Lecture 9: DFT data analysis, power spectrum

©Christopher S. Bretherton

Winter 2015

9.1 Application of DFT to a Gaussian function

Before diving into use of the DFT for data analysis, we consider the DFT of a familiar function, the Gaussian

$$u(x) = \exp(-x^2/2).$$

Note that in this case, x rather than t is the coordinate variable. The function $u(x)$ is not naturally periodic, but since it rapidly goes to zero for large $|x|$, we can use a large enough domain $-L/2 < x < L/2$ so $u(x)$ is nearly zero at its edges, and imagine a periodic extension of the resulting truncated Gaussian. That function will have a slight derivative discontinuity at $|x| = L/2$ but for the example $L/2 = 4$ that is shown, u is so small at the boundaries that this causes negligible contamination of the DFT by spectral ringing.

The script **DFT_gauss** (class web page) goes through the steps of making a grid with spacing $\Delta x = 1/8$ starting at $x_1 = -L/2$, and doing the DFT of u sampled on this grid. Note that for efficiency of calculating the DFT, the vector $\mathbf{u}$ of sampled values has been chosen to have length $N = L/\Delta x = 64$, which is a power of 2. For convenience, we have chosen the first grid point to be at $-L/2$ rather than zero; this has no impact on the calculations. In the script, note the plots of the periodic extension and the presentations of the power spectrum. A semi-log or log-log plot is often useful to capture the dynamic range in the spectral power. In general, it is more informative to plot the power spectrum $|\hat{u}_m/N|^2$ vs. the corresponding harmonic M rather than vs. the component index m.

The spectral power is localized in the low harmonics with small $|M|$ and drops off to a negligibly small 'floor' of $O(10^{-10})$ that is due to the derivative discontinuity in the periodic extension of u introduced by the finite domain size; we could make this floor lower by increasing L. The script verifies Parseval's theorem for this example, and also shows that the average of $|u|^2$ over the domain is a good approximation to the integral of the square of the Gaussian over $-\infty < x < \infty$ (which is essentially all from the area inside the domain $-L/2 < x < L/2$), divided by the domain length.

The last part of the script explores taking the DFT of the first half of $\mathbf{u}$, which has a large endpoint discontinuity in its periodic extension. The resulting spectral ringing contaminates the power spectrum, causing much more of the spectral power to be in high harmonics M.

9.2 An example time series

We load the file **nino.mat** (class web page) into Matlab. It contains three variables, **year**, **month** (month) and **SST**. The variable SST is the average sea-surface temperature in the 'Nino3.4' region 5°S - 5°N, 120°W - 170°W in the central Equatorial Pacific Ocean, for each month in 1950-2012. This is a commonly used measure for El Nino, a slow irregular multiyear oscillation of ocean temperatures, wind patterns and rainfall whose influences reach around the world.

The script **nino1.m** first plots the Nino3.4 SST. We can calculate the sample mean (23.1 °C) and standard deviation (2.2 °C).

There is variability on top of an obvious annual cycle. Before we do any Fourier analysis, we should look whether the time series has a large discontinuity between its endpoints, i. e. how bad is the assumption that the time series can be periodically extended. One way to assess this is to compare the difference $T_1 - T_N = 1.07$°C of its first and last elements with $T_2 - T_1 = 1.09$°C and $T_N - T_{N-1} = 0.85$°C; since $T_1 - T_N$ is of comparable magnitude to adjacent differences the periodicity assumption seems quite acceptable here. If the time series had an obvious linear trend, it would be advisable to remove it before doing Fourier analysis.

For a general data vector $\mathbf{u}$ with DFT $\hat{\mathbf{u}}$, the **periodogram** is

$$S_m = |\hat{u}_m/N|^2. \qquad (9.2.1)$$

The elements S_m of the periodogram are often called the *spectral power* of index m or its corresponding harmonic M_m and sum to the second moment of $\mathbf{u}$ (normalized by dividing by N rather than $N-1$).

For the given time series $\mathbf{u} = \mathbf{T}$, the script finds the DFT $\hat{\mathbf{T}}$ and plots its periodogram vs. M. The time series length $N = 12 \times 63$ is not efficient for the FFT, but we don't do enough computation to make this an issue. Since the time series is real, the Fourier amplitude of wavenumber $-M$ is the complex conjugate of that at M, so their spectral power $|\hat{u}_{\pm M}|^2/N$ are equal. Hence we really need only plot the power spectrum for $M \geq 0$ (though here we plot all M for clarity).

The x-axis of a periodogram is often presented in units of inverse period $f_m = M/L$ (or, less commonly, angular frequency $\omega_m = 2\pi M/L$), rather than M. In this case, we would plot the **power spectral density**

$$PSD(f_m) = S_m/\Delta f, \qquad \Delta f = L^{-1}. \qquad (9.2.2)$$

(or $S_m/\Delta\omega$, if using angular frequencies), such that the integral under the PSD curve equals the data variance. This has the approach that the plotted quantities are more 'physical' and not sensitive to L and N. Be careful to specify correct units of variance $\times$ L for the PSD.

The DFT is dominated by its first component ($m = 1$, corresponding to harmonic $M = 0$, i. e. the mean). As expected $\hat{T}_1/N = 23.1$, the sample mean. There are three other prominent harmonics in the DFT, $M = \pm 63, \pm 126, \pm 189$. Since the dataset has length 63 years, this is the fundamental period L of its DFT, so these are the first three harmonics of the annual cycle.

The script zeros these harmonics and the time-mean (**SSTAhat**), and takes the inverse DFT to get an *SST anomaly* time series **SSTA** with the annual cycle of SST removed ('filtered out'). SSTA mainly shows irregular variability on multiyear timescales. Its mean is zero by construction, and its std is 1.1 °C (half that of SST).

9.3 Analyzing low-frequency variability in Nino3.4 SSTA

The SSTA periodogram has a broad maximum in spectral power for small harmonics M, which correspond to long periods $T = L/|M|$. The spectral power is *broadband*; that is, there is no obvious concentration of power in individual discrete periods now we have removed the annual cycle. Because the spectral power is concentrated at low frequencies and long periods (like red light in the visible part of the electromagnetic spectrum), the power spectrum is called *red*, by analogy to the visible light, whose lowest frequencies and longest wavelengths are red. Similarly, a *blue* spectrum would have power concentrated at high frequencies, and a *white* spectrum has similar power at all frequencies.

The periodogram is quite noisy, with a cloud of values for nearby harmonics. We will later discuss *spectral estimation* methods that often provide much less noisy estimates of the underlying power spectral density

Script **nino2.m** plots the power spectrum of the lowest frequencies with periods longer than a year ($|M| < 63$), which account for 88% of the variance (i. e. power) of SSTA. Within the noise, there seems to be a peak at $M = 1$ and a concentration of power for $11 < M < 31$ (2-6 year periods).

Peaks at the lowest frequencies must be interpreted with caution because of the periodicity assumption made by the DFT, which can *alias* variability or trends too slow to resolve with the 63-year time series into the lowest harmonics of the DFT. For now, we just note that the $M = 1$ (63-year period) in SSTA does not visually stand out and we don't consider it further.

Harmonics $11 < M < 31$ account for 57% of the variance of SSTA. Is this chance, or is there really a broad spectral peak in this 2-6 year band? We need stronger analysis tools to pursue this further.