

9.4 Properties of the power spectrum

For input white Gaussian noise of unit variance, the FFT delivers real and imaginary parts which are independent and each of variance $1/2$ (modulo normalization factors in the particular implementation of the transform, be careful of these!).

Why is this? A common definition of the transform of data x_i , with variance σ^2 is

$$X_k = \frac{1}{\sqrt{N}} \sum_i x_i \exp\left(\frac{2\pi i(i-1)(k-1)}{N}\right)$$

and so

$$\text{var}(X_k) = \frac{1}{N} \sum_{i,j} \langle x_i x_j^* \rangle \exp\left(\frac{2\pi i(i-j)(k-1)}{N}\right)$$

and for white noise, the correlation is zero except when $i = j$, so $\langle x_i x_j^* \rangle = \sigma^2 \delta_{ij}$ and

$$\text{var}(X_k) = \sigma^2.$$

A little extra work shows that the real and imaginary parts of X_k are equal and uncorrelated, so they must each have variance $\sigma^2/2$.

The power spectrum, being the sum of squares of these, has each component distributed like chi-square, with two degrees of freedom. Calling the power spectrum P_k , the k denoting the Fourier variable, we have

$$\text{prob}(P_k) = e^{-P_k}$$

with cumulant

$$C(P_k) = 1 - e^{-P_k}.$$

The distribution of the minimum $Y = P_k$, in a spectrum N long, is easily found to be

$$N e^{-NY}$$

and so clearly N does not have to be very large for the typical smallest value in the spectrum to be very small indeed. This is why plotting the log of a power spectrum needs some attention to the scaling!

The distribution of the maximum X is likewise

$$N e^{-X} (1 - e^{-X})^{N-1}$$

and taking the derivative and setting equal to zero gives the most likely X to be $\log N$. But what is N ? Is it the length of the input data, or the length of the power spectrum? Clearly the latter – but for real input data, the power spectrum is symmetrical and so

only half of the values are distinct. So in this case “ N ” means *half* the length of the input data train. This is a slightly subtle point, and in fact the meaning of N has shifted half-way through this solution (because I only realized this myself as I prepared the solution!). What matters for the statistics is the number of samples from which you can select the maximum.

Note that it is fundamental that the input data and the transform are of equal length. Padding the data out with zeroes, as is sometimes done for other reasons, means that the spectral components are no longer statistically independent.