Lecture 9:

Reading: Bendat and Piersol, Ch. 11.5.1-11.5.2

Recap

We've now examined how to compute spectra and how to use the χ^2 distribution to put error bars on our spectra. We've also talked about detrending and a bit about windowing. Now, we'll need to fill in some details to really understand windowing.

Spectral resolution

Suppose I have 1000 points collected over 1 day. Will I be able to tell the difference between the M2 (12.42 h) and solar semi-diurnal S2 (12 h) tides? What are the frequencies of these peaks? For M2, we have 24 h/12.42 h/1000 = 1.9324 cycles/1000 pts = 0.0019 cycles/measurement. For S2, we have 24 h/12 h/1000 = 2 cycles/1000 pts = 0.0020 cycles/measurement

The lowest frequency we can measure is 1 cycle/1000 points, and we count by integers: 2 cycles/1000 points, 3 cycles/1000 points, etc. So can we distinguish these peaks? How many measurements would we need, or how long a measurement period, to separate M2 and S2? Can we measure more frequently to separate M2 and S2? More density won't do the trick. We definitely need more data point?

So how long a record do we need? To separate two frequencies, we require that the record length be greater than the frequency resolution. Since $\Delta f = 1/N$, then $N = 1/\Delta f$. So to separate 12.42 h and 12 h, what does that require? We need to be able to separate 1.9324 cycles/day from 2 cycles/day. The frequency separation Δf is 0.0676 cycles/day. So

$$T = \frac{1}{\Delta f} = \frac{1}{2 - 1.9324} = 14.79 \text{ days.} \tag{1}$$

Some further convolution examples

Let's return to convolution for one more example. We've pointed out that convolution in the time domain is equivalent to multiplication in the frequency domain.

What if I compute the convolution of a boxcar filter with itself?

$$\begin{aligned}
y(\tau) &= \int_{-\infty}^{\infty} h(t)h(\tau - t)dt \quad \text{... \textit{definition of convolution}} && (2)\\
&= \int_0^1 h(\tau - t)dt \quad \text{for } 0 < \tau - t < 1 \quad \text{... \textit{adjust integration limits to represent} } h(t) && (3)\\
&= \int_0^1 h(\tau - t)dt \quad \text{for } -\tau < -t < 1 - \tau \quad \text{... \textit{subtract} } \tau && (4)\\
&= \int_0^1 dt \quad \text{for } \tau - 1 < t < \tau \quad \text{... \textit{reverse sign from} } -t \text{ \textit{to} } +t && (5)\\
&= \begin{cases} \int_0^\tau dt = \tau, & \text{for } 0 < \tau \le 1 \\ \int_{\tau-1}^1 dt = 1 - (\tau - 1) = 2 - \tau & \text{for } 1 < \tau \le 2 \end{cases} && (6)
\end{aligned}$$

The final step requires some examination of the limits of integration. We require that t be between 0 and 1, but τ can have a broader range. We can impose limits on t or on τ, somewhat separately. First, we worry only about the upper limit. If $t < \tau$, and $0 < t \le 1$, then we can consider a case when $0 < \tau \le 1$, while imposing the lower limit of the integral. This gives us limits of integration from 0 to τ for $0 < \tau \leq 1$. If we consider only the lower limit, so that $\tau - 1 < t$ with $0 < t \leq 1$ then τ can be between 1 and 2, with integration limits from $\tau - 1$ to 1. In any case, the results produce a triangle.

This is definitely easier to consider graphically, and it will show that boxcar filter convolved with itself becomes a triangle filter.

What if we repeat this again and again? It's actually analogous to what we considered for the central limit theorem. Applying a filter to a filter to a filter will eventually give us a Gaussian.

The Fourier Transform of a Boxcar

When we deal with data records of finite length, we're always taking finite sized segments of data, sort of like Fourier transforming a boxcar filter. Usually we think of a Fourier transform of this form:

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-i2\pi ft}\, dt. \tag{7}$$

If $x(t)$ is a boxcar filter, this becomes:

$$\int_{-\infty}^{\infty} h(t) e^{-i2\pi ft}\, dt = \int_{-1/2}^{1/2} e^{-i2\pi ft}\, dt = \left. \frac{e^{-i2\pi ft}}{-i2\pi f} \right|_{-1/2}^{1/2} = \frac{e^{i2\pi f/2} - e^{-i2\pi f/2}}{i2\pi f} = \frac{\sin(\pi f)}{\pi f} \tag{8}$$

This is the sinc function, which can be written as:

$$\text{sinc}\,(x) = \frac{\sin(x)}{x} \tag{9}$$

or (for digital signal processing)

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}. \tag{10}$$

We'll try to use the second definition, the "normalized sinc function", since we're trying to use a frequency f that needs to be multiplied by 2π.

Remember that convolution in the time domain corresponds to multiplication in the frequency domain. In this case, we're doing the opposite. We multiplied our data by a boxcar window in the time domain, and that's equivalent to convolving with a sinc function in the frequency domain:

$$\hat{X}(f) = \int_{-\infty}^{\infty} X(g) W(f-g)\, dg \tag{11}$$

where

$$W(f) = \frac{\sin(2\pi fT)}{2\pi fT} = \text{sinc}\,(2fT)\,, \tag{12}$$

for a record length of 2T.

Corresponding, the spectrum is convolved with $W(f)^2 = \text{sinc}(2fT)^2$. This has a central peak of width $\Delta f = \pm 1/(2T)$. The window will have a maximum value

$$|W(f)|^2 < (2\pi fT)^{-2} \tag{13}$$

The side lobes of the window are definitely problematic. And it turns out that we're not stuck with them. If we can widen the central peak of W (the Fourier transform of our window), we can lower the impact of the side lobes. To do this, we'll want to forego the rectangular window in the time domain and choose something that lets us attenuate the beginnings and ends of each segment of our data. What if we chose a triangle window? That will already give us fewer side lobes.

But we can keep going to find a window that looks more like an exponential. Leading possibilities:

1. Cosine taper:
$$w(t) = \cos^{\alpha}\left(\frac{\pi t}{2T}\right) \tag{14}$$
with $\alpha = [1, 4]$.

2. Hanning window or "raised cosine" window (developed by von Hann):
$$w(t) = \cos^2\left(\frac{\pi t}{2T}\right) = \frac{1 + \cos(\pi t/T)}{2} = 0.5 + 0.5\cos(\pi t/T) \tag{15}$$

3. Hamming window. This variant of the Hanning window was developed by Hamming.
$$w(t) = 0.54 + 0.46\cos(\pi t/T) \tag{16}$$
The Hamming window has less energy in the first side lobe but more in the distant side lobes.

Some other options include a Blackman-Harris window or a Kasier-Bessel window, and Harris (1978, Use of windows for harmonic analysis, *Proc. IEEE*) provides detailed discussion of options.

So how do you use a window?

1. First you must demean your data—otherwise, the window will shift energy from the mean into other frequencies. If you're working in segments, you should demean (and detrend) each segment before you do anything further.

2. Second, for a segment with N points, multiply by a window that is N points wide.

3. Since the window attentuates the impact of the edge of each segment, you can use segments that overlap (typically by 50%). This will give you (almost) twice as many segments, so instead of ν degrees some larger number.

4. Now Fourier transform, scale appropriately (e.g. by $\sqrt{8/3}$ for a Hanning window, to account for energy attenuation) and compute amplitudes.

Will the window preserve energy in your system? Not necessarily. You can normalize it appropriately, but windowing can shift the background energy level of your spectrum relative to the spectral peaks, and you'll want to keep track of this.

How many degrees of freedom do you have for overlapping windows. Not 2ν but close to that. Bendat and Piersol usefully say that overlapping by 50% will recover about 90% of the stability lost due to tapering.

So a quick recap. When we filter, we convolve the filter with our data in the time domain, which is equivalent to multiplying in the frequency domain. When we window, we multiply by a tapered window in the time domain, which is equivalent to convolving in the frequency domain.

Once you've created overlapping, windowed segments, then you'll need to figure out how many independent segments you really have. Clearly at a minimum you should have the equivalent of the number of segments that you would have if you did no overlapping. If you have N data points divided into segments that are $2M$ wide, then the minimum number of segments is $N/(2M)$. But with windowing, the end points of each segment are used less than the middle, making the overlapping segments more independent, so perhaps you have N/M segments.

Thomson and Emery's book provides the following table (Table 5.5) identifying equivalent degrees of freedom. (They lifted the table from Priestley (*Spectral Analysis and Time Series*, 1981, Table 6.2.)

Window type	Equivalent degrees of freedom (ν)	multiplier $\times$ double number segments
Truncated peridogram (boxcar)	N/M	m/2
Bartlett (triangle)	$3N/M$	1.5 m
Daniell (sinc)	$2N/M$	m
Parzen	$3.708614N/M$	1.354 m
Hanning	$8/3N/M$	4/3 m
Hamming	$2.5164N/M$	1.25 m

Window type	$N/M \text{var}\hat{h}/h^2$
Truncated peridogram (boxcar)	2
Bartlett (triangle)	2/3
Daniell (sinc)	1
Parzen	0.539285
Hanning	0.75
Hamming	0.7948

Priestley includes another table (Table 6.1) that contains a different set of numbers.

Extracting phase information from the Fourier coefficients

After all this effort to square Fourier coefficients, you might wonder what the real and imaginary parts are really good for. They are useful for sorting out the phasing of your sinusoidal oscillations. When is the amplitude at a maximum? To do this you can keep in mind that

$$A\cos(2\pi ft + \phi) = a\cos(2\pi ft) + b\sin(2\pi ft). \tag{17}$$

This can be rewritten:

$$\cos(2\pi ft)\cos(\phi) - \sin(2\pi ft)\sin(\phi) = \frac{a}{A}\cos(2\pi ft) + \frac{b}{A}\sin(2\pi ft), \tag{18}$$

which means that

$$\frac{a}{A} = \cos(\phi) \tag{19}$$

$$\frac{b}{A} = -\sin(\phi) \tag{20}$$

so

$$\phi = \text{atan}\left(-\frac{b}{a}\right). \tag{21}$$

Actually there's more information in the Fourier coefficients than this conveys, since you know the signs of both a and b, and not just their relative magnitudes. The arctangent function doesn't distinguish +45°from -135°, but we can. In some numerical implementations, you can address this using a function called atan2.

```
phi = atan2(-b,a);
```