Lecture 2

FOURIER TRANSFORMS

The Heart of the FFT

The FFT is a great tool for producing numbers but -- and you probably heard this before -- a major purpose of doing computations is to get insight into a problem, not numbers. So -- once you have a FFT, what does it mean? What understanding does it offer?

Let's take a very simple example: a heartbeat, as heard through a stethoscope, beat.wav. It's been sampled at 8012Hz. Import it into Matlab, take the FFT, and plot the absolute value. While you're at it, why not plot the signal?

If you look at the signal, you see it has length 32465. Sampled at 8012Hz, it represents 32465/8012 = 4.0520469 seconds of sound. Within those four seconds, you have five full heart beats (which consist of two major sounds, hence ten peaks). So, the patient has a heart rate of some 74 beats per minute -- well within normal ranges.

However, this is about FFT, so we prefer to ask, what frequency is the whole beat? Of course, it's 5/4.05= 1.23 beats per second. If we're more generous, and consider ten cycles instead of five, we see the beats are at a frequency of 2.46 bps. Here's the FFT: where's the peak at 2.46?

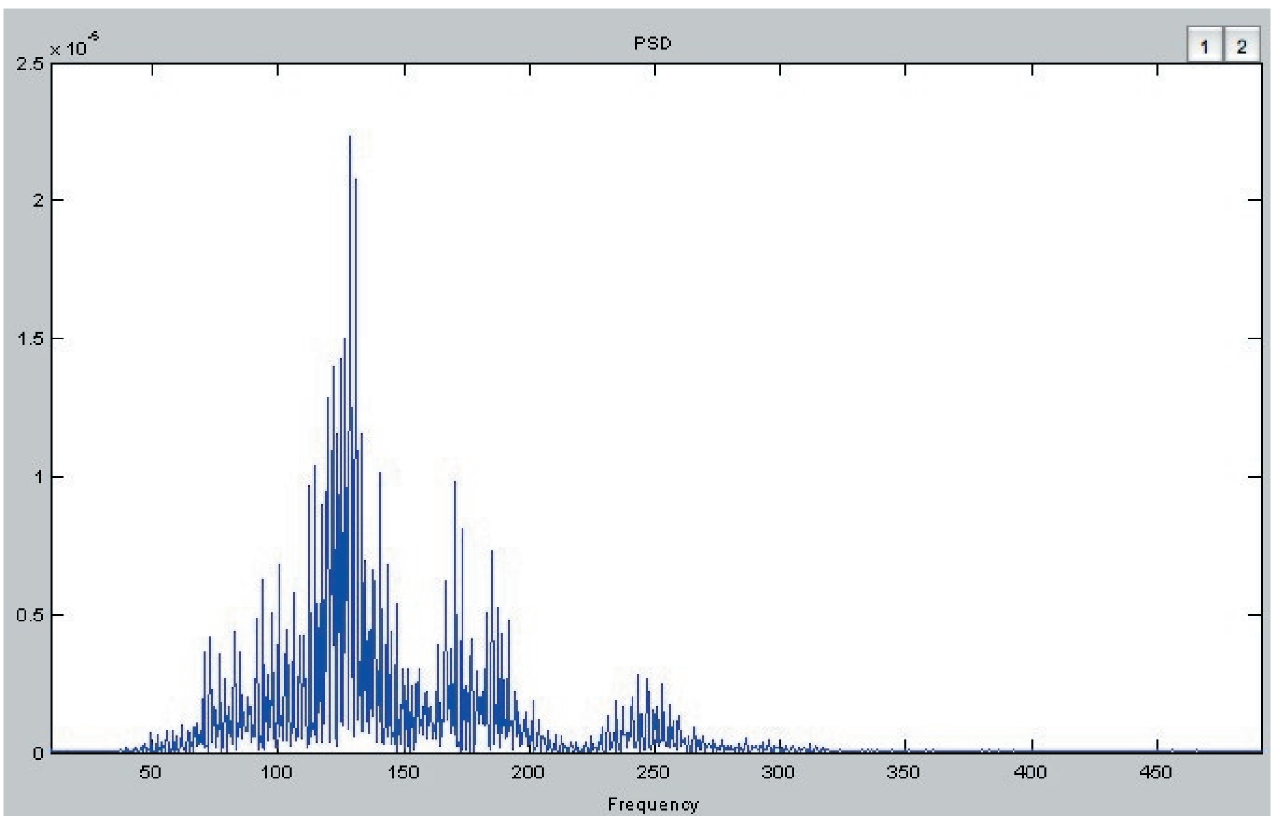

Ummm -- there isn't any. There's a grand old peak at around 125Hz, and a smaller one at 175Hz, but nothing for the main beat. What gives?

Let’s look closer at a beat. Here’s the two peaks in one beat, beat1.wav and beat2.wav.

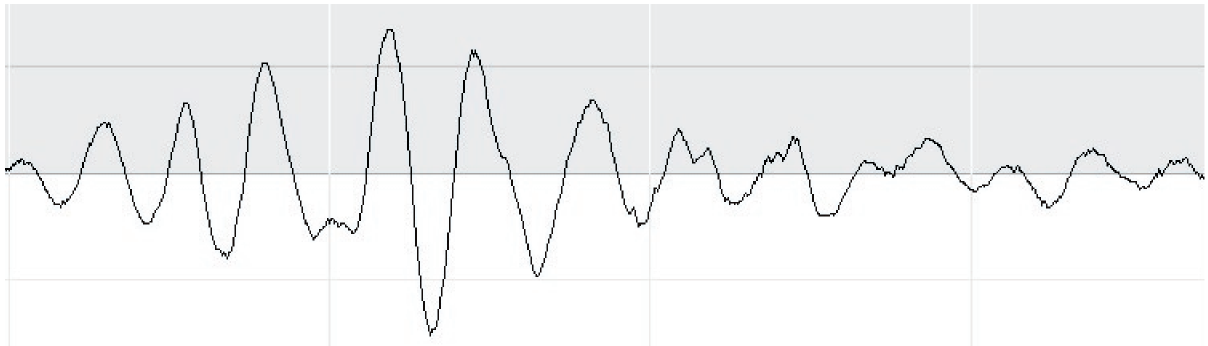

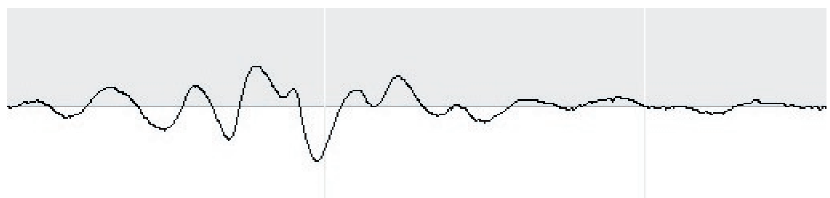

Neither is a pure tone, what they remind me of most is a modulated signal:

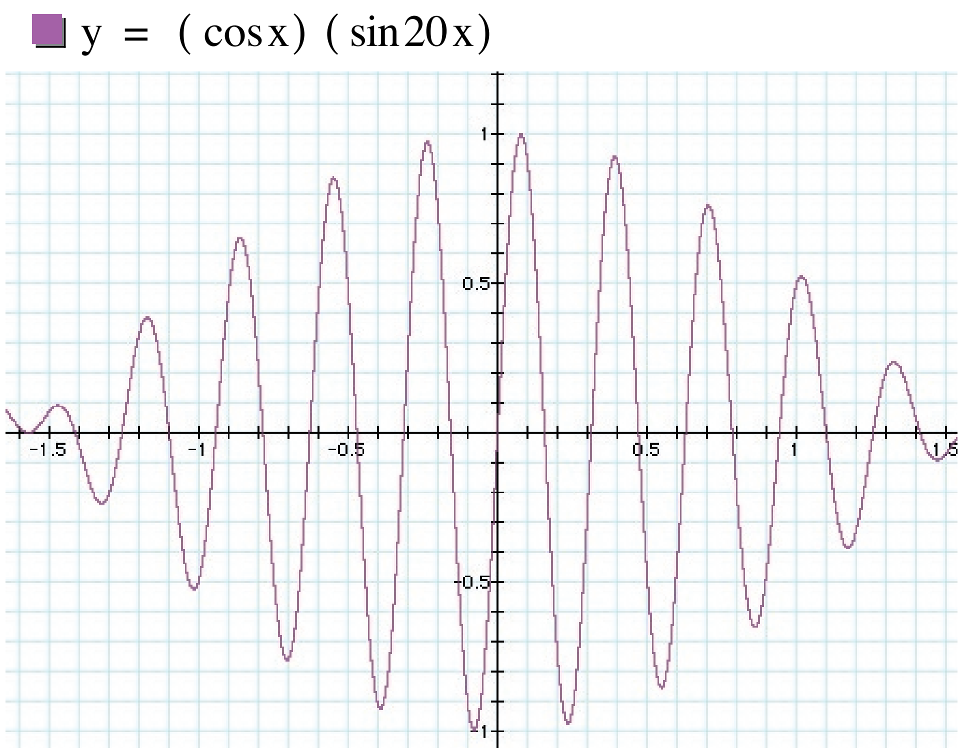

If I tried to pick out the modulating cosine in the above graph, I’d integrate y against the cosine. But the signal y oscillates up and down so quickly, the various positive and negative parts of the integral cancel, and I get no contribution from the cosine at all. The only frequencies that could possibly contribute to the FFT would be those close to sin(20x). Or, if we want to be rigorous instead of intuitinve, using trig identities:

$$\cos(x)\sin(20x) = \frac{1}{2}\left[\sin(20+1)x + \sin(20-1)x\right]$$

The modulated sine wave has frequencies close to the original. This is why the heart beat appears to have no frequencies at 1.23 or 2.46Hz. Those are the modulating frequencies of some much higher sine.

Let's try using that approach. We can find the frequency of the sine that's being modulated in beat1 by counting the number of cycles in beat1: half the number of zero crossings. Beat1 has 14 cycles in 772 points, which gives me a sine of frequency 145.29Hz. Let's confirm that by taking the FFT of beat1:

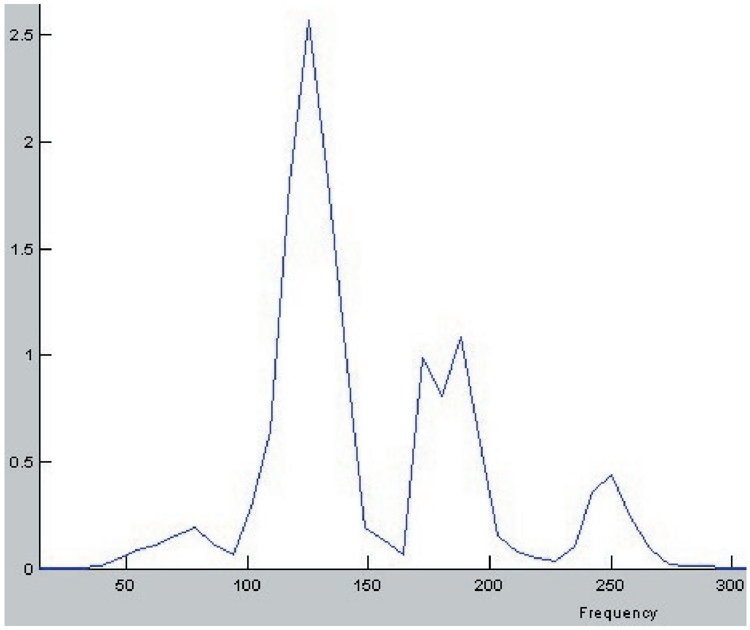

Note the major peak between 100 and 150Hz; our computations aren't too far off. Similarly, beat2 has 9 cycles in 491 points, giving a frequency of 73.42Hz. And here's the FFT of beat2:

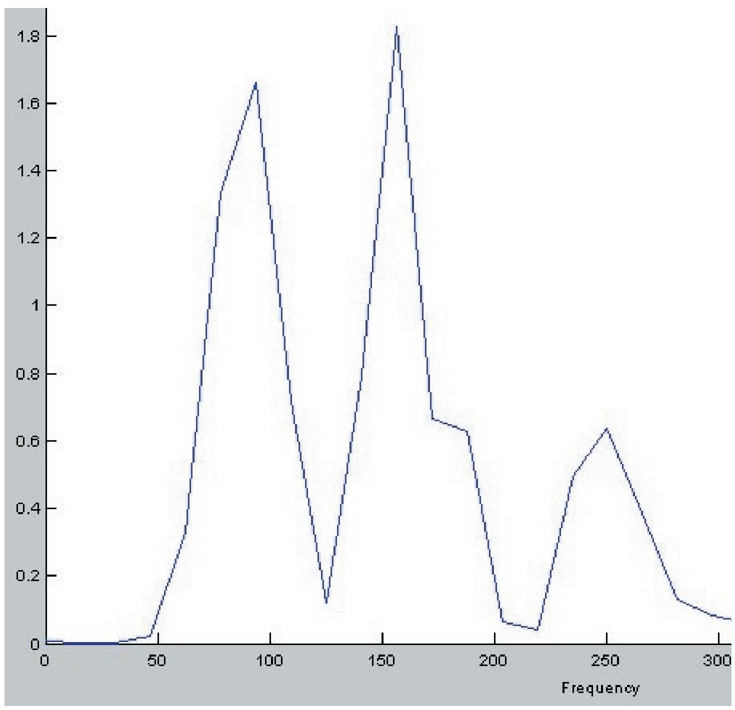

Again, a peak about where we'd expect it. Although let's not get too thrilled with ourselves; there are other peaks as well, suggesting we haven't got the full story of this FFT. But it's the kind of start that tells us we are beginning to understand what's going on in this signal.

Is there no hope, then, of recovering the modulating frequency? we'd have to find a way to eliminate the cancellations. And the simplest way to do that is to take the absolute value of the signal. Try it -- take the absolute value of beat.wav, then FFT, and inspect near 1Hz. What do you see?