

Spectral Analysis in Matlab

Introduction

This lab covers the following topics. Corresponding exercises and questions to be answered are contained in the sections.

1. **The Spectral Estimation Problem**
2. **Discrete Fourier Transform**
3. **Power Spectral Density (PSD) Estimation**
4. **Spectral Methods - Summary & Exercises**
5. **Time-varying Spectra**
6. **Wavelets**

1 The Spectral Estimation Problem

The spectrum of a signal is fundamentally a description of the energy of a signal as a function of frequency [1], [2]. It may also be a function of time and frequency which leads to the time-varying or the so-called time-frequency spectrum description. Spectral analysis, which has a long history dating back to the late 19th century, refers to the various methods deployed to determine spectra. These include non-parametric methods such as fourier and wavelet analysis and parametric ones that entail fitting data with finite rational filter models (with filter coefficients constituting the parameters). Furthermore, signals may be deterministic or random processes. In both cases, they may have time-independent or time-dependent spectral content. Even so, the spectra can still be described by non-parametric or parametric time-frequency methods. Spectra are represented by their amplitude and phase or more meaningfully, by the squared magnitude, which is referred to as the Power Spectrum. Thus, any process that quantifies the various amounts of different frequencies in a signal, qualifies as a spectral analysis method.

It should be noted at the outset that in the spectral estimation problem (of deterministic or random data), the practical constraint is that one only has a finite data record, $x(0), x(1), \dots, x(N-1)$. It is from this record that one has to estimate the true spectrum, the so-called power spectral density. This is the central issue in spectral estimation. Size N is a critical quantity. It may be that we wish to look at short data records to maintain “stationarity” or that only short data records are available. Using short data records has its problems. We will see that it affects the so-called bias-variance trade off in spectral estimators. This is discussed in the following sections.

A signal is (wide-sense) *stationary* if its distribution at any n is independent of time shifts, i.e., if $x(n)$ and $x(n + N)$ have the same distribution for every integer N . Stationary signals are associated with steady states in the system producing the signal, and are characterized by constant averages and variances along the signal. When statistics computed along any one realization of a signal are identical to those computed across ensembles of different realizations of the signal, the system is *ergodic*. Ergodicity is assumed in most practical situations.

Matlab has functions for computing the mean (*mean*), variance (*var*), standard deviation (*std*), covariance (*cov*), and correlation coefficient (*corrcoeff*) of signal values sampled across time or across ensembles. The Statistics Toolbox provides many additional statistical functions.

The Signal Processing Toolbox provides a *crosscorrelation* function (*xcorr*) and a *crosscovariance* function (*xcov*). Crosscorrelation of two signals is equivalent to the convolution of the two signals with one of them reversed in time. The *xcorr* function adds several options to the standard Matlab *conv* function. The *xcov* function simply subtracts the mean of the inputs before computing the crosscorrelation. The crosscorrelation of a signal with itself is called its *autocorrelation*. It measures the *coherence* of a signal with respect to time shifts.

Try:

```
x = randn(1,100);
w = 10;
y = conv(ones(1,w)/w,x);
avgs = y(10:99);
plot(avgs)
```

Ensemble averages:

```
w = 10;
for i = 1:w;
X(i,:) = randn(1,100);
end
AVGS = mean(X);
plot(AVGS)
```

Q1. What principle is being illustrated in the above two examples?

```
x = [-1 0 1];
y = [0 1 2];
xcorr(x,y)
conv(x,fliplr(y))
```

```
xcov(x,y)
xcov(x,x)
xcov(x,y-1)
```

Q2. Explain results in each of the command above exercise.

Example: Crosscorrelation

In a simple target ranging system, an outgoing signal x is compared with a returning signal y . We can model y by

$$y(n) = \alpha x(n - d) + \beta$$

where α is an attenuation factor, d is a time delay, and β is channel noise. If T is the return time for the signal, then x and y should be correlated at $n = T$. The target will be located at a distance of vT , where v is the channel speed of the signal.

Try:

```
x = [zeros(1,25),1,zeros(1,25)];
subplot(311), stem(x)
y = 0.75*[zeros(1,20),x] + 0.1*randn(1,71);
subplot(312), stem(y)
[c lags] = xcorr(x,y);
subplot(313), stem(lags,c)
```

Q3. Does this example show the expected behavior?
Why or why not?

2 Discrete Fourier Transform (DFT)

The computational basis of *classical* spectral analysis is the *Discrete Fourier Transform* (DFT). The DFT of a vector y of length N is another vector Y of length N :

$$Y(k) = \sum_{n=0}^{N-1} W^{nk} y(n), \quad k = 0, 1, 2, \dots, (N - 1) \quad (1)$$

where W is a complex n^{th} root of unity:

$$W = e^{-2\pi j/N}$$

The graphical user interface `fftgui` allows you to explore properties of the DFT. If y is a vector,

```
fftgui(y)
```

plots `real(y)`, `imag(y)`, `real(fft(y))`, and `imag(fft(y))`. You can use the mouse to move any of the points in any of the plots, and the points in the other plots respond.

(You need ML01 for some of the functions below).

Try:

Roots of unity

```
edit z1roots
```

```
z1roots(3);
```

```
z1roots(7);
```

Q4. Express the above roots in exponential form.

Explore the DFT:

```
delta1 = [1 zeros(1,11)];
```

```
fftgui(delta1)
```

Q5. Confirm the DFT from the formula.

```
delta2 = [0 1 zeros(1,10)];
```

```
fftgui(delta2)
```

Q6. Confirm the DFT any way you can.

```
deltaNyq = [zeros(1,6),1,zeros(1,5)];
```

```
fftgui(deltaNyq)
```

```
square = [zeros(1,4),ones(1,4),zeros(1,4)];
```

```
fftgui(square)
```

```
t = linspace(0,1,50);
```

```
periodic = sin(2*pi*t);
```

```
fftgui(periodic)
```

2.1 Fast Fourier Transform (FFT)

The Matlab function `fft`, called by `fftgui`, uses the *fast* Fourier transform algorithm to compute the DFT.

DFTs with a million points are common in applications. For modern signal and image

processing applications, and many other applications of the DFT, the key is the ability to do such computations rapidly. Direct application of the definition of the DFT requires N multiplications and N additions for each of the N coefficients - a total of $2N^2$ floating-point operations. This number does not include the generation of the powers of W . To do a million-point DFT, a computer capable of doing one multiplication and addition every microsecond would require a million seconds, or about 11.5 days.

Modern FFT algorithms have computational complexity $O(N \log_2 N)$ instead of $O(N^2)$. If N is a power of 2, a one-dimensional FFT of length N requires less than $3N \log_2 N$ floating-point operations. For $N = 2^{20}$, that's a factor of almost 35,000 times faster than $2N^2$.

When using the FFT, note that there is a distinction between a *window length* and an *FFT length*. The window length is the length of the input. It might be determined by, say, the size of an external buffer. The FFT length is the length of the output, the computed DFT. The command

```
Y=fft(y)
```

returns the DFT Y of y . Here, the window length `length(y)` and the FFT length `length(Y)` are the same.

The command

```
Y=fft(y,n)
```

returns the DFT Y with length n . If the length of y is less than n , y is padded with trailing zeros to length n . If the length of y is greater than n , the sequence y is truncated. The FFT length is then the same as the padded/truncated version of the input y .

Try:

Vector data interpolation and the origins of the FFT

```
edit fftinterp
```

```
fftinterp
```

```
hold off
```

2.2 Spectral Analysis with the FFT

The FFT allows you to estimate efficiently the component frequencies in data from a discrete set of values sampled at a fixed rate. The following list shows the basic relationships among the various quantities involved in any spectral analysis. References to time can be replaced by references to space.

y	Sampled data
$N = \text{length}(y)$	Number of samples

<code>Fs</code>	Samples/unit time
<code>Ts = 1/Fs</code>	Time increment
<code>t = (0:N-1)/Fs</code>	Time range
<code>Y = fft(y)</code>	Discrete Fourier Transform (DFT)
<code>abs(Y)</code>	Amplitude of the DFT
<code>abs(Y).^2/N</code>	Power of the DFT
<code>Fs/n</code>	Frequency increment
<code>f = (0:N-1)*(Fs/n)</code>	Frequency range
<code>Fs/2</code>	Nyquist frequency

A plot of the power spectrum (fourth from the bottom and equation (6) below) is called a *periodogram*. The first half of the principal frequency range (from 0 to the Nyquist frequency $Fs/2$) is sufficient, because the second half is a complex conjugate reflection of the first half.

Spectra are sometimes plotted with a principal frequency range from $-Fs/2$ to $Fs/2$. Matlab provides the function `fftshift` to rearrange the outputs of `fft` and convert to a 0-centered spectrum.

Try:

Periodograms

```
edit pgrams
pgrams
```

Whale call

```
edit whalefft
whalefft
```

FFT Demos

```
sigdemo1
playshow fftdemo
phone
playshow sunspots
```

(Make sure you activate the sound box in the Phone pad).

2.3 Aliasing

Discrete time signals sampled from time-periodic analog signals need not be time periodic, but they are always periodic in frequency, with a period equal to the sampling frequency. The resulting *harmonics* show up as spectral copies in frequency domain plots like the pe-

riodogram. If the sampling rate is too low, these spectral copies can overlap within the principal range leading to aliasing error.

3 Power Spectral Density (PSD) Estimation

This section gives a bit more technical description of spectral analysis. For a deterministic signal, there are various methods for determining its spectrum. These are described in the following sections. Along with the spectrum, another important property of deterministic signals $x(n), 0 \leq n \leq N - 1$ is the autocorrelation function (ACF) that is correlation with itself. It is defined as,

$$R_{xx}(\tau) = \sum_{n=0}^{N-\tau} x(n)x(n + \tau)$$

When the signal that we have is a random process, then the matter of determining spectra becomes a bit more subtle. Recall that the ACF of a discrete-time random process $X(n, \xi)$ is the *deterministic* function,

$$R_{XX}(m) = E\{x(n)x(n + m)\} \quad (2)$$

where $x(n)$ is a realization of the random process $X(n, \xi)$. Hence we have,

$$R_{XX}(0) = E\{x(n)x(n)\}$$

where E is the expectation operator. Hence the right hand side is the *expected* i.e. *average* value of the square of the sequence $x(n)$. This leads to the characterization of $R_{XX}(0)$ as being the *average power* of the random process. (In the characterization above, we have of course assumed that we have a real, stationary random process.)

Now, since $R_{XX}(m)$ is a deterministic sequence, we can take its Fourier transform (actually the Discrete-time Fourier Transform, DTFT) which is defined as,

$$S_{XX}(\omega) = \sum_{m=-\infty}^{+\infty} R_{XX}(m)e^{-j\omega m}. \quad (3)$$

Then $R_{XX}(m)$ is the inverse DTFT of $S_{XX}(\omega)$ which is given by,

$$R_{XX}(m) = \frac{1}{2\pi} \int_{-\pi}^{+\pi} S_{XX}(\omega)e^{j\omega m} d\omega.$$

This leads to,

$$R_{XX}(0) = \frac{1}{2\pi} \int_{-\pi}^{+\pi} S_{XX}(\omega) d\omega$$

which leads to the interpretation that $S_{XX}(\omega)$ is the power spectral density (PSD) of the random process since $R_{XX}(0)$ is average power. Hence the PSD of a random process describes the distribution of its power with frequency.

We note from equation(2) that the PSD depends on an infinite number of observations of $R_{XX}(m)$. Hence determination of PSD is an impossible task. Accordingly, we try to determine an *estimate* of the PSD. Furthermore, $R_{XX}(m)$ in equation(1) is defined as the expectation of $x(n)x(n+m)$, that is, an *ensemble average*. Not only does this require an ensemble of realizations but also a knowledge of the joint pdf, both of which are generally impossible tasks. Accordingly, we aim to form an estimate $R_{XX}(m)$ to permit an estimation of $S_{XX}(m)$.

Typically, we only have *one* realization of a random process and that too for a finite duration. Hence the goal is to form an estimate of the ACF from the single finite duration realization of the random process. Under certain conditions [3] the *temporal* ACF can be shown to converge to the true ensemble ACF. That is,

$$\text{Lim}_{M \rightarrow \infty} \frac{1}{2M+1} \sum_{n=-2M}^{2M} x(n)x(n+k) = E(x(n)x(n+k)) = R_{XX}(k) \quad (4)$$

Random processes for which this holds are said to be ergodic, or more accurately second-order ergodic. Proving ergodicity is not an easy task, and more often than not, it is just assumed.

3.1 Alternative definition of PSD

It can be shown [3] that a nearly equivalent definition of the PSD, equation 3 is,

$$S_{XX}(\omega) = \text{Lim}_{M \rightarrow +\infty} E\left\{ \frac{1}{2M+1} \left| \sum_{n=-M}^M x(n)e^{-j\omega n} \right|^2 \right\} \quad (5)$$

This says that you first take the magnitude squared of the FT (DTFT) of the data $\{x(n)\}$ and then divide by the data record length. Then, since the Fourier Transform will be a random variable (each new realization $x(n)$ will be different), the expected value is taken. Finally, since the random process is generally of infinite duration, the limit is taken. Equality of equations 3 and 5 is established in [3], (p.59).

3.2 Classical Techniques of Spectral Estimation

The two most common classical spectral (that is, power spectral density) estimation techniques are the *Periodogram* and *Blackman-Tukey* methods. These are based on Fourier analysis (a non-parametric method) as opposed to so-called modern spectral estimation methods

which are parametric methods. The latter are discussed in the following sections. The fundamental distinguishing feature of the classical method is the *bias-variance* trade off. Bias occurs when the mean value of an estimator is different to the true value of the variable being estimated. Variance is the expected square deviation from its mean value. A brief description of the two classical methods is given here. Detailed description may be found in many texts, such as [3].

3.2.1 Periodogram

The periodogram spectral estimator stems from the PSD already defined in equation(4),

$$S_{XX}(\omega) = \lim_{M \rightarrow +\infty} E \left\{ \frac{1}{2M+1} \left| \sum_{n=-M}^M x(n) e^{-j\omega n} \right|^2 \right\}.$$

Removing the expectation operator and using only a finite data size, the periodogram spectral estimator is defined as,

$$\hat{S}_{XX}(\omega) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x(n) e^{-j\omega n} \right|^2 \quad (6)$$

which of course is the DTFT squared and averaged. The immediate question is whether the periodogram estimator is a *consistent estimator* of the PSD, that is: Does $\hat{S}_{XX}(\omega) \rightarrow S_{XX}(\omega)$, $N \rightarrow \infty$? The fact is that it is not [3], p.66. However, for $N \rightarrow \infty$, the estimator is unbiased although the variance does not approach zero.

3.2.2 Averaged Periodogram

The variance in the periodogram estimator of equation(not decreasing with increasing data record length) may be due to the removal of the expectation operator, that is, a lack of averaging. To improve the statistical properties of the the periodogram, the expectation operator can be approximated by averaging a set of periodograms. Assume that K independent records of the random process are available for the interval $0 \leq n \leq L - 1$. Then the averaged periodogram estimator is defined as,

$$\hat{S}_{AVPER}(\omega) = \frac{1}{K} \sum_{m=0}^{K-1} \hat{S}_{PER}^{(m)}(\omega)$$

where $\hat{S}_{PER}^{(m)}(\omega)$ is the periodogram for the m th data set,

$$\hat{S}_{PER}^{(m)}(\omega) = \frac{1}{L} \left| \sum_{n=0}^{L-1} x_m(n) e^{-j\omega n} \right|^2 \quad (7)$$

This results in a reduction of the variance by a factor of K (recall that we have K independent data sets). In practice, we seldom have independent data sets but rather, only one data record of length N . And we need to base the estimator on this data set. A standard approach is to segment the data into K nonoverlapping blocks of length L , where $N = KL$. Accordingly, for equation (4), one has data sets,

$$x_m(n) = x(n + mL), \quad n = 0, 1, \dots, L - 1; \quad m = 0, 1, \dots, K - 1$$

The associated problem here might be that, given the blocks are contiguous, they may not be uncorrelated let alone independent. (Only white noise would have that property). Accordingly, the variance reduction would be less than that for independent records.

3.2.3 Welch method

The Welch method provides a variation of the averaged periodogram. It differs in two ways: First, the data is windowed and secondly, data blocks are overlapped. The data window reduces spectral leakage (we see this later) and by overlapping blocks of data, typically by 50 or 75%, some extra variance reduction is achieved.

3.3 Blackman-Tukey Estimation

The periodogram equation (6) is in general a poor estimator of the PSD. This may be seen as follows: The equivalent form of equation (6) can be shown to be,

$$\hat{S}_{PER}(\omega) = \sum_{k=-(N-1)}^{N-1} \hat{r}_{xx}(k) e^{-j\omega k} \quad (8)$$

where,

$$\hat{r}_{xx}(k) = \begin{cases} \frac{1}{N} \sum_{n=0}^{N-1-k} x(n)x(n+k), & k = 0, 1, \dots, N-1 \\ \hat{r}_{xx}(-k), & k = -(N-1), -(N-2), \dots, -1 \end{cases} \quad (9)$$

The periodogram given by equation (8) depends on the autocorrelation function $r_{xx}(k)$ which is *estimated* by $\hat{r}_{xx}(k)$ of equation (9). Accordingly it is not surprising to obtain a poor estimator of the PSD, given that it depends on an estimate of the autocorrelation function $r_{xx}(k)$, which estimate itself may be poor. The Blackman-Tukey (BT) spectral estimator

improves on the performance of the periodogram by weighing the estimate $\hat{r}_{xx}(k)$ less at higher lags (where it performs poorly). Accordingly, the BT spectral estimator is defined as,

$$\hat{P}_{BT}(\omega) = \sum_{k=-(N-1)}^{N-1} w(k)\hat{r}_{xx}(k)e^{-j\omega k} \quad (10)$$

where the window $w(k)$ weighs the higher lags less (see [3], p.80). Note that it is equal to the periodogram if $w(k) = 1$, for all k . Again, bias-variance trade-off can be observed here depending on $w(k)$.

4 Spectral Methods - Summary & Exercises

PSD estimates of noisy analog signals from a finite number of its samples are based on three fundamentally different approaches:

- **Non-parametric methods**

Make no assumptions about the data in the sample and work directly with the DFT.

Welch: `pwelch`

Multitaper: `pmtm`

- **Parametric methods**

Model the data in the sample as the output of a linear system excited by white noise (noise with zero mean and constant PSD), estimate the filter coefficients, and use these to estimate the PSD.

Burg: `pburg`

Yule-Walker: `pyulear`

- **Subspace methods**

Based on an eigenanalysis or eigendecomposition of the correlation matrix associated with the data in the sample.

EV: `peig`

MUSIC: `pmusic`

Try:

```
Fs = 100;
```

```
t = 0:1/Fs:10;
```

```
y = sin(2*pi*15*t) + sin(2*pi*30*t);
```

```
nfft = 512;
```

```

Y = fft(y,nfft);
f = Fs*(0:nfft-1)/nfft;
Power = Y.*conj(Y)/nfft;
plot(f,Power)
title('Periodogram')

```

Q7. Explain how the PSD (periodogram) is derived, as described above.

```

figure
ryy = xcorr(y,y);
Ryy = fft(ryy,512);
plot(f, abs(Ryy))
title('DFT of Autocorrelation')

```

Q8. Which equations in this Lab. description do the two PSD estimates (Q7 and Q8) correspond to? Any thoughts as to why the scales are so off?

4.1 Non-parametric Methods

As described earlier, non-parametric methods estimate the PSD directly from the signal itself. The simplest such method is the periodogram. An improved version of the periodogram is Welch's method. A more modern technique is the multitaper method (MTM).

The following functions estimate the PSD P_{xx} in units of power per radians per sample. The corresponding vector of frequencies w is computed in radians per sample, and has the same length as P_{xx} .

- **Periodogram method**

```
[Pxx w] = periodogram(x)
```

Estimates the PSD using a periodogram. Optional inputs specify windows (default is rectangular), FFT length, PSD sample frequencies, and output frequency range.

- **Welch method**

```
[Pxx w] = pwelch(x)
```

Estimates the PSD using Welch's averaged periodogram method. The vector x is segmented into equal-length sections with overlap. Trailing entries not included in the final segment are discarded. Each segment is windowed. Optional inputs specify windows (default is Hamming), overlap, FFT length, and PSD sample frequencies.

- **Multitaper method**

```
[Pxx w] = pmtm(x, nw)
```

Estimates the PSD using a sequence of $2*nw - 1$ orthogonal tapers (windows in the frequency domain). The quantity nw is the time-bandwidth product for the discrete prolate spheroidal

sequences specifying the tapers. Optional inputs specify taper frequencies, FFT length, and PSD sample frequencies.

Try:

```
t = 0:1/100:10-1/100;
x = sin(2*pi*15*t) + sin(2*pi*30*t);
periodogram(x, [], 512, 100);
figure
pwelch(x, [], 512, 100);
figure
pmtm(x, [], 512, 100);
```

The above three estimator examples are only for your observation. No comments are necessary. Only the observations that the DFT can be applied in many ways in spectral estimation.

4.2 Parametric Methods

Parametric methods can yield higher resolution than non-parametric methods in cases where the signal length is short. These methods use a different approach to spectral estimation: instead of estimating the PSD directly from the data, they model the data as the output of a linear system driven by white noise (an adaptive filter), and then attempt to estimate the parameters of that linear system.

The most commonly used linear system model is the all-pole model, a system with all of its zeros at the origin in the z -plane. The output of such a system for white noise input is an autoregressive (AR) process. These methods are sometimes referred to as *AR methods*.

AR methods give accurate spectra for data that is “peaky,” that is, data with a large PSD at certain frequencies. The data in many practical applications (such as speech) tends to have peaky spectra, so that AR models are used there. In addition, the AR models lead to a system of linear equations that is relatively simple to solve for the parameters of the unknown system.

The following methods are summarized on the next page. The input p specifies the order of the autoregressive (AR) prediction model.

- **Yule-Walker AR method**

[Pxx f] = pyulear(x,p,nfft,fs)

- **Burg method**

[Pxx f] = pburg(x,p,nfft,fs)

- **Covariance and modified covariance methods**

```
[Pxx f] = pcov(x,p,nfft,fs)
[Pxx f] = pmcov(x,p,nfft,fs)
```

Example: Here is an example of using the AR parametric method to estimate the PSD of the signal.

Because the Yule-walker method estimates the spectral density by fitting an AR prediction model of a given order to the signal, first generate a signal from an AR (all-pole) model of a given order. You can use `freqz` to check the magnitude of the frequency response of your AR filter. This will give you an idea of what to expect when you estimate the PSD using `pyulear`:

```
% AR filter coefficients
a = [1 -2.2137 2.9403 -2.1697 0.9606];
```

```
% AR filter frequency response
freqz(1,a)
title('AR System Frequency Response')
```

Now generate the input signal `x` by filtering white noise through the AR filter. Estimate the PSD of `x` based on a fourth-order AR prediction model, since in this case, we know that the original AR system model `a` has order 4.

```
randn('state',1);
x = filter(1,a,randn(256,1)); % AR system output
pyulear(x,4) % Fourth-order estimate
```

The spectral estimate returned by `pyulear` is the squared magnitude of the frequency response of this AR model.

Q9. Explain, very briefly, the AR method of spectral estimation

The Section 4 exercises below are only for your review. No questions need to be answered. Go to **Section 5** for **Q10**.

Try:

```
edit pmethods
pmethods('pyulear',25,1024)
pmethods('pburg',25,1024)
pmethods('pcov',5,512)
pmethods('pmcov',5,512)
```

Some of the factors to consider when choosing among parametric methods are summarized in the following table. See the documentation for further details.

Burg	Covariance	Modified Covariance	Yule-Walker
Characteristics			
Does not apply window to data	Does not apply window to data	Does not apply window to data	Applies window to data
Minimizes forward and backward prediction errors	Minimizes forward prediction error	Minimizes forward and backward prediction errors	Minimizes forward prediction error
Advantages			
High resolution for short data records	Better resolution than Y-W for short data records	High resolution for short data records	As good as other methods for large data records
Always produces a stable model	Extract frequencies from mix of p or more sinusoids	Extract frequencies from mix of p or more sinusoids	Always produces a stable model
		No spectral line-splitting	
Disadvantages			
Peak locations highly dependent on initial phase	Can produce unstable models	Can produce unstable models	Performs relatively poorly for short data records
Spectral line-splitting for sinusoids in noise, or when order is very large	Frequency bias for estimates of sinusoids in noise	Peak locations slightly dependent on initial phase	Frequency bias for estimates of sinusoids in noise
Frequency bias for sinusoids in noise		Minor frequency bias for sinusoids in noise	
Conditions for Nonsingularity			
	p must be $\leq 1/2$ input frame size	p must be $\leq 2/3$ input frame size	Autocorrelation matrix always positive-definite, nonsingular

Subspace Methods

Subspace methods, also known as *high-resolution methods* or *super-resolution methods*, generate PSD estimates based on an eigenanalysis or eigendecomposition of the correlation matrix. These methods are best suited for *line spectra* - i.e., spectra of sinusoidal signals - and are effective for detecting sinusoids buried in noise, especially when signal to noise ratios are low.

The following functions estimate the *pseudospectrum* S (an indicator of the presence of sinusoidal components in a signal) of the input signal x , and a vector w of normalized frequencies (in rad/sample) at which the pseudospectrum is evaluated. The input p controls the dimensions of the signal and noise subspaces used by the algorithms.

- **Eigenvector method**

```
[S f] = peig(x,p,nfft,fs)
```

- **Multiple Signal Classification (MUSIC) method**

```
[S f] = pmusic(x,p,nfft,fs)
```

The MUSIC algorithm uses Schmidt's eigenspace analysis method. The eigenvector uses a weighted version of the MUSIC algorithm.

Try:

```
edit ssmethod  
ssmethod(3)  
ssmethod(4)
```

Spectrum Viewer in SPTool

The SPTool allows you to view and analyze spectra using different methods. To create a spectrum in SPTool,

1. Select the signal in the **Signals** list in SPTool.
2. Select the **Create** button under the **Spectra** list.

Use the **View** button under the **Spectra** list in the SPTool GUI to display one or more selected spectra.

Try:

```
t = 0:1/100:10-1/100;  
x = sin(2*pi*15*t) + sin(2*pi*30*t);
```


Import this signal into SPTool and view the spectrum using various methods.

5 Time-Varying Spectra

The spectral estimation methods described so far are designed for the analysis of signals with a constant spectrum over time. In order to find time-varying spectra, different methods of analysis and visualization must be used.

The *time-dependent Fourier transform* of a signal is a sequence of DFTs computed using a sliding window. A *spectrogram* is a plot of its magnitude versus time.

```
[B f t] = specgram(x,nfft,fs>window,numoverlap)
```

calculates the time-dependent Fourier transform for the signal in vector x and returns the DFT values B , the frequency vectors f , and the time vectors t . The spectrogram is computed as follows:

1. The signal is split into overlapping sections and applies to the window specified by the *window* parameter to each section.
2. It computes the discrete-time Fourier transform of each section with a length $nfft$ FFT to produce an estimate of the short-term frequency content of the signal; these transforms make up the columns of B . The quantity $length(window) - numoverlap$ specifies by how many samples *specgram* shifts the window.
3. For real input, *specgram* truncates to the first $nfft/2 + 1$ points for $nfft$ even and $(nfft + 1)/2$ for $nfft$ odd.

When called with no outputs, *specgram* displays the spectrogram.

Try:

Time-constant spectrum

```
t = 0:1/100:10-1/100;  
x = sin(2*pi*15*t) + sin(2*pi*30*t);  
specgram(x,256,100,hann(21),15)  
colorbar
```

Time-varying spectrum

```
load handel  
sound(y,Fs)  
specgram(y,512,Fs,kaiser(100,5),75)  
colorbar
```

Q10. Give a very brief explanation of the spectrogram function.
(Remainder of the Lab. has additional exercises. There are no questions to be answered).

Spectrogram Demos

```
specgramdemo  
xpsound
```

Example: Reduced Sampling Rate

This example compares the sound and spectrogram of a speech signal sampled at progressively reduced rates.

If you resample a signal at a fraction of the original sampling frequency, part of the signal's original frequency content is lost. The down-sampled signal will contain only those frequencies less than the new Nyquist frequency. As down-sampling continues, the words in the signal remain recognizable long after the original spectrogram has become obscured. This is a tribute to the human auditory system, and is the basis of signal compression algorithms used in communications.

Try:

```
edit HAL9000  
HAL9000
```

6 Wavelets

One of the drawbacks of the Fourier transform is that it captures frequency information about a signal without any reference to time. For stationary signals this is unimportant, but for time-varying or “bursty” signals, time can be critical to an analysis.

The time-dependent Fourier transform computed by the *specgram* function is one solution to this problem. By applying a DFT to a sliding window in the time domain, *specgram* captures a signal's frequency content at different times. The disadvantage of this method is that it is *uniform* on all time intervals: it does not adjust to local idiosyncrasies in the frequency content. As a result, more subtle nonstationary characteristics of a signal can go undetected.

Wavelet analysis uses a more adaptable method, and is capable of revealing trends, breakdown points, discontinuities in higher derivatives, and self-similarity that the DFT might miss.

A wavelet is a waveform of effectively limited duration that has an average value of zero.

The *discrete wavelet transform* (DWT) computes *coefficients of similarity* between a signal and a sliding wavelet. The coefficients are found with wavelets of different *scales* (widths that approximate different frequencies) to analyze the signal at different resolutions. In a wavelet analysis, a signal is iteratively decomposed into the sum of a lowpass *approximation* and a progressive sequence of highpass *details*.

The Wavelet Toolbox adds wavelet analysis techniques to the signal processing techniques available in the Signal Processing Toolbox.

`wavemenu`

brings up a menu for accessing graphical tools in the Wavelet Toolbox.

Try:

`wavemenu`

Select **Wavelet 1-D**

and then **File** → **Example Analysis** → **Basic Signals** → **Frequency breakdown**

S is the signal

a_5 is the approximation

d_n are the details

What is happening in this signal and how does the wavelet analysis show it?

References

- [1] Naidu, Prabhakar S. *Modern Spectrum Analysis of Time Series*. Florida: CRC Press, 1996.
- [2] Naidu, Prabhakar S. *Modern Digital Signal Processing: An Introduction*. Alpha Science International Ltd, 2003.
- [3] Kay, Steven M. *Modern Spectral Estimation: Theory & Application*. Prentice Hall Signal Processing Series, 1987.

(The material in this lab handout was initially put together by Paul Beliveau. It derives from the MathWorks training document “MATLAB for Signal Processing”, 2006 and from [3].)