Lecture 12: Windowing and Tapering for Robust Spectral Estimation

Winter 2015

References:

Hartmann Atm S 552 notes, Chapter 6.1-2.

Percival, D. B., and A. T. Walden : *Spectral Analysis for Physical Applications*, Cambridge Univ. Press, 1993.

12.1 Stationary time series

A time series or random process is called **stationary** if its statistical properties do not change with time, e. g.. white and red noise. Fourier spectral analysis is particularly useful for stationary random processes because they do not have systematic trends that violate the periodicity assumptions inherent in a DFT.

Many time series are not stationary when viewed as a whole, but segments of them may be approximated as stationary. For instance, speech or music are made of sounds. Each sound is composed of a set of frequencies sustained for a period of time, but the next sound will have a different frequency content.

12.2 Statistical properties of the periodogram of a stationary Gaussian random process

Consider a stationary zero-mean Gaussian random process $U_j, -\infty < j < \infty$ sampled at regularly-spaced time intervals Δt. Here we use a capital U to denote a random variable, i. e. that we will talk about statistics describing an infinite ensemble of realizations of a random process, rather than just one particular set of observations. Denote the true lagged autocovariance sequence (acvs) of this process by $a_p^U = E[U_j U_{j+p}], -\infty < p < \infty$. The true power spectral density of this random process is defined as

$$P_U(f) = \Delta t \sum_{p=-\infty}^{\infty} a_p^U \exp[-2\pi i p f \Delta t], \qquad -\frac{1}{2\Delta t} < f < \frac{1}{2\Delta t} = f_N. \quad (12.2.1)$$

Here f_N is the Nyquist frequency. The interpretation is that the true acvs gives the complex Fourier series representation of $P^U(f)$, which is a periodic function of f with period $1/\Delta t$, just like the sample acvs is the IDFT of the periodogram. We wish to estimate the true power spectral density $P^U(f)$ robustly from finite samples from the random process; this is a statistical estimation problem.

From the stationarity assumption and the definition of the DFT as a linear unitary matrix transformation, we can show (Percival and Walden Ch. 6):

1. The real and imaginary part of the DFT components $\hat{U}_m$ are each normally distributed with zero mean and equal variance. If N is large, this is approximately true even if the random process does not have a Gaussian PDF.

2. Define the frequency-normalized periodogram as

$$P^{(N)}(f_m) = S_m/\Delta f, \qquad S_m = |\hat{U}_m|^2/N^2, \qquad f_m = M/(N\Delta t) \quad (12.2.2)$$

   where $\Delta f = 1/(N\Delta t)$. Then the RV $P(f_m)$ is the sum of the squares of two identical normally-distributed RVs, which has a $\chi^2(2)$ or exponential PDF, whose *relative standard deviation* (ratio of its standard deviation to its mean) is 1, with a pronounced tail of large positive values. This explains the large random scatter in the periodogram (e. g. upper right panel of Fig. 1 of previous lecture).

3. More favorably, $P^{(N)}(f_m)$ is uncorrelated with $P^{(N)}(f'_m)$ for any other frequency $m' \neq m$ (except for indices corresponding to harmonic pairs M and $-M$, which always have equal spectral power). Roughly, this allows estimates S_m from nearby harmonics to be averaged to get an estimate of the true power spectral density with less random scatter. In practice, this effect is better achieved using *windowing and tapering* of the time series.

12.3 Windowed Fourier analysis

In windowed Fourier analysis, we divide or **window** our time series of length N into shorter segments of length N_w, each covering a time interval $T = N_w\Delta t$. For a stationary time series, the periodogram of each window of data gives an independent unbiased estimate of the power spectrum (if the time series were not stationary, the power spectrum from each window would not be expected to have the same statistical character so this averaging would be invalid). If there are $n = N/N_w$ windows, we can average these estimates. The resulting estimator (which is proportional to a χ^2_{2n} distribution), has a relative standard deviation of $n^{-1/2}$ instead of 1 for the unwindowed periodogram, giving a more robust estimate of the true power spectrum. It takes 100 windows to bring down the relative standard deviation of each spectral estimate down to 10%!

The tradeoff is that since each window is a factor n shorter, the frequency separation between power spectral estimates Δf is n times as coarse. In this sense, there is a close analogy between windowed Fourier analysis and averaging the periodogram across blocks of n adjacent frequencies.

If the data has substantial low-frequency variability, a problem with windowing is that there may be substantial endpoint discontinuities that contaminate the spectral estimates. The usual strategy to remove possible endpoint discontinuities is to multiply each window of the time series by a weight function or **taper** that smoothly goes to zero at the end points. This can be shown to affect the estimated power spectrum at the lowest frequencies (where it is fairly unreliable anyway) but not at frequencies corresponding to periods much less than the window period T.

A common choice of window and taper is the **Hann window**

$$w(t) = \begin{cases} 1 - \cos(2\pi t/T), & 0 \le t < T \\ 0 & \text{other } t \end{cases} \tag{12.3.1}$$

The smoothness of the cosine taper minimizes its distortion of the power spectrum, aka *spectral leakage*. The power of $w(t)$ averaged across the window $0 \le t < T$ is

$$\sigma_w^2 = 3/8.$$

We divide our power spectral estimate from the windowed DFT of the tapered time series by σ_w^2 to compensate for the reduction in time series power due to multiplying by the taper.

In addition, since the taper downweights data near the ends of each window, we use *overlapping windows* of length N_w, each starting $N_w/2$ samples apart. This ensures that all data is near the center of some window. The power spectral estimates from overlapping windows are not independent; their correlation coefficient depends on the taper being used, and this needs to be factored into uncertainty estimates for the power spectrum; the information to do this is given by Matlab spectral estimation functions from its signal processing toolbox.

The next lecture will give an example of overlapped windowing and tapering for spectral estimation applied to the Nino3.4 SSTA time series.