Lecture 16:

Reading: Bendat and Piersol, Ch. 9.1-9.2

Recap

We've looked at a couple examples of coherence calculations along with some (incomprehensible) figures from published cases. The key feature of coherence is that it allows you to decide if two records vary in a consistent way at any given frequency. This is a subtle point: if you view the world from the perspective of a Fourier transform, everything is sinusoidal, and naturally all data records vary sinusoidally, although the phasing of record x could differ from the phasing of record y. When we compute coherence, we ask whether the phasing (at frequency σ or wavenumber k) between x and y is consistent between different chunks of the data records. To compute coherence we need to segment our data; for the same reasons that we detrend and window when we compute spectra, we should detrend and window when we compute coherence.

Coherence calculations produce two results: a coherence that varies between 0 and 1, and a phase that varies between $-\pi$ and π (or equivalently from 0 to 360 degrees or -180 to 180 degrees.)

Uncertainties of coherence and phase: What do we believe?

We talked about coherence uncertainties last time, but we didn't finish sorting out phase uncertainties, and I gave you three different formulations. Which one do we trust?

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

$$\text{var}[\tilde{A}] = E[\tilde{A}^2] - A^2, \tag{1}$$

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:

$$\epsilon^2 = \frac{\text{var}\tilde{A}}{A^2}. \tag{2}$$

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 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).

Uncertainty for phase is often reported with the formula I provided last time

$$\delta_\phi = \sin^{-1}\left[t_{\alpha,2n_d}\sqrt{\frac{1-\gamma_{xy}^2}{2n_d\gamma_{xy}^2}}\right] \tag{3}$$

where $t_{\alpha,2n_d}$ is identified as the "Student t distribution". 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:

$$\text{std}\left[\phi_{xy}(f)\right] \approx \frac{\left[1-\gamma_{xy}^2(f)\right]^{1/2}}{|\gamma_{xy}(f)|\sqrt{2n_d}} \tag{4}$$

Zwiers and Von Storch quote Hannan and provide:

$$\delta_\phi = \sin^{-1}\left[t_{(1+p)/2,2n_d-2}\frac{\gamma_{xy}^{-2}-1}{2n_d-2}\right], \tag{5}$$

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. In Matlab, these become:

```
% cab is covariance between a and b
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;
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, (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.

More on degrees of freedom for windowed segments

When we talked about the sinc function, windowing, and the use of overlapping segments, we also noted that when you overlap segments you reduce the number of degrees of freedom. How much do you reduce the number of degrees of freedom? If you look this up in the 2nd edition of Emery and Thomson, you find a nice table (their Table 5.6.4), ostensably lifted from Priestley. Priestley's results are nicely discussed by Koopmans (*The Spectral Analysis of Time Series*, Academic Press, 1974). Emery and Thomson describe the table as representing the equivalent degrees of freedom for windowed "block averaged" spectra.

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

Table 1: Table 5.6.4 from Emery and Thomson showing "Equivalent degrees of freedom for spectra calculated using different windows."

The windowed spectra that we've used as a basic approach in class are referred to by a number of names in the literature. Some textbooks refer to windowing as "tapering". And the formal method is sometimes called the Welch method. Percival and Walden (*Spectral Analysis for Physical Applications*, Cambridge University Press, 1993) provide a detailed discussion of this approach under the name "Welch's Overlapping Segment Averaging (WOSA)".

A little further exploration of the literature shows that the values in the table are incorrectly labeled and actually represent degrees of freedom for spectra determined by filtering or averaging adjacent frequencies from an initial spectral estimate. As we pointed out in class, if you filter in frequency, you can increase your degrees of freedom just the way you increase degrees of freedom by using multiple segments. And you can refine your filtering to reduce spectral ringing effects by using a carefully constructed filter. While people use the same types of filters for time domain windowing/tapering and for frequency domain filtering, the effect is not identical. (If we wanted an identical effect, we'd need to use the Fourier transform of the time domain window to filter in the frequency domain. You could do that, but it's not what we usually envision when we talk about windowing or tapering, and it's not what Emery and Thomson seem to describe.) What this means is that the tables of "equivalent degrees of freedom" for the frequency domain don't actually work for windowed time domain data.

However, all is not lost. Percival and Walden provide a full derivation of the actual degrees of

freedom for overlapping segments. They define the following terms:

$$\begin{aligned} N &= \text{total length of record} \\ N_B &= \text{number of blocks} \\ N_S &= \text{segment length or block size} \\ n &= \text{shift factor or number of points of overlap between segments} \\ h &= \text{window, normalized so that } h^2 \text{ sums to 1.} \end{aligned}$$

They point out that the covariance between adjacent segments determines the adjustment to the degrees of freedom, and this depends on h.

In this framework, the variance of the spectral estimate is:

$$\text{var}\left\{\hat{S}^{(WOSA)}(f)\right\} = \frac{1}{N_B}\sum_{j=0}^{N_B-1}\text{var}\left\{\hat{S}_{jn+1}(f)\right\} + \frac{2}{N_B}\sum_{j<k}\text{cov}\left\{\hat{S}_{jn+1}(f), \hat{S}_{kn+1}(f)\right\}, \quad (6)$$

where j and k are indices for separate but overlapping segments. The variance of the jth spectrum should converge to the canonical spectrum:

$$\text{var}\left\{\hat{S}_{jn+1}(f)\right\} \approx S^2(f). \quad (7)$$

The covariance depends on the overlap of the tapers or windows:

$$\text{cov}\left\{\hat{S}_{jn+1}(f), \hat{S}_{kn+1}(f)\right\} \approx S^2(f)\left|\sum_{t=1}^{N_S} h_t h_{t+|k-j|n}\right|^2, \quad (8)$$

with h_t defined to be zero when t is out of range (i.e. $t > N_S$). This means that:

$$\begin{aligned} \text{var}\left\{\hat{S}^{(WOSA)}(f)\right\} &\approx \frac{S^2(f)}{N_B}\left(1 + \frac{2}{N_B}\sum_{j<k}\left|\sum_{t=1}^{N_S} h_t h_{t+|k-j|n}\right|^2\right) \quad (9) \\ &= \frac{S^2(f)}{N_B}\left(1 + 2\sum_{m=1}^{N_B-1}\left(1 - \frac{m}{N_B}\right)\left|\sum_{t=1}^{N_S} h_t h_{t+mn}\right|^2\right). \quad (10) \end{aligned}$$

Thus for a full record with arbitrary overlap:

$$\nu \approx \frac{2N_B}{1 + 2\sum_{m=1}^{N_b-1}\left(1 - \frac{m}{N_B}\right)\left|\sum_{t=1}^{N_S} h_t h_{t+mn}\right|^2}. \quad (11)$$

This formulation allows for arbitrary levels of overlap, so you could imagine starting a new segment every data point and having to contend with with lots of complicated covariances between adjacent segments.

For practical purposes, we typically work with 50% overlap, so $n = N_S/2$. In this case, Percival and Walden show that the equation for the effective degrees of freedom simplifies to

$$\nu \approx \frac{2N_B}{1 + 2\left(1 - \frac{1}{N_B}\right)\left|\sum_{t=1}^{N_S/2} h_t h_{t+N_S/2}\right|^2}. \quad (12)$$

In the limit of large N_B and many samples, it's relatively straightforward to find an analytic solution:

$$\nu \approx \frac{2N_B}{1 + 2\left|\int_0^{L/2} h(t)h(t + L/2)\, dt\right|^2}. \tag{13}$$

subject to the requirement that the window normalization is:

$$\int_0^L h(t)^2\, dt = 1. \tag{14}$$

Thus for a boxcar filter, $h(t) = 1/\sqrt{L}$, and

$$\int_0^{L/2} h(t)h(t + L/2)\, dt = \int_0^{L/2} \frac{1}{L}\, dt = fractL|_0^{L/2} = \frac{1}{2}. \tag{15}$$

Thus

$$\nu \approx \frac{2N_B}{1 + 2\left|\frac{1}{2}\right|^2} = \frac{2N_B}{1 + \frac{1}{2}} = \frac{4N_B}{3}. \tag{16}$$

Either analytically, or by plugging in normalized discrete window values h_t, we can compute the adjustments to our effective degrees of freedom shown in Table 2. You'll see that these values provide a fairly effective match to the values that you obtained from Monte Carlo simulation.

Window type	Equivalent degrees of freedom (ν)
Boxcar	4/3
Triangle	16/9
Hanning	36/19 $\approx$ 1.90
Hamming	$\sim$1.80

Table 2: Effective number of degrees of freedom relative to the total number of segments, using 50% overlap. (With no overlap, the equivalent degrees of freedom would be double the number of segments.)

So what of the other texts? The 2014 edition of Thomson and Emery is as misleading as the earlier editions. Von Storch and Zwiers, who are usually fairly lucid on data analysis, strongly favor filtering in the frequency domain so don't consider the impact of windowing or tapering in the time domain. Priestley also focuses largely on spectra computed from the autocovariance and spectra computed by filtering the periodogram. Their published tables are intended to provide guidance on the "lag window" (e.g. $\lambda(t)$) for spectra computed from the autocovariance, and the "spectral window", $W(f)$, which is Fourier transform of the lag window. When $\lambda(t)$ and $W(f)$ are used as a Fourier transform pair, they should have equivalent impacts on the degrees of freedom.

Finally, Percival and Walden note that we can also consider overlaps other than 50%, by adjusting m in their original equation:

$$\nu \approx \frac{2N_B}{1 + 2\sum_{m=1}^{N_b-1}\left(1 - \frac{m}{N_B}\right)\left|\sum_{t=1}^{N_S-mn} h_t h_{t+mn}\right|^2}. \tag{17}$$

Their Figure 293 shows degrees of freedom as a function of overlap for the Hanning window. We can code this in Matlab to consider other windows as well, as illustrated in Figure 1:

```
Ns=512;
n=256;
N=Ns*100;
Nb_theory=N/Ns;

h=ones(Ns,1)/sqrt(Ns);
for n=1:Ns-1
Nb=round((N-Ns)/n+1);
sumh=[];
for m=1:Nb-1
 if(Ns-m*n>=1)
  sumh(m)=(1-m/Nb)*abs(sum(h(1:Ns-m*n).*h(1+m*n:Ns)))^2;
 end
end
denom=1+2*sum(sumh);
nu_boxcar(n)=2*Nb/denom;
end

h=sqrt(2/3/Ns)*(1-cos(2*pi*(1:Ns)/Ns));

for n=1:511
Nb=floor((N-Ns)/n+1);
sumh=[];
for m=1:Nb-1
 if(Ns-m*n>=1)
  sumh(m)=(1-m/Nb)*abs(sum(h(1:Ns-m*n).*h(1+m*n:Ns)))^2;
 end
end
denom=1+2*sum(sumh);
nu_hanning(n)=2*Nb/denom;
end


hold off
plot(1-(1:2:Ns-1)/Ns,nu_boxcar(1:2:end)/Nb_theory,'LineWidth',3);
hold on
plot(1-(1:2:Ns-1)/Ns,nu_hanning(1:2:end)/Nb_theory,'LineWidth',3);
set(gca,'FontSize',16)
xlabel('Fractional overlap between segments','FontSize',16)
ylabel('Effective dof relative to # non-overlapping segments',...
   'FontSize',16)
legend('Boxcar','Hanning')
```

Appendix: More detail on variance of cross-spectra

If we have two degrees of freedom, the cross spectrum is

$$
\begin{aligned}
|\hat{G}_{XY}(f)|^2 &= \hat{G}^*_{XY}\hat{G}_{XY} && (18)\\
&= |X^*(f)Y(f)|^2 && (19)\\
&= X(f)Y^*(f)X^*(f)Y(f) && (20)
\end{aligned}
$$

So trying all combinations to get the sum of the 4-term product:

$$
\begin{aligned}
\langle|\hat{G}_{XY}(f)|^2\rangle &= \langle XX^*\rangle\langle YY^*\rangle + \langle XY^*\rangle\langle X^*Y\rangle + \langle XY\rangle\langle X^*Y^*\rangle && (21)\\
&= G_{XX}G_{YY} + |G_{XY}|^2 && (22)
\end{aligned}
$$

where G here is the total cross-spectrum. The variance is then

$$
\begin{aligned}
\text{var}[\hat{G}_{XY}] &= \langle|\hat{G}_{XY}(f)|^2\rangle - |\hat{G}_{XY}(f)|^2 && (23)\\
&= G_{XX}G_{YY} && (24)\\
&= \frac{|G_{XY}|^2}{\gamma^2_{xy}} && (25)
\end{aligned}
$$

where γ_{xy} is the coherence.

With more degrees of freedom, error scales with the square root of the number of samples, just like the standard error:

$$
\begin{aligned}
\text{var}G_{XX} &= \frac{G^2_{xx}}{n_d} && (26)\\
\text{var}G_{YY} &= \frac{G^2_{YY}}{n_d} && (27)\\
\text{var}G_{XY} &= \frac{|G_{XY}|^2}{\gamma^2_{xy}n_d} && (28)
\end{aligned}
$$

This scaling gives us the uncertainty for the coherence and phase after some manipulation. (See Bendat and Piersol, Ch. 9 for details.)

This means that the normalized uncertainty of G_{XY} is

$$\epsilon[|\hat{G}_{XY}|] = \frac{\text{std}[G_{XY}]}{G_{XY}} = \frac{1}{|\gamma_{xy}|\sqrt{n_d}} \quad (29)$$

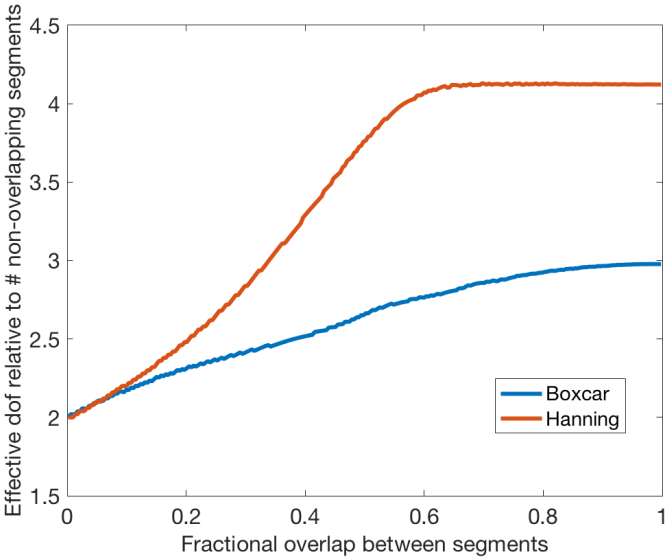

Figure 1: Ratio of degrees of freedom ν relative to nominal number of segments available if no overlapping is used for Hanning window and boxcar window.