Lecture 15:

Reading: Bendat and Piersol, Ch. 5.2.5, 5.2.6, 9.2

Recap

Last time we took a general look at correlation (and correlation coefficients) and their analog in spectral space: coherence. Coherence tells us how effectively two time series resemble each other at any given frequency.

We defined the cross-spectrum:

$$\hat{S}_{XY}(f_m) = \frac{\langle X_m^* Y_m \rangle}{T}. \tag{1}$$

This is complex: the real part is the co-spectrum ($C(f)$) and the imaginary part is the quadrature spectrum ($Q(f)$)—consistent with the terminology we use to describe cosine and sine being "in quadrature" with each other.

From that, squared coherence is:

$$\gamma_{xy}^2(f_k) = \frac{C^2(f_k) + Q^2(f_k)}{S_{xx}(f_k) S_{yy}(f_k)}, \tag{2}$$

where we needed S_{xx}, S_{yy} and S_{xy} to represent averages of multiple segments. Coherence is 1 if two data sets consistently oscillate in the same way in all segments we consider.

The phase $\phi(f_k) = \tan^{-1}(-Q(f_k)/C(f_k))$ tells us the timing difference between the two time series. If $\phi = 0$, changes in x and y happen at the same time. If $\phi = \pi$, then x is at a peak when y is at a trough. And a value of $\phi = \pi/2$ or $\phi = -\pi/2$ tells us that the records are a quarter cycle different.

Digression: 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

$$A\cos(\sigma t + \phi) = a\cos(\sigma t) + b\sin(\sigma t). \tag{3}$$

This can be rewritten:

$$\cos(\sigma t)\cos(\phi) - \sin(\sigma t)\sin(\phi) = \frac{a}{A}\cos(\sigma t) + \frac{b}{A}\sin(\sigma t), \tag{4}$$

which means that

$$\frac{a}{A} = \cos(\phi) \tag{5}$$

$$\frac{b}{A} = -\sin(\phi) \tag{6}$$

so

$$\phi = \text{atan}\left(-\frac{b}{a}\right). \tag{7}$$

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.

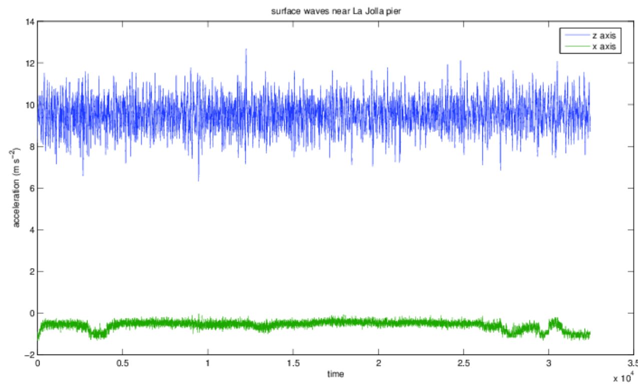

Figure 1: Time series of vertical acceleration and x-axis accleration for free-floating accelerometer near Scripps pier.

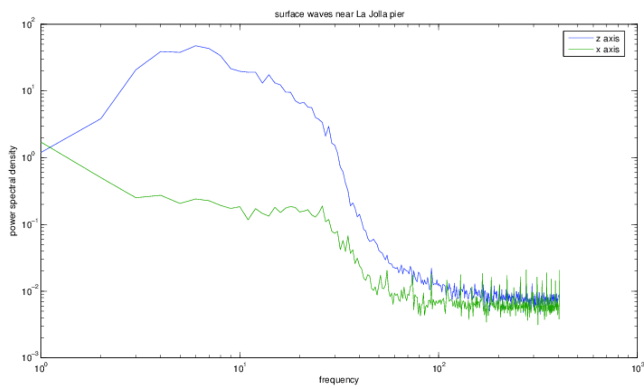

Figure 2: Spectra for vertical and x acceleration of free-floating accelerometer near Scripps pier.

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

Example: Coherence and Wave Spectra

So let's see whether surfboard acceleration measurements show any signs of coherence. We'll start by comparing vertical and horizontal accelerations of the free floating accelerometer, as shown in Figure 1. These two records have rather different spectra as shown in Figure 2. The two records are coherent, as shown in Figure 3 with a phase difference of roughly π radians, implying that they are 180° out of phase, at least at the frequencies at which they are actually coherent. In contrast, the vertical acceleration for the free floating accelerometer is not coherent with vertical acceleration from the shortboard.

Coherence: Examples

The power of coherence comes because it gives us a means to compare two different variables. With spectra we can ask, is there energy at a given frequency? With coherence we can ask whether wind energy at a given frequency drives an ocean response at a given frequency. Does the ocean respond to buoyancy forcing? Does momentum vary with wind? Does one geographic location vary with another location? Coherence is our window into the underlying physics of the system.

Let's put this to work, starting with an idealized case: Suppose we want to estimate currents entering and leaving the Mission Bay Channel. How do waves travel through the channel? We can

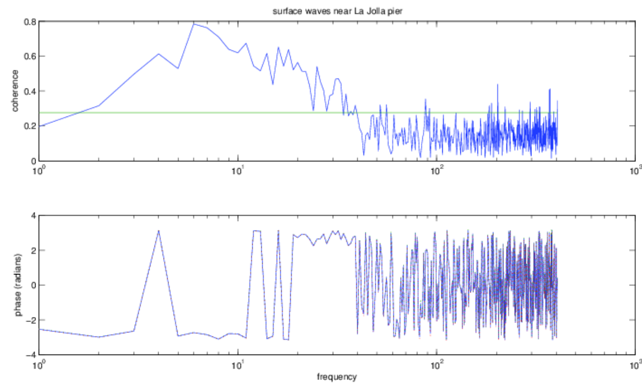

Figure 3: (top) Coherence of vertical and x acceleration of free-floating accelerometer near Scripps pier. (bottom) Phase difference between vertical and x acceleration components.

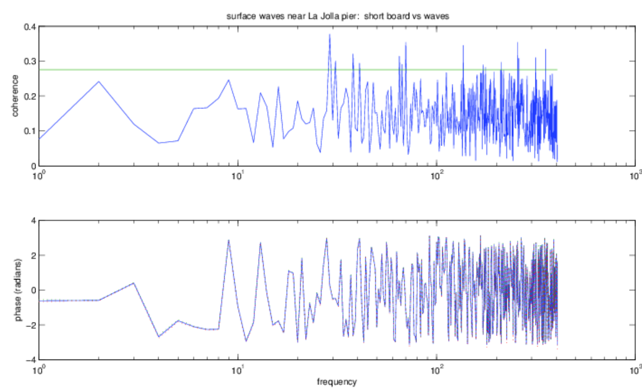

Figure 4: (top) Coherence of vertical acceleration of free-floating accelerometer versus shortboard accelerometer near Scripps pier. (bottom) Phase difference.

represent this with a dispersion relationship describing the dominant propagation in frequency-wavenumber space: $k = K(f)$.

You could imagine measuring Mission Bay by installing one current meter (with a cost of \$10-\$20,000), but another approach is to install a couple of pressure recorders along the axis of the channel (at a cost of \$1000 each). Let's assume all waves come from the ocean, and travel along the channel axis at speed $V = c + U_{current}$, where c is the wave speed and $U_{current}$ the background current speed. If the waves are surface gravity waves, $c = \sqrt{gD}$. The sensors measure time series of pressure only, so provide frequency information f. How does f relate to velocity? If we have a wavenumber $2\pi k = 2\pi/\lambda$, what does the pressure sensor see? It will detect frequencies $f = kV = k(c + U_{current})$. So we can compute the cross spectrum between our two records.

Let's test this out. We'll define a hypothetical data set:

```
lambda=10; % 10 m wavelength
V=0.3; % 0.3 m/s propagation
n2s=0.2; % noise-to-signal ratio
time=(1:5000)';
x=n2s*randn(5000,1)+cos(2*pi/lambda*V*time);
y=n2s*randn(5000,1)+cos(2*pi/lambda*V*(time)+pi/2);
```

What happens if you Fourier transform without bothering to segment? Then the data end up being unrevealing. We can demonstrate this:

```
fx=fft(x);
fy=fft(y);
sx=abs(fx(1:end/2)).^2; sx(2:end)=2*sx(2:end);
sy=abs(fy(1:end/2)).^2; sy(2:end)=2*sy(2:end);
cxy=conj(fx(1:end/2)).*fy(1:end/2); cxy(2:end)=2*cxy(2:end);
C=abs(cxy)./sqrt(sx.*sy);
plot(C)
```

In this case, the coherence is 1 everywhere. Why is that? Because without averaging, we're merely computing:

$$\gamma_{xy}^2 = \frac{(X^*Y)^*(X^*Y)}{X^*XY^*Y} \quad (8)$$

$$= \frac{XY^*X^*Y}{X^*XY^*Y} \quad (9)$$

$$= \frac{X^*XY^*Y}{X^*XY^*Y} \quad (10)$$

$$= 1 \quad (11)$$

We need the averaging to find out if the phase relationship between x and y is repeatable. With only one segment, both x and y are guaranteed to have information at each frequency with a definable phase relationship between x and y. The multiple segments allow us to test whether this phase relationship is relatively stable in time: does x always lead y by about the same fraction of a cycle?

To do the coherence calculation more constructively, we determine the frequency-space relationship between two data sets x_n and y_n, by first dividing them into segments and then Fourier transforming them, so that we have a set of X_k's and a set of Y_k's. When we computed spectra, we found the amplitude of each X_k and then summed over all our segments. Now we're going to do something slightly different. For each segment pair, we'll compute the product of X times the complex conjugate of Y: $X_k Y_k^*$. Then we'll sum over all the segments. In Matlab this becomes:

```
segment_length=500;
N=length(x);
M=segment_length/2; % define this value
Nseg=N/segment_length;
x_use=[reshape(x,segment_length,Nseg) ...
   reshape(x(M+1:end-M),segment_length,Nseg-1)];
y_use=[reshape(y,segment_length,Nseg) ...
  reshape(y(M+1:end-M),segment_length,Nseg-1)];
fx=fft(x_use); % should window and detrend here, but we're
               % skipping that for now
fy=fft(y_use);
sx=sum(abs(fx(1:M+1,:)).^2,2)/N; % sum over all spectra
                                 % (sum over 2nd index)
sx(2:end)=sx(2:end)*2;
sy=sum(abs(fy(1:M+1,:)).^2,2)/N; % sum over all spectra
                                 % (sum over 2nd index)
```

```
sy(2:end)=sy(2:end)*2;
cxy=sum(fx(1:M+1,:).*conj(fy(1:M+1,:)),2)/N;
cxy(2:end)=cxy(2:end)*2;  % since we multiplied the spectra by 2,
                          % we also need to multiply the cospectrum by 2

nd=size(x_use,2);
C=abs(cxy)./sqrt(sx.*sy);
phase_C = atan2(-imag(cxy),real(cxy));
```

The phase difference that emerges from this is only relevant at the phase where there is coherence energy (15 cycles/1000 points in the example above), and in that case the phase is a quarter cycle different, with relatively small error bars. If we reverse the order of x and y, we'll find negative phase, so a lead will turn into a lag.

Uncertainties of coherence and phase: What do we believe?

No estimate is complete without an uncertainty. We compute a significance level for coherence several ways. The standard approach that we discussed previously is to set a threshold for evaluating whether a calculated coherence exceeds what we might expect from random white noise. We started with the uncertainty for the squared coherence, γ^2:

$$\beta = 1 - \alpha^{1/(n_d-1)}, \tag{12}$$

where n_d is the number of segments, α is the significance level and is tyically 0.05 for a 95% significance level (see Thomson and Emery). In Matlab, the threshold for γ is:

```
gamma_threshold= sqrt(1-alpha^(1/(nd-1)));
```

An alternate formulation is presented by Bendat and Piersol (Table 9.6), who report the standard deviation of the squared coherence (γ^2) to be:

$$\delta_{\gamma^2_{xy}} = \frac{\sqrt{2}(1-\gamma^2_{xy})}{|\gamma_{xy}|\sqrt{n_d}}. \tag{13}$$

These are different metrics. One tells us whether the derived coherence is statistically different from zero; the second evaluates the range of values that would be consistent with an observed coherence.

The phase difference that emerges from this is only relevant at the phase where there is coherence energy (15 cycles/1000 points in the example above), and in that case the phase is a quarter cycle different. If we reverse the order of x and y, we'll find negative phase, so a lead will turn into a lag.

The phase error can seem a little murky. Formally, the uncertainty in the phase is often reported as

$$\delta_\phi = \sin^{-1}\left[t_{\alpha,2n_d}\sqrt{\frac{1-\gamma^2_{xy}}{2n_d\gamma^2_{xy}}}\right] \tag{14}$$

where $t_{\alpha,2n_d}$ is the "Student t distribution". Actually this definition comes from Koopmans (*The Spectral Analysis of Time Series*, Academic Press, 1974, p. 285), who specifically states

$$\sin(\delta_\phi) = \left\{\frac{1-\gamma^2_{xy}}{\gamma^2_{xy}(2n-2)}\right\}^{1/2} t_{2n-2}\left(\frac{\alpha}{2}\right), \tag{15}$$

where $\alpha/2$ is the upper cutoff point for the t-distribution, and we assume $2n-2$ degrees of freedom.

Watch out! What we actually want is not the "Student t distribution" but the inverse of the distribution. Given an upper cut-off point of $\alpha/2 = 0.975$ for the cdf of the t-distribution, we're looking for the corresponding value of the function. In case you have doubts, check Table A9.3 of Koopmans, which shows, for example, that $t(0.975, 20) = 2.086$. This can be computed in Matlab using tinv.

But when we plot this up, for our white noise case, it seems to be a complex number, since we've ended up with some out of range values for the arcsine—perhaps this isn't surprising since the phase is ill-defined for white noise. Bendat and Piersol provide a different formulation, which has the virtue of producing a real number:

$$\text{std}\left[\phi_{xy}(f)\right] \approx \frac{\left[1-\gamma_{xy}^2(f)\right]^{1/2}}{|\gamma_{xy}(f)|\sqrt{2n_d}} \tag{16}$$

Zwiers and Von Storch quote Hannan (1970) and provide:

$$\delta_\phi = \sin^{-1}\left[t_{(1+p)/2,2n_d-2}\frac{\gamma_{xy}^{-2}-1}{2n_d-2}\right] \tag{17}$$

where p is the confidence interval (e.g. 0.95), so $(1+p)/2$ and $(1-p)/2$ provide the limits for $p\%$ significance levels. This is clearly inconsistent with the uncertainty suggested by other authors, and it also appears to be a mistranscription from Hannan (1970). More on that later. In Matlab, these become:

```
% cab is coherence between a and b
% cab=abs(mean(fab,2)) ./sqrt(abs(mean(faa,2)) .* abs(mean(fbb,2)));
alpha = .05;
nd=10; % # of segments
p=1-alpha;
delta_phase = asin(tinv(.95,2*nd)*...
  sqrt((1-abs(cab).^2)./(abs(cab).^2*(2*nd))));
delta_phase2 = sqrt((1-cab.^2)./(abs(cab).^2*2*nd));
delta_phase3 = asin(tinv(.975,2*nd-2)*(1 ./cab.^2-1)/(2*nd-2));
```

The expressions are similar, though not identical. Which is most appropriate?