

Lecture 8:

Reading: Bendat and Piersol, Ch. 4.2.2, 5.2.2, 8.5.4

Recap

Last time we looked at Parseval's theorem as a starting point for computing spectra, and we used white noise to compute spectra. We noted that the basic recipe for computing spectra is simple:

1. Subdivide data into segments
2. Compute $\text{fft}(\text{data})$
3. Compute squared amplitude for half of fft
4. Multiply by factor of 2
5. Normalize by $N = \text{length of record}$
6. Average multiple segments
7. Plot

But we only considered spectra for white noise. Now we need to consider some real world issues. What are uncertainties? What happens when there's a background trend? What about all of those frequencies that we don't resolve?

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 $24 \text{ h}/12.42 \text{ h}/1000 = 1.9324 \text{ cycles}/1000 \text{ pts} = 0.0019 \text{ cycles/measurement}$. For S2, we have $24 \text{ h}/12 \text{ h}/1000 = 2 \text{ cycles}/1000 \text{ pts} = 0.0020 \text{ cycles/measurement}$

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 $\Delta\sigma = 1/N$, then $N = 1/\Delta\sigma$. 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

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

Exercises

What are the Fourier transforms of the following functions: $x(t) = 1$, $x(t) = \cos(\omega t)$, $x(t) = \sin(2\omega t)$, $x(t) = \exp(-t^2/2\sigma^2)$, $x(t) = t$?

In our examples, note that the Fourier transforms of single frequency sine and cosine give a single peak. (This uses the Kronecker delta, δ_{nm} .) The Gaussian (f), has a transform of a Gaussian, though it's a bit distorted in this finite-length log-log domain. Formally if $f(t) = e^{-\pi t^2}$, then

$F(\sigma) = e^{-\sigma^2/(4\pi)}$ or with coefficients $f(t) = \alpha/\pi e^{-\alpha t^2}$ corresponds to $F(\sigma) = e^{-\sigma^2/(4\alpha^2)}$, which says in essence that the Fourier transform of a Gaussian is still a Gaussian. The normalization for this is dependent on our exact notation for the Fourier transform. And the linear pattern should really be thought of as a repeating sawtooth. It's Fourier transform will be a dramatically red spectrum. What does this mean for the Fourier transform of any long-term trend?

Spectral Uncertainties

The problem with all of the spectra that we've computed so far (the amplitude of the Fourier Transform) is that we have no way to evaluate the uncertainty. By eye we can see that it's fairly noisy. We know from computing means that the error of the mean decreases as we average more quantities together.

How do we incorporate more data into our spectra? You might imagine that you could improve your spectrum by extending the input time series from N to $2N$ data points for example. Unfortunately, although adding data points will change your spectrum, it won't reduce your noise or make the spectrum more precise at any individual frequency. Instead it will increase the number of frequencies for which you obtain results from $N/2$ to N .

Error bars for spectra rely on a principle similar to what we used when we estimated the standard error of the mean. Our uncertainties in our spectra decrease as we average more spectra together. The challenge is to figure out how to obtain more spectra that can be averaged together. Typically what we do is to break our time series into segments, compute spectra for each of the segments, and average these to get a mean spectrum. Since we're averaging squared quantities (the spectral amplitudes) the distribution of the sum will follow a χ^2 distribution. And we'll use this to estimate the uncertainties. When we compute our estimate of the spectrum of our data, we'll compute coefficients $a(\sigma)$, and for M segments, at frequency σ , the estimated spectrum will be

$$\hat{E}(\sigma) = \langle |a(\sigma)|^2 \rangle = \frac{1}{n} \sum_{i=1}^n |a(\sigma)|^2 \quad (2)$$

(Since we actually had two degrees of freedom for each segment, this will give us $\nu = 2n$ degrees of freedom.)

Now to estimate errors the whole argument stems from a notion that there is a true spectrum $E(\sigma)$, which we don't know, and our best estimate $\hat{E}(\sigma)$, that we've computed, but which might be noisy. We could consider the difference between E and \hat{E} , but this gets hairy quickly. Besides we're going to look at our results in log/log space anyway. So let's look at the ratio of the two. Recall that $\ln(\hat{E}/E) = \ln(\hat{E}) - \ln(E)$. We can define a variable that represents this sum of squares normalized by the true spectrum

$$y = \nu \frac{\hat{E}_\nu(\sigma)}{E(\sigma)} \quad (3)$$

where \hat{E}_ν is our best estimate of the spectrum from real observations, when we have ν degrees of freedom. This variable should have a χ^2 distribution with ν degrees of freedom. Thus we'll always consider the ratio $\hat{E}(\sigma)/E(\sigma)$, where we use $\nu/2$ data segments.

So even though we don't know E , we know that the ratio of $\nu\hat{E}$ to E ($\nu\hat{E}/E$) is a χ^2 quantity. The expectation value of this ratio is ν , and the standard deviation is $\langle y^2 \rangle - \langle y \rangle^2 = 2\nu$ (see Bendat and Piersol, section 4.2.2). From that we can infer that with two degrees of freedom, the standard deviation $\langle \hat{E}^2 \rangle - \langle \hat{E} \rangle^2 = 2/\nu E_0^2 = E_0^2$, which means that the standard deviation is equal to the original value. Clearly we need more samples.

Now in reality, we don't care about the details of χ^2 but rather the probability that our estimate of the spectrum \hat{E} is close to or far from the unknown true spectrum E , so we're going to look at the probability that $\nu\hat{E}(\sigma)/E(\sigma)$ falls within a fixed range. So we want to use the χ^2 distribution as a probability distribution and ask about the probability that $\hat{E}(\sigma)/E(\sigma)$ should be within \pm some range of 1. Formally, the probability that the estimated spectrum should be close in value to the true spectrum is:

$$P\left(\chi_{\nu,1-\alpha/2}^2 < \nu\frac{\hat{E}(\sigma)}{E(\sigma)} < \chi_{\nu,\alpha/2}^2\right) \quad (4)$$

so if we want to find a 95% significance level, we set α to 0.05. We'll want to know the value of χ^2 that corresponds to a given point on the cdf, so for this we use the inverse χ^2 function ("chi2inv" in Matlab). So we can invert this relationship to provide the probability that the true value is within a particular range of the observed value:

$$P\left(\frac{\chi_{\nu,1-\alpha/2}^2}{\nu} < \frac{\hat{E}(\sigma)}{E(\sigma)} < \frac{\chi_{\nu,\alpha/2}^2}{\nu}\right) \quad (5)$$

$$P\left(\frac{\nu}{\chi_{\nu,1-\alpha/2}^2} > \frac{E(\sigma)}{\hat{E}(\sigma)} > \frac{\nu}{\chi_{\nu,\alpha/2}^2}\right) \quad (6)$$

$$P\left(\frac{\nu\hat{E}(\sigma)}{\chi_{\nu,1-\alpha/2}^2} > E(\sigma) > \frac{\nu\hat{E}(\sigma)}{\chi_{\nu,\alpha/2}^2}\right) \quad (7)$$

$$P\left(\frac{\nu}{\chi_{\nu,\alpha/2}^2} < \frac{E(\sigma)}{\hat{E}(\sigma)} < \frac{\nu}{\chi_{\nu,1-\alpha/2}^2}\right) \quad (8)$$

Whether you use χ^2 or its reciprocal, the ratio between the high and low error bars should be the same, and the ratio will be what matters.

This error formulation differs from the usual error bars that we're used to seeing where we say for example that the true temperature should be the measured temperature plus or minus an uncertainty: $T = \hat{T} \pm \delta_T$. We can develop a similar expression for the true spectrum: $E(\sigma)$ is in the range between $\nu\hat{E}(\sigma)/\chi_{\nu,\alpha/2}^2$ and $\nu\hat{E}(\sigma)/\chi_{\nu,1-\alpha/2}^2$, where ν is twice the number of segments. This expression isn't very easy to interpret, since it varies as a function of frequency, and the estimated value $\hat{E}(\sigma)$ is not at the mid-point of the range.

Instead we'll keep in mind that we've computed the uncertainty for the ratio $\hat{E}(\sigma)/E(\sigma)$, and the probabilities for this ratio does not depend on frequency. On a log plot, error bars defined by the range between $\nu/\chi_{\nu,\alpha/2}^2$ and $\nu/\chi_{\nu,1-\alpha/2}^2$ are the same size at all frequencies, so we can easily compare spectral peaks at different frequencies.

Some statistics books include look-up tables for χ^2 , but we can compute it directly in Matlab. For $\nu = N/2$ data segments, the error limits are:

```
err_high = N/chi2inv(.05/2,N);
err_low = N/chi2inv(1-.05/2,N);
```

We can plot these values as:

```
semilogy([f f],[err_low err_high]*A);
```

where we set the frequency f and the amplitude A , so that the error bar ends up positioned in a convenient spot on the plot.

Now to have $\nu = N/2$ data segments, we have to split our long data record into shorter segments. We can do this by taking M data points at a time:

```
N=length(data);
M=segment_length; % define this value
for n=1:floor(N/M)
    d=data((n-1)*M+1:n*M); %select data for the nth segment
    fd(:,n)=fft(d);          % compute fft
end
sd=sum(abs(fd(1:M/2+1,:)).^2,2)/N; % sum over all spectra
                                   % (sum over 2nd index)
sd(2:end)=sd(2:end)*2;

nu=2*floor(N/M);
err_high = nu/chi2inv(.05/2,nu);
err_low  = nu/chi2inv(1-.05/2,nu);

semilogy(0:M/2,sd,[M/4 M/4],[err_low err_high]*sd(M/4))
```