known exactly. Thus windowing is not necessary and trend artifacts can be otherwise dealt with, as shown elsewhere in this treatise. Windowing introduces artifacts, the overspill, and reduces the amplitude estimate at the response frequency. We suggest not to use windowing whenever possible.

128, 256, 512, . . . – Not necessary

Many experimental parameters determine the number of data points per sweep, including time sampling resolution (e.g., 1 ms), stimulus frequency and the number of stimuli per sweep. This will often lead to an arbitrary number (e.g. 40, 203, 1000). However, the fast algorithms to do the DFT (FFT, [7]; Fast Hartley transform, [8], excepting [9]) only operate when the number of points per analysis interval is a power of two (e.g., 64, 128, 256, . . .). Textbooks here suggest to fill up with zeros, that is to add zero values at the end of the sweep until the next larger power-of-2 number is reached. However, this plays havoc with the spectrum of periodic data (Figure 6B). A much better solution is to ‘resample’ the record: The number of data points is increased to the next higher power-of-2 number by interpolation (Figure 6C). When the original record samples the data well above the Nyquist limit (that is: the sampling frequency is more then twice the highest signal frequency) this interpolation does not introduce additional noise. Interpolation is numerically very easy and efficient to do and should always be preferred to ‘zero fill-up’. In Figure 3, where we wanted to avoid overspill by achieving an integer number of stimulus cycles per analysis interval, the same interpolation was performed.

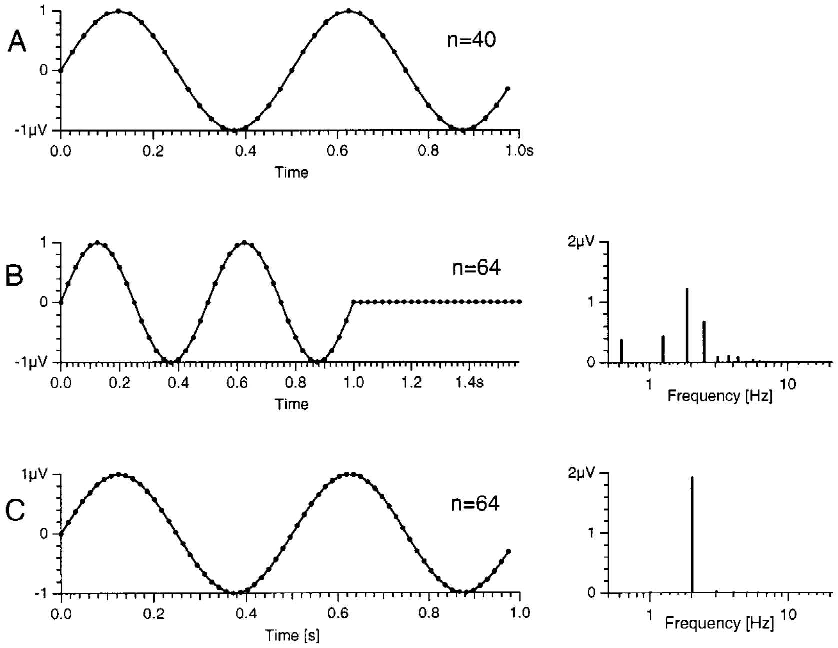

Figure 6. ‘Power of 2’-problems. Trace A shows a sinusoid, sampled with 40 data points. Many FFT algorithms require a ‘power-of-two’-number of data points, e.g. 64, 128, 256, ... When this requirement is fulfilled by adding trailing zeros (trace B), the spectrum contains many spurious signals (right of trace B) due to the non-cyclic extension. A better approach is to interpolate the trace, such that the next higher power-of-two (here 64) is reached, the appropriate single-line spectrum results (right of trace C).

Response and noise

Unfortunately, few studies employing steady-state recording assess the signal-to-noise ratio (SNR), although various methods are available. A technique with high power is the ‘t_{CIRC} statistic’ [10] which requires the non-averaged raw data set. A promising new approach was recently published by Sieving et al. [11] called ‘cycle-by-cycle Fourier analysis’. We have calculated the statistics of a simple post-hoc test, where SNR is defined as magnitude at the stimulus frequency over the average of the magnitude at the two neighboring frequencies [12]. This yields simple rules of thumb: SNR=2 corresponds to P=0.12; SNR=2.8 to p=0.05; and SNR=3 corresponds to p=0.04. Or, for the inverse question, a p-value of 0.01 requires SNR$\geq$4.84; p of 0.001 requires SNR$\geq$8.55. As the frequency resolution can be improved by enlarging the analysis interval, it may be beneficial when the analysis interval contains 10 or more stimulus cycles. Then the two neighboring frequencies are close enough to serve as reliable noise estimates for the stimulus frequency.

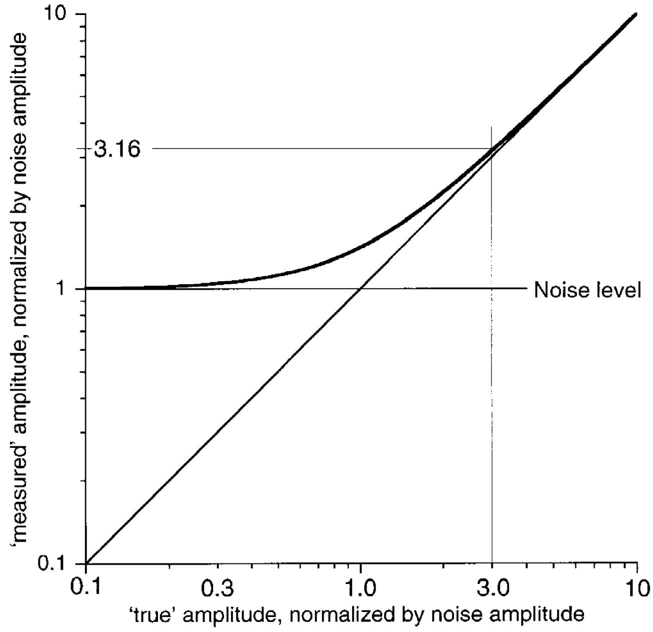

Figure 7. 'Response and noise'. The ordinate represents the 'true' amplitude, expressed in multiples of noise. The ordinate plots the spectral magnitude, which is extracted from the spectrum. It can never be less than noise, but rapidly approaches the veridical amplitude with increasing signal-to-noise ratio. When the response is three times as large as the noise there is only a 5% error in the spectral estimate [13, 14].

Noise and true response (magnitude in the frequency domain) do not add linearly [13, 14]. This means that small amplitudes will likely be noise only. However, if the amplitude is three times the noise estimate [as defined above after [12]], there is little noise (5.3%, see Figure 7: measured amplitude 3.16 for a true amplitude of 3.0), and the response is significant on the 5% level [12].

Number of stimuli per sweep

Some researchers select their sweep length in such a way that it contains exactly one cycle of the stimulus. We suggest that it is beneficial to use a higher number of stimulus periods per analysis interval which allows to recognize 'trend artifacts' that may arise from blink rejection and electrode drift. A linear trend will generate a sawtooth-spectrum. If the response frequency is high as compared to the lowest possible frequency (representing the fundamental of the sawtooth), the response will not be markedly affected by sawtooth harmonics. If there is only one response period per sampling interval, the entire low frequency noise will ‘fold’ into the response frequency. When the statistical significance of steady-state response is to be evaluated as described above (response and noise) the number of stimulus periods per analysis interval should be larger than 10 [12].

Conclusions

Discrete Fourier analysis (DFT) can greatly enhance analysis of our results. As we have much a-priori knowledge about the signal when looking at evoked responses of the visual system, we can:

- Choose an integer number of stimuli per sweep, avoiding overspill
- Choose an unconstrained number of points per sweep and interpolate for fast HFT or FFT
- Choose an integer relationship between stimulus period and frame rate
- Avoid windowing
- Disambiguate the response from trend artifacts
- Make use of noise estimates
- Estimate the significance of our responses

Hopefully, some of the techniques described here are beneficial to further development of our exciting field, the clinical electrophysiology of vision.

References

1. Marmor MF, Zrenner E. Standard for clinical electroretinography (1999 update). Doc Ophthalmol 1999; 97: 143–156.
2. Marmor M, Holder G, Porciatti V, Trick G, Zrenner E. Guidelines for basic pattern electroretinography. Recommendations by the International Society for Clinical Electrophysiology of Vision. Doc Ophthalmol 1996; 19: 291–8.
3. Simon F. The phase of PVEP in Maxwellian view: Influence of contrast, spatial and temporal frequency. Vision Res 1992; 32: 591–9.
4. Bach M, Meigen T, Strasburger H. Raster-scan cathode ray tubes for vision research – limits of resolution in space, time and intensity, and some solutions. Spatial Vision 1997; 10: 403–14.
5. Press WH, Flannery BP, Teukolsky SA, Vetterling WT. Numerical recipes. The art of scientific computing. Cambridge: Cambridge University Press, 1986: Pages.
6. Papoulis A. Probability, random variables, and stochastic processes. New York: McGraw-Hill, 1984: Pages.
7. Cooley, JW, Tukey J. An algorithm for the machine calculation of complex fourier series. Math Comput 1965; 19: 297–301.
8. O’Neill MA. Faster than fast Fourier. Byte 1988; 4: 293–300.
9. Winograd S. On Computing the Discrete Fourier Transform. Math Comput 1978; 141: 175–99.

10. Victor JD, Mast J. A new statistic for steady-state evoked potentials. Electroenceph Clin Neurophysiol 1991; 78: 378–88.
11. Sieving PA, Arnold EB, Jamison J, Liepa A, Coats C. Submicrovolt flicker electroretinogram: Cycle-by-cycle recording of multiple harmonics with statistical estimation of measurement uncertainty. Invest Ophthalmol Vis Sci 1998; 39: 1462–9.
12. Meigen T, Bach M. On the statistical significance of electrophysiological steady-state responses. Doc Ophthalmol 1999; 98: 207–232.
13. Strasburger H. The analysis of steady state evoked potentials revisited. Clin Vision Sci 1987; 1: 245–56.
14. Norcia AM, Tyler CW, Hamer RD, Wesemann W. Measurement of spatial contrast sensitivity with the swept contrast VEP. Vision Res 1989; 29: 627–37.

Address for correspondence: M. Bach, Universitäts Augenklinik, Killianstrasse 5, D-79106 Freiburg, Germany
Phone +49 (761) 270-4060; Fax: +49 (761) 270-4052; E-mail: bach@uni-freiburg.de