

Lecture 10:*Reading: Bendat and Piersol, Ch. 5.2.1**Recap*

Last time we looked at the sinc function, windowing, and detrending with an eye to reducing edge effects in our spectra. We've got a full recipe, but in this case, there's more than one way to bake a cake. Is there another way to do this? First a quick diversion to Monte Carlo simulation....

Monte Carlo simulation: How to avoid the traps imposed by standard statistical assumptions (and how to fake your way as a statistician through computational inefficiency rather than clever mathematics)

Most of the time, we estimate spectral error bars using basic statistical assumptions—that data are normally distributed, that we have enough samples for the central limit theorem to apply, that statistics are stationary. These assumptions make our statistical models tractable—we end up with equations we can manipulate, allowing us (or clever statisticians 100 years ago) to derive simple equations that give us rules for how systems should behave. But what happens when those assumptions break down? Or what happens when we have little doubts about the validity of the statistical model. We can always resort to a Monte Carlo process. In Monte Carlo methods, we throw theory on its head and use an empirical approach to generate many realizations of our data set, with noise appropriate to our problem.

As an example, consider the problem of determining the standard error of the mean. When we discussed it in class, we did a quick derivation to show that the standard error of the mean is σ/\sqrt{N} , where σ is the standard deviation and N is the effective degrees of freedom. But what if I didn't trust this realization? I could generate a large number of realizations of my data with noise typical of the real data, compute means for each realization, and look at the statistics of those values.

So let's put this to work. Suppose I'm computing the mean of $N = 500$ data points. With one sample, I can compute the mean μ and standard deviation σ , and standard error $\sigma/\sqrt{500}$. But I might wonder if μ is really representative. So I can generate an ensemble of fake data, perhaps 100 data sets based on adding Gaussian white noise (or non-Gaussian white noise) to the real data. Each of these data sets will have a mean μ_i and a standard deviation σ_i . And I can look at the standard deviation of all of the μ_i values. I can also look at the pdf of my μ_i 's and other higher order statistics. For example:

```
A=randn(500,100);
mu=mean(A);
sigma_A=std(A);
std_A=sigma_A/sqrt(500);
[std(mu) mean(std_A)] % compare standard deviation of means
                    % vs standard error
```

Now we could expand on our example and ask, what if our noise were non-Gaussian or gappy or had other problems, and we could adjust our Monte Carlo process appropriately.

Getting the units right

So far we've been a tiny bit sloppy about the units in our spectra, but we do need to make this right. Our fundamental principle is that we want Parseval's theorem to work. But this gets a tiny bit messy when we average multiple frequencies. Still the basic rule of Parseval's theorem is not

that the sum of the squares equals the sum of the squared Fourier coefficients, but rather that the integrated variance equals the integral under the spectrum.

When Matthew lectured, his notes had the spectral estimate divided by N^2 and then divided by $df = 1/T$, so he was multiplying by T/N^2 . But my sample code so far has only had a division by N . Why?

In the discrete Fourier transform, we have:

$$\sum_{n=0}^{N-1} x_n^2 = \frac{1}{N} \sum_{k=0}^{N-1} |X_k|^2, \quad (1)$$

or after spectral normalizations, maybe we write

$$\sum_{n=0}^{N-1} x_n^2 = \frac{1}{N} \left[|X_0|^2 \Delta f + \sum_{k=1}^{(N/2-1)} 2|X_k|^2 \Delta f + |X_{N/2}|^2 df \right] \quad (2)$$

$$\approx \frac{1}{N} \sum_{k=0}^{(N-1)/2} 2|X_k|^2. \quad (3)$$

For this discussion, to keep the equations compact, we'll use the approximation in the final line, neglecting the fact that the mean and the $k = N/2$ value should not be doubled. Many times we work with this form, since it gives us meaningful spectral slopes. But in continuous form, we had a form more like this:

$$\int_{-\infty}^{\infty} x^2(t) dt = \int_{-\infty}^{\infty} |X(f)|^2 df \quad (4)$$

where f is frequency in cycles per unit time (or we might use $\sigma = 2\pi f$ sometimes, where σ is frequency in radians per unit time.) If we want this integral form to work for our real data, then we have to be a bit careful with our normalizations. We're going to want the area under the curve in our spectrum to be equal to the total variance integrated over time. So if total integrated variance is

$$variance = \sum_{n=0}^{N-1} x_n^2 \Delta t \quad (5)$$

where $\Delta t = T/N$. Then the integrated spectrum should be

$$variance = \frac{\alpha}{N} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f, \quad (6)$$

where $\Delta f = 1/(N\Delta t) = 1/T$, and we'll need to figure out α to ensure that the spectral estimator that we compute still properly adheres to Parseval's theorem. This implies that we might imagine normalizing our spectra to have:

$$\sum_{n=0}^{N-1} x_n^2 \Delta t = \frac{\alpha}{N} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f = \frac{\Delta t}{N\Delta f} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f = \frac{T^2}{N^2} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f \quad (7)$$

Or maybe this makes for units that aren't easily compared, so we could normalize our spectra to represent the average energy per unit time in the time domain, and adjust the frequency domain

accordingly:

$$\frac{1}{T} \sum_{n=0}^{N-1} x_n^2 \Delta t = \frac{1}{N \Delta t} \sum_{n=0}^{N-1} x_n^2 \Delta t \quad (8)$$

$$= \frac{T}{N^2} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f \quad (9)$$

$$= \frac{(\Delta t)}{N} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f \quad (10)$$

$$= \frac{1}{N^2 \Delta f} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f \quad (11)$$

$$= \frac{1}{N 2 f_{Nyquist}} \sum_{k=0}^{N/2-1} 2|X_k|^2 \Delta f \quad (12)$$

so we could divide our spectrum by twice the Nyquist frequency to have energy in units appropriate for comparing if we wanted to have our integrals match.

This isn't always the way we think about this, but it serves as our reminder that we should think about the units of our spectrum. What we know is that integral of our spectrum over a certain frequency range should give a measure of the signal variance:

$$\text{variance in a band} = \int_{f-\Delta f/2}^{f+\Delta f/2} |X(f)|^2 df \quad (13)$$

So if we expand this out, this implies that the units of $|X(f)|^2$ should be equivalent to variance divided by frequency, so it's our reminder that we'll label the y-axis units as the squared units of x divided by frequency, with a normalization to account for the units of time in our data.

Filtering in the frequency domain

Last time, when we talked about windowing, we noted that windowing in the time domain is equivalent to convolution in the frequency domain (and filtering in the time domain is equivalent to multiplication in the frequency domain.) This could lead you to an interesting conclusion. What if you skipped all the windowing and just did convolutions (i.e. filtering) in the frequency domain? In the limit in which you choose the same filter, these options should be the same.

This approach was originally developed by Daniell and is nicely discussed by von Storch and Zwiers (see their section 12.3.11). Daniell's original idea was to run a moving average over the Fourier transform of the full record. In this case the confidence intervals are determined by:

$$P \left(\chi_{\nu, 1-\alpha/2}^2 < \nu \frac{\hat{E}(f)}{E(f)} < \chi_{\nu, \alpha/2}^2 \right) \quad (14)$$

where ν in this case is $2 \times$ the number of frequencies averaged together.

The advantages of this approach are that it provides an unbiased estimate of the true spectrum. The width of our averaging forces us to tradeoff bias (minimized if we do less averaging) vs variance (minimized with more averaging). One virtue of averaging in the frequency domain is that we can apply different levels of averaging (with different error bars) depending on the frequency.

Using the auto-covariance to think about spectra.

Now let's look at spectra from a different perspective. When we talked about Parseval's theorem, we took a look at autocovariance. That was the convolution of $x(t)$ with its time reversal, $x(-t)$.

$$y(\tau) = \int_{-\infty}^{\infty} x(t)x(\tau + t)dt. \quad (15)$$

More formally, we might write this autocovariance as $R_{xx}(\tau)$.

$$R_{xx}(\tau) = \int_{-\infty}^{\infty} x(t)x(\tau + t)dt. \quad (16)$$

Now, what if we Fourier transform R ?

$$S_{xx}(f) = \int_{-\infty}^{\infty} R_{xx}(\tau)e^{-i2\pi f\tau} d\tau. \quad (17)$$

Formally, this and its inverse transform are the Wiener-Khinchine relations.

Now let's think about starting with two functions, $x(t)$ and $y(t)$. We can write their Fourier transforms:

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft} dt \quad (18)$$

$$Y(f) = \int_{-\infty}^{\infty} y(t)e^{-i2\pi ft} dt. \quad (19)$$

So now let's define X times the complex conjugate of Y . (Why do we consider the complex conjugate? Because it's how we always multiply vectors.) So

$$X(f)Y^*(f) = \int_{-\infty}^{\infty} k(t)e^{-i2\pi ft} dt. \quad (20)$$

For the moment, we have no idea what $k(t)$ should be, but we should be able to figure it out. If $X(f)Y^*(f)$ is a product in the frequency domain, then $k(t)$ should be a convolution in the time domain:

$$k(t) = \int_{-\infty}^{\infty} x(u)y(u - t) du. \quad (21)$$

We can plug $k(t)$ into our equation to check this.

$$\int_{-\infty}^{\infty} k(t)e^{-i2\pi ft} dt = \int_{-\infty}^{\infty} \left\{ \int_{-\infty}^{\infty} x(u)y(u - t) du \right\} e^{-i2\pi ft} dt \quad (22)$$

$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u)y(u - t)e^{i2\pi f(u-t)}e^{-i2\pi fu} dt du \quad (23)$$

$$= \int_{-\infty}^{\infty} x(u)e^{-i2\pi fu} \left\{ \int_{-\infty}^{\infty} y(u - t)e^{i2\pi f(u-t)} dt \right\} du \quad (24)$$

$$= \int_{-\infty}^{\infty} x(u)e^{-i2\pi fu}Y^*(f) du \quad (25)$$

$$= Y^*(f) \int_{-\infty}^{\infty} x(u)e^{-i2\pi fu} du \quad (26)$$

$$= Y^*(f)X(f). \quad (27)$$

So we can think about what happens when $x(t) = y(t)$, so that

$$k(t) = \int_{-\infty}^{\infty} x(u)x(u-t) du. \quad (28)$$

This means that $k(t)$ is the autocovariance of x . and

$$|X(f)|^2 = \int_{-\infty}^{\infty} k(t)e^{-i2\pi ft} dt. \quad (29)$$

This says that the Fourier transform coefficients squared (what we use when we compute spectra) are equivalent to the Fourier transform of the autocovariance.

Using the auto-covariance to compute spectra requires averaging, just as we did by segmenting our data and using the fft, but there's one tidy little trick. Let's use some white noise again, and take a look at our options:

1. Suppose we start with a big matrix of white noise, and we compute the autocovariance for each column of our matrix, then Fourier transform, and use these to compute a spectrum. We'll end up doing something along these lines:

```
A=randn(1000,100);
for i=1:100
    AcA(:,i)=xcov(A(:,i),A(:,i),'unbiased'); % autocovariance for
end
fAcA=fft(AcA(500:1500,:)); % Fourier transform of autocovariance
frequency=(0:500)/1000;
loglog(frequency,abs(mean(fAcA(1:501,:),2)),'LineWidth',3)
set(gca,'FontSize',16)
xlabel('Frequency (cycles per data point)','FontSize',16)
ylabel('Spectral energy','FontSize',16)
```

2. Alternatively, we could average all of the autocovariances, and then Fourier transform:

```
mean_AcA=mean(AcA,2);
fmean_AcA=fft(mean_AcA(500:1500));
hold on
loglog(frequency,abs(fmean_AcA(1:501,:))*1.1,'r','LineWidth',3)
legend('average of FFTs of many autocovariances',...
      'FFT of averaged autocovariance (scaled by 1.1)')
```

3. For comparison, the periodogram-based determined from the fft of the data:

```
fA=fft(A);
ampA=abs(fA(1:501,:)).^2/1000; ampA(2:500,:)=2*ampA(2:500,:);
loglog(frequency,mean(ampA,2),'LineWidth',3)
```

In the results, shown in Figure 1, the curves are identical, though the red line has been scaled up by 10% to make both visible. There are some normalizations here that we haven't properly confronted. Notably we are missing factor of 2 for the autocovariance cases. As with the periodogram, we're

only plotting half the spectrum, so we need double the energy. If we were dealing with real data, we'd also need a factor of N or Δt to properly normalize our fft. Details can be sorted out later, and Thomson and Emery provide a bit of guidance on this.

You shouldn't be surprised that averaging before or after the FFT leads to the same results, since averaging has no impact on the FFT. But this might give you an idea of how you can take advantage of the autocovariance to compute spectra from gappy data.

All of this means that we could compute spectra without needing to chunk our data and compute lots of ffts, provided that we had a good estimate of the autocovariance. In the days before the development of the FFT, the autocovariance was a natural pathway to determining the spectrum, since it was clean and easy to compute. And now, with modern computing, you might not feel like there's any need to take advantage of the FFT anymore. If you can obtain the best possible estimate of the autocovariance, by whatever means necessary, then you should be able to compute one FFT and obtain reasonable estimate of the spectrum, without concern for data gaps or computational speed.

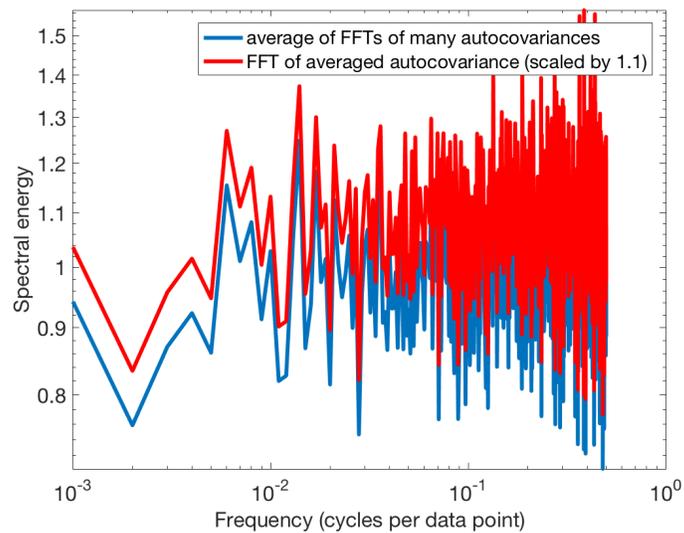


Figure 1: Spectra for white noise, computed by Fourier transforming 100 realizations of the autocovariance function (blue), or by Fourier transforming a smoothed autocovariance function computed from 100 realizations of the data (red). The red line is scaled upward by a factor of 1.1.