

11. Spectral Estimation

We have seen that the Fourier coefficients (periodogram) of a normal white noise process are uncorrelated (and thus independent) random variables, whose phase has no information except that its underlying probability density is uniform. The squares of the absolute values are distributed in χ_2^2 such that the mean values $\langle a_n^2 + b_n^2 \rangle = \sigma_\theta^2/N$ (if one prefers to use the complex Fourier series, $\langle |\alpha_n|^2 \rangle = \sigma_\theta^2/2N$, with the understanding that $\langle |\alpha_n|^2 \rangle = \langle |\alpha_{-n}|^2 \rangle$, that is, half the variance is at negative frequencies, but one could decide to double the power for positive n and ignore the negative values). Its variance about the mean is given by (10.93).

Suppose we wished to test the hypothesis that a particular time series is in fact white noise. Figure 10b is so noisy, that one should hesitate in concluding that the Fourier coefficients have no structure inconsistent with white noise. To develop a quantitative test of the hypothesis, we can attempt to use the result that $\langle |\alpha_n|^2 \rangle$ should be a constant independent of n . A useful, straightforward, approach is to exploit the independence of the neighboring Fourier coefficients. Let us define a power spectral estimate as

$$\tilde{\Psi}^\nu(s_n) = \frac{1}{[\nu/2] + 1} \sum_{p=n-[\nu/4]}^{n+[\nu/4]} |\alpha_p|^2 \quad (11.1)$$

where for the moment, ν is taken to be an odd number and $[\nu/4]$ means “largest integer in $\nu/4$ ”. That is the power spectral estimate is a local average over $[\nu/2]$ of the squared Fourier Series coefficients surrounding frequency s_n . Fig. 11a shows the result for a white noise process, where the average was over 2 local values surrounding s_n (4 degrees of freedom). Fig. 11b shows an average over 6 neighboring

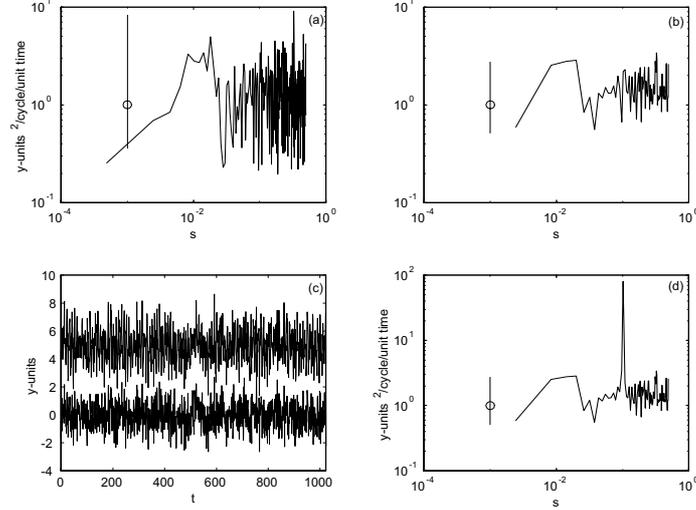


FIGURE 11. (a) Power density spectral estimate of a white noise process averaged over 2 frequency bands (4 degrees-of-freedom), and (b) over 6 frequency bands (12 degrees of freedom). An approximate 95% confidence interval is shown as obtained from (10.76). (c) Lower curve is the white noise process whose spectra are shown in (a), (b), and the upper curve is the same white noise ($\sigma_\theta^2 = 1$) plus a unit amplitude sine wave, displaced upwards by 5 units. Visually, the presence of the sine wave is difficult to detect. (d) Power density of upper curve in (c), making the spectral peak quite conspicuous, and much larger, relative to the background continuum than the 95% confidence limit.

frequency estimates (12 degrees-of-freedom). The local averages are obviously a good deal smoother than the periodogram is, as one expects from the averaging process. The probability density for the local average $\Psi^\nu(s_n)$ is evidently that for the sum of $[\nu/2]$ variables, each of which is a χ_2^2 variable. That is, $\tilde{\Psi}^\nu(s_n)$ is a χ_ν^2 random variable with ν degrees-of-freedom (2 degrees-of-freedom come from each periodogram value, the sine and cosine part, or equivalent real and imaginary part, being uncorrelated variables). The mean $\langle \tilde{\Psi}^\nu(s_n) \rangle = \sigma_\theta^2/N$ and the variance is

$$\langle (\tilde{\Psi}^\nu(s_n) - \Psi^\nu(s_n))^2 \rangle = \frac{\sigma_\theta^4}{2N^2\nu}, \quad (11.2)$$

which goes to zero as $\nu \rightarrow \infty$ for fixed N . Visually, it is much more apparent from Fig. 11 that the underlying Fourier coefficients have a constant value, albeit, some degree of variability cannot be ruled out as long as ν is finite.

This construct is the basic idea behind spectral *estimation*. The periodogram is numerically and visually extremely noisy. To the extent that we can obtain an estimate of its mean value by local frequency band averaging, we obtain a better-behaved statistical quantity. From the probability density

of the average, we can construct expected variances and confidence limits that can help us determine if the variances of the Fourier coefficients are actually independent of frequency.

The estimate (11.1) has one obvious drawback: the expected value depends upon N , the number of data points. It is often more convenient to obtain a quantity which would be independent of N , so that for example, if we obtain more data, the estimated value would not change; or if we were comparing the energy levels of two different records of different length, it would be tidier to have a value independent of N . Such an estimate is easily obtained by dividing $\tilde{\Psi}^\nu(s_n)$ by $1/N$ (multiplying by N), to give

$$\tilde{\Phi}^\nu(s_n) = \frac{1}{([\nu/2] + 1)/N} \sum_{p=n-[\nu/4]}^{n+[\nu/4]} |\alpha_p|^2. \quad (11.3)$$

This quantity is called the estimated “power spectral density”, and it has a simple interpretation. The distance in frequency space between the individual periodogram values is just $1/N$ (or $1/N\Delta t$ if one puts in the sampling interval explicitly). When averaging, we sum the values over a frequency interval $([\nu/2]/N\Delta t)$. Because $\sum_{p=n-[\nu/4]}^{n+[\nu/4]} |\alpha_p|^2$ is the fractional power in the range of averaging, $\tilde{\Phi}^\nu(s_n)$, is just the power/unit frequency width, and hence the *power density*. (If one works with the Fourier transform the normalization factors change in the obvious way.) For a stochastic process, the power density is independent of the data length.

Exercise. Show analytically that the power density for a pure sinusoid is *not* independent of the data length.

A very large number of variations can be played upon this theme of stabilizing the periodogram by averaging. We content ourselves by mainly listing some of them, leaving the textbooks to provide details.

The average over the frequency band need not be uniform. One may prefer to give more weight to frequencies towards the center of the frequency band. Let W_n be any set of averaging weights, then (11.3) can be written very generally as

$$\tilde{\Phi}^\nu(s_n) = \frac{1}{([\nu/2] + 1)/N} \frac{\sum_{p=n-[\nu/4]}^{n+[\nu/4]} W_p |\alpha_p|^2}{\sum_{p=n-[\nu/4]}^{n+[\nu/4]} W_p}, \quad (11.4)$$

and the bands are normally permitted to overlap. When the weights are uniform, and no overlap is permitted, they are called the “Daniell” window. When the averaging is in non-overlapping bands the values of $\tilde{\Phi}^\nu(s_n)$ in neighboring bands are uncorrelated with each other, asymptotically, as $N \rightarrow \infty$. With overlapping bands, one must estimate the expected level of correlation. The main issue is the determination then of ν —the number of degrees-of-freedom. There are no restrictions on ν being even or odd. Dozens of different weighting schemes have been proposed, but rationale for the different choices is best understood when we look at the Blackman-Tukey method, a method now mainly of historical interest.

An alternative, but completely equivalent estimation method is to exploit explicitly the wide-sense stationarity of the time series. Divide the record up into $[\nu/2]$ non-overlapping pieces of length M (Fig.

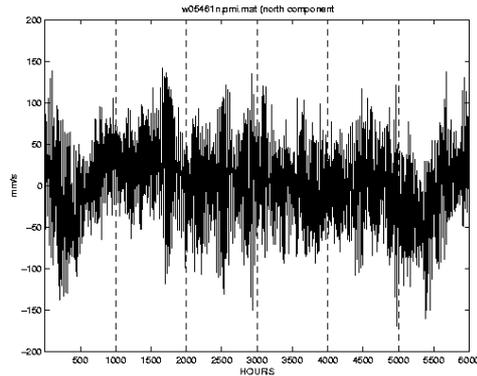


FIGURE 12. North component of a current meter record divided into 6 non-overlapping segments. One computes the periodogram of each segment and averages them, producing approximately 2 degrees of freedom from each segment (an approximation dependent upon the spectrum not being too far from white).

12) and form the periodogram of each piece $|\alpha_n^{(p)}|^2$ (The frequency separation in each periodogram is clearly smaller by a factor $[\nu/2]$ than the one computed for the entire record.) One then forms

$$\tilde{\Phi}^\nu(s_n) = \frac{1}{([\nu/2] + 1)M} \sum_{p=1}^{[\nu/2]} |\alpha_n^{(p)}|^2 \quad (11.5)$$

For white noise, it is possible to prove that the estimates in (11.3, 11.5) are identical. One can elaborate these ideas, and for example, allow the sections to be overlapping, and also to do frequency band averaging of the periodograms from each piece prior to averaging those from the pieces. The advantages and disadvantages of the different methods lie in the trade-offs between estimator variance and bias, but the intent should be reasonably clear. The “method of faded overlapping segments” has become a common standard practice, in which one permits the segments to overlap, but multiplies them by a “taper”, W_n prior to computing the periodogram. (See Percival and Walden).