

Lecture 16:*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:

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

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:

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

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 coherence phase is:

$$\phi(f_k) = \tan^{-1}(-Q(f_k)/C(f_k)), \quad (3)$$

where Q is the imaginary part of the co-spectrum S_{xy} , and C is the real part of the co-spectrum. The phase 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.

Now our task is to look at some examples to gain a bit more intuition, and then to figure out what is significant.

Example: Coherence and Wave Spectra

Let's consider a real world example. A number of years ago, students in the first-year seminar on the Physics of Surfing took a pair of accelerometers out into the waves near the pier. They used one as a free floater and the second mounted onto a surfboard. The accelerometers measured vertical and horizontal acceleration.

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 surfboard.

Interpreting Phase

Let's consider a little thought experiment. What happens if you compute coherence between two data sets which are essentially the same, aside from a little noise, except that one is offset in time relative to the other. For example:




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




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

$$a = \tau_i^x + n_i \quad (4)$$

$$b = \tau_{i+7}^x + m_i, \quad (5)$$

where τ_i^x is zonal wind at time step i , n_i is one type of noise, and m_i is another noise that is uncorrelated with n_i . Assuming the noise to be fairly small, what should the coherence and phase be between a and b ?

To figure this out, we can estimate the cross-spectrum:

$$S_{ab} = S_{\tau_i, \tau_{i+7}} + S_{\tau_i, m_i} + S_{\tau_{i+7}, n_i} + S_{n_i, m_i}. \quad (6)$$

Since the noise is uncorrelated with the data τ and uncorrelated with other noise, with a large enough sample this becomes:

$$S_{ab} \approx S_{\tau_i, \tau_{i+7}}. \quad (7)$$

The wind τ is coherent with itself, albeit with a little phase lag, so we expect to find:

$$\gamma^2 = \frac{|S_{ab}|^2}{S_{aa}S_{bb}} \approx \frac{|S_{aa}|^2}{S_{aa}S_{aa}} = 1 \quad (8)$$




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.

And the phase is

$$\phi(f) = \text{atan2}(\Im G_{ab}(f), \Re G_{ab}(f)). \quad (9)$$

In this case, the phase will simply reflect the 7 timestep shift between a and b . For a frequency of n cycles per N points, one cycles is N/n time units, and the offset will represent a fraction of a cycle: $7n/N$. For higher frequencies, the phase shift will represent a linearly increasing phase.

Coherence uncertainty

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 :

$$\beta = 1 - \alpha^{1/(n_d-1)}, \quad (10)$$

where n_d is the number of segments, α is the significance level and is typically 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:

$$\delta_{\gamma_{xy}^2} = \frac{\sqrt{2}(1 - \gamma_{xy}^2)}{|\gamma_{xy}|\sqrt{n_d}}. \quad (11)$$

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.

What is n_d ?

We have a formulation for coherence uncertainty that depends on the number of segments. What if we want to use overlapping segments, just as we did for the Welch method? You can test this through a Monte Carlo process. If you set n_d equal to the total number of segments, ignoring the fact that some overlap, your error bars will be visibly too small. You can run Monte Carlo tests with overlapping segments to figure out how many effective segments you really have. And perhaps not surprisingly, the results are equivalent to what we found in the Welch method (albeit scaled by a factor of 2, since we're now counting segments and not degrees of freedom):

Window type	Equivalent number of segments (n_d)
Boxcar	2/3
Triangle	8/9
Hanning	18/19 $\approx$ 0.95
Hamming	$\sim$ 0.90

Table 1: Effective number of independent segments relative to the total number of segments, using 50% overlap. (With no overlap, assume n_d segments.)

Uncertainties of phase: What do we believe?

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.

First a little terminology. Bendat and Piersol provide a good discussion of bias and uncertainties in spectral estimators. As a starting point, the variance of the quantity that we want to estimate is

$$\text{var}[\tilde{A}] = E[\tilde{A}^2] - A^2, \quad (12)$$

where A is the true value, and $\tilde{A}$ is the unbiased estimate (so $E[\tilde{A}] = A$). For spectral estimators we tend to talk about the normalized error:

$$\epsilon^2 = \frac{\text{var}\tilde{A}}{A^2}. \quad (13)$$

Bendat and Piersol first derive relationships for the variance of the spectrum and cross-spectrum in the case of one segment and two degrees of freedom (see their appendix). They then note that variance scales with $1/n$, where n is the number of degrees of freedom, so that variance

can be inferred for spectra and cross-spectra with any number of degrees of freedom (by dividing by n_d the number of segments).

Phase error is sometimes reported as a standard error and sometimes as a 95% confidence range.

Bendat and Piersol suggest the following as a standard error of the phase:

$$\text{std} [\phi_{xy}(f)] \approx \frac{[1 - \gamma_{xy}^2(f)]^{1/2}}{|\gamma_{xy}(f)| \sqrt{2n_d}} \quad (14)$$

```
% define fab as the cross-spectrum, faa as the spectrum of the first record,
% and fbb as the spectrum of the second record.
cab=abs(mean(fab,2))./sqrt(abs(mean(faa,2)).*abs(mean(fbb,2)));
% nd = # of segments
phase_std = sqrt((1-cab.^2)./(abs(cab).^2*2*nd));
```

For a 95% range, Hannon (1970) and Koopmans (1974) provide this expression:

$$\delta_\phi = \sin^{-1} \left[t_{\alpha, 2n_d} \sqrt{\frac{1 - \gamma_{xy}^2}{2n_d \gamma_{xy}^2}} \right] \quad (15)$$

where $t_{\alpha, 2n_d}$ is identified as the “Student t distribution”, and is actually the inverse of the Student t distribution (“tinv” in Matlab). This is computed for $p=0.95$ to obtain a 95% limit.

```
delta_phase = asin(tinv(.975,2*nd)*...
sqrt((1-abs(cab).^2)./(abs(cab).^2*(2*nd))));
```

See the sidebar for more commentary as well as code

Sidebar: Deciphering phase errors from the textbooks

If you plow through spectral analysis textbooks, you’ll find a range of different contradictory guidelines for phase errors. Here’s a quick synopsis, and some guidance for deciding what you trust.

The phase error can seem a little murky. One common formulation for phase uncertainty is:

$$\delta_\phi = \sin^{-1} \left[t_{\alpha, 2n_d} \sqrt{\frac{1 - \gamma_{xy}^2}{2n_d \gamma_{xy}^2}} \right] \quad (16)$$

where $t_{\alpha, 2n_d}$ is identified as the “Student t distribution”, and is actually the inverse of the Student t distribution (“tinv” in Matlab). 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 $t(0.975, 20) = 2.086$.

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:

$$\text{std} [\phi_{xy}(f)] \approx \frac{[1 - \gamma_{xy}^2(f)]^{1/2}}{|\gamma_{xy}(f)| \sqrt{2n_d}} \quad (17)$$

Zwiers and Von Storch quote Hannan (1970)¹ and provide:

$$\delta_\phi = \sin^{-1} \left[t_{(1+p)/2, 2n_d-2} \frac{\gamma_{xy}^{-2} - 1}{2n_d - 2} \right], \quad (18)$$

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. However, Zwiers and Von Storch have misquoted Hannan (1970), who actually has a form equivalent to this:

$$\delta_\phi = \sin^{-1} \left[t_{(1+p)/2, 2n_d-2} \left\{ \frac{\gamma_{xy}^{-2} - 1}{2n_d - 2} \right\}^{1/2} \right], \quad (19)$$

which is exactly equivalent to Koopmans (1974). In Matlab, these become:

```
% cab is squared coherence between a and b
% for example:
%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? We can test this out by creating a fake data set with a known phase relationship:

```
a=randn(100,1000)+ cos(2*pi/10*(1:100)')*ones(1,1000);
b=randn(100,1000) + sin(2*pi/10*(1:100)')*ones(1,1000);
fa=fft(a);
fb=fft(b);
fab=conj(fa).*fb;
faa=conj(fa).*fa;
fbb=conj(fb).*fb;

cab=abs(mean(fab,2)) ./sqrt(abs(mean(faa,2)) .* abs(mean(fbb,2)));

m=10;
clear phase_c
for i=1:1000/m
    phase_c(:,i)=atan2(-imag(mean(fab(:,(i-1)*m+1:i*m),2)),...
        real(mean(fab(:,(i-1)*m+1:i*m),2)));
end

nd=m;
```

¹Hannan, 1970, *Multiple Time Series*, John Wiley & Sons, 536 pp. (See p. 257, equation 2.11)

```

delta_phase = asin(tinv(.95,2*nd)*...
    sqrt((1-abs(cab).^2)./(abs(cab).^2*sqrt(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));
% compare results
[delta_phase(11) delta_phase2(11) delta_phase3(11)]
std(phase_c(11,:))

```

It's clear from these tests that (a) the distribution of the phases should be roughly Gaussian, but it isn't exactly Gaussian; (b) Bendat and Piersol's representation for the standard deviation of the phase (delta_phase2) is relatively reliable; (c) the inverse sine formulations should produce phase errors representing the 95th percentile. If we plug in 0.68 in place of 0.95 for the Koopmans formulation above, we'll get results consistent with the Bendat and Piersol and Monte Carlo results.

So the bottom line is:

1. If you want a one standard deviation error bar, the Bendat and Piersol formula is pretty good.
2. If you want to be a little more accurate, and especially if you want an error bar representing a 95% probability band, then it's probably better to use the formula from Koopmans.

Coherence: The autocovariance perspective

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.

We can write the cross-spectrum for x and y :

$$\hat{S}_{XY}(\sigma_m) = \frac{\langle X_m^* Y_m \rangle}{\Delta\sigma} \quad (20)$$

where here we've normalized by frequency resolution, and we have represented frequency as σ_m rather than f_m , meaning that we've incorporated 2π into the frequency.

Just as we considered spectra as the Fourier transform of the autocovariance, we can now think about the Fourier transform of the lagged co-variance.

$$R_{xy}(\tau) = \frac{1}{2T} \int_{-T}^T x^*(t)y(t+\tau) dt. \quad (21)$$

We can rewrite this:

$$R_{xy}(\tau) = \frac{1}{2T} \int_{-T}^T \sum_{n=-\infty}^{\infty} X_n^* e^{-i\sigma_n t} \sum_{m=-\infty}^{\infty} Y_m e^{i\sigma_m(t+\tau)} dt \quad (22)$$

$$= \frac{1}{2T} \sum_{n=-\infty}^{\infty} X_n^* \sum_{m=-\infty}^{\infty} Y_m e^{i\sigma_m \tau} \int_{-T}^T e^{i(\sigma_m - \sigma_n)t} dt \quad (23)$$

$$= \sum_{n=-\infty}^{\infty} X_n^* \sum_{m=-\infty}^{\infty} Y_m e^{i\sigma_m \tau} \delta_{nm} \quad (24)$$

$$= \sum_{n=-\infty}^{\infty} X_n^* Y_n e^{i\sigma_n \tau} \quad (25)$$

$$= \Delta\sigma \sum_{n=-\infty}^{\infty} S_{XY} e^{i\sigma_n \tau}, \quad (26)$$

where we used the Kronecker delta δ_{nm} to extract only frequencies for which $n = m$, since all other modes are orthogonal. The result tells us that the lagged covariance is the inverse Fourier transform of the cross spectrum. In other words,

$$S_{XY}(\sigma_n) = \int_{-T}^T R_{xy}(\tau) e^{-i\sigma_n \tau} d\tau = \frac{X_n^* Y_n}{\Delta\sigma} \quad (27)$$

Thus we could determine the cross-spectrum from the lagged covariance.