

Computer labs

INSTRUCTION

Similarly to the homeworks, the computer-lab assignments are based on exercises from the course textbook (see listing below). The exercises are also appended to this instruction. Solutions should be reported to the lab supervisors before the end of each lab session, either orally or in written format. Preferably, try to present you findings to the lab supervisor as the lab progresses, to avoid waiting for getting passed in the end. Please come well prepared for the labs, e.g. study your notes and/or read about the methods in advance; this will help you finish the lab in time.

MATLAB GUIs that are helpful when solving the exercises are available under `G:\Program\Systemteknik\Spektralanalys\` in the computer labs; however, the MATLAB scripts for the estimation methods used can also be found at the course webpage or at <http://www.prenhall.com/stoica/>. If the G: drive isn't mounted, you can mount the folder needed yourself using "Map Network Drive" and inputting the server location: `\\Atena\Program\Systemteknik\Spektralanalys`.

As opposed to the homeworks, little or no extra MATLAB programming by the student will be needed. Some further instructions are given below, including information on how to run the GUIs created for each lab.

An UpUnet-S (the university computer network for students) account is required for the PC-labs. Information about UpUnet-S can be found at: <http://www.uu.se/student/it-for-studenter/studentnatet/>

Note that passing all 4 labs is a requirement for passing the course.

Lab 1. Periodogram Methods:

- Exercise C2.19: Zero Padding Effects on Periodogram Estimators.
- Exercise C2.20: Resolution and Leakage Properties of the Periodogram.

Complementing instructions:

For exercise C2.19, run `zeropadding.m`. The blue curve shows the zero-padded periodogram. The black curve shows the “theoretical” periodogram (padded with many zeros). The red lines show the true frequencies of the signal. The bottom plot shows the data samples (red) including the amount of zeropadding (black horizontal bar).

For exercise C2.20, Resolution properties, run `resolution.m`. The script shows the normalized periodogram in both linear and dB-scale. The noise level is set to zero (0), the amplitudes of the both sinusoids, a_1 and a_2 , are set to one (1), and the phase of the first sinusoid, ϕ_1 , is set to zero (0). To change the frequency spacing α and the phase of the second sinusoid, ϕ_2 , move the sliders and press the “next” button. Check the “Hamming window” check button and press next to apply the window, otherwise the rectangular window is used. In part (d), it is sufficient to run the `resolution.m` script, i.e., to use $(a1,a2) = (1,1)$ and no noise. To look at the signal and estimated auto-correlation sequence (could be interesting), press the “Plot signal and ACS” button. Note that C2.20a refers to $W_B(\omega)$, which is the triangular, or Bartlett, window; however, the `resolution.m` script uses the rectangular window for simplicity, as the properties are similar.

For exercise C2.20, Spectral leakage, run `leakage.m`. The script shows the normalized periodogram in both linear and dB-scale. To change the frequency spacing α and the amplitude of the second sinusoid, a_2 , move the sliders and press the “next” button. Check the “60dB Cheb Win” check button and press next to apply the window, otherwise the rectangular window is used. The right plot with dB-scale is typically more useful, as the sidelobes are shown more clearly. In part (d), it is sufficient to use the “60dB Cheb Win” to see whether the sinusoids can be resolved or not; no comparison with the Blackman-Tukey estimate is necessary. Again, the rectangular window is used in `leakage.m`, as opposed to the Bartlett window mentioned in C2.20c-d.

Lab 2. Parametric Methods for Rational Spectra:

- Exercise C3.17: Comparison of AR, ARMA and Periodogram Methods for ARMA Signals

Complementing instructions:

Run the file `lab_interface.m`. The interface shows the PSDs (true and estimated) in linear and dB scale, as well as the corresponding pole-zero plot. Choose process, estimation method, and model orders, then click the next button to compute the estimates and update the plots. If you encounter a bug in the interface, such that you cannot choose the MA-order for the LSARMA estimator, please choose the MYW estimator, set the MA-order to zero (0), and then choose the LSARMA estimator (you can now select the MA-order). If this does not work, you may need to restart `lab_interface.m`.

Lab 3. Parametric Methods for Line Spectra:

- Exercise C4.12: Resolution Properties of Subspace Methods for Estimation of Line Spectra.

Complementing instructions:

Run the file `lab_interface.m`. Select input $y(t)$ to use estimated auto-correlation sequence, ACS (from finite and potentially noisy data); “True ACS” makes use of the theoretical ACS. Choose method, order, and noise variance, then press next to compute the estimates and update the plots. Note that for part (d), you do not need to compare with the results from C3.18 (which is a part of the homeworks); however, think about the problem and the differences between the techniques.

Lab 4. Filter Bank Methods:

- Exercise C5.13: The Capon Method.

Complementing instructions:

Run the file `lab_interface.m`. Select which part of C5.13 is being studied, (a) or (b)+(c). In part (a), choose Slepian refined filter bank (RFB) or Capon method (CM); in part (b)+(c), choose RFB, Capon, or LSAR. Options for RFB are the number of filters K (also filter bandwidth or number of averages), and the center frequency ω . For Capon, the options are the filter length m , and the center frequency ω . For LSAR, set the model order n . Red is the true spectrum of the data, and blue is the frequency response if the filter(s) in (a), and the estimated PSD in (b)+(c).

Exercise C2.19: Zero Padding Effects on Periodogram Estimators

In this exercise we study the effect zero padding has on the periodogram.

Consider the sequence

$$y(t) = 10 \sin(0.2 \cdot 2\pi t + \phi_1) + 5 \sin((0.2 + 1/N)2\pi t + \phi_2) + e(t), \quad (2.9.20)$$

where $t = 0, \dots, N - 1$, and $e(t)$ is white Gaussian noise with variance 1. Let $N = 64$ and $\phi_1 = \phi_2 = 0$.

From the results in Chapter 4, we find the spectrum of $y(t)$ to be

$$\begin{aligned} \phi(\omega) = & 50\pi [\delta(\omega - 0.2 \cdot 2\pi) + \delta(\omega + 0.2 \cdot 2\pi)] \\ & + 12.5\pi [\delta(\omega - (0.2 + 1/N) \cdot 2\pi) + \delta(\omega + (0.2 + 1/N) \cdot 2\pi)] + 1 \end{aligned}$$

Plot the periodogram for the sequence $\{y(t)\}$, and the sequence $\{y(t)\}$ zero padded with N , $3N$, $5N$, and $7N$ zeroes.

Explain the difference between the five periodograms. Why does the first periodogram not give a good description of the spectral content of the signal? Note that zero padding does not change the resolution of the estimator.

Exercise C2.20: Resolution and Leakage Properties of the Periodogram

We have seen from Section 2.4 that the expected value of the periodogram is the convolution of the true spectrum $\phi_y(\omega)$ with the Fourier transform of a Bartlett window, denoted $W_B(\omega)$ (see equation (2.4.15)). The shape and size of the $W_B(\omega)$ function determines the amount of *smearing* and *leakage* in the periodogram. Similarly, in Section 2.5 we introduced a windowed periodogram in (2.6.24) whose expected value is equal to the expected value of a corresponding Blackman–Tukey estimate with weights $w(k)$ given by (2.6.31). Different window functions than the rectangular window could be used in the periodogram estimate, giving rise to correspondingly different windows in the correlogram estimate. The choice of window affects the resolution and leakage properties of the periodogram (correlogram) spectral estimate.

Resolution Properties: The amount of smearing of the spectral estimate is determined by the width of the main lobe, and the amount of leakage is determined by the energy in the sidelobes. The amount of smearing is what limits the resolving power of the periodogram, and is studied empirically below.

We first study the resolution properties by considering a sequence made up of two sinusoids in noise, where the two sinusoidal frequencies are “close”. Consider

$$y(t) = a_1 \sin(f_0 \cdot 2\pi t + \phi_1) + a_2 \sin((f_0 + \alpha/N)2\pi t + \phi_2) + e(t), \quad (2.9.21)$$

where $e(t)$ is real-valued Gaussian white noise with zero mean and variance σ^2 . We choose $f_0 = 0.2$ and $N = 256$, but the results are nearly independent of f_0 and N .

- (a) Determine empirically the 3 dB width of the main lobe of $W_B(\omega)$ as a function of N , and verify equation (2.4.18). Also determine the peak sidelobe height (in dB) as a function of N . Note that the sidelobe level of a window function is generally independent of N . Verify this by examining plots of the magnitude of $W_B(\omega)$ for several values of N ; try both linear and dB scales in your plots.
- (b) Set $\sigma^2 = 0$ (this eliminates the statistical variation in the periodogram, so that the bias properties can be isolated and studied). Set $a_1 = a_2 = 1$ and $\phi_1 = \phi_2 = 0$. Plot the (zero-padded) periodogram of $y(t)$ for various α and determine the resolution threshold (*i.e.*, the minimum value of α for which the two frequency components can be resolved). How does this value of α compare with the predicted resolution in Section 2.4?

- (c) Repeat part (b) for a Hamming-windowed correlogram estimate.
- (d) For reasonably high signal-to-noise ratio (SNR) values and reasonably close signal amplitudes, the resolution thresholds in parts (b) and (c) above are not very sensitive to variations in the signal amplitudes and frequency f_0 . However, these thresholds are sensitive to the phases ϕ_1 and ϕ_2 , especially if α is smaller than 1. Try two pairs (ϕ_1, ϕ_2) so that the two sinusoids are in phase and out of phase, respectively, at the center of the observation interval, and compare the resolution thresholds. Also, try different values of a_1 , a_2 , and σ^2 to verify that their values have relatively little effect on the resolution threshold.

Spectral Leakage: In this part we analyze the effects of leakage on the periodogram estimate. Leakage properties can be clearly seen when trying to estimate two sinusoidal terms that are well separated but have greatly differing amplitudes.

- (a) Generate the sinusoidal sequence above for $\alpha = 4$, $\sigma^2 = 0$, and $\phi_1 = \phi_2 = 0$. Set $a_1 = 1$ and vary a_2 (choose $a_2 = 1, 0.1, 0.01$, and 0.001 , for example). Compute the periodogram (using a rectangular data window), and comment on the ability to identify the second sinusoidal term from the spectral estimate.
- (b) Repeat part (a) for $\alpha = 12$. Does the amplitude threshold for identifiability of the second sinusoidal term change?
- (c) Explain your results in parts (a) and (b) by looking at the amplitude of the Bartlett window's Fourier transform at frequencies corresponding to α/N for $\alpha = 4$ and $\alpha = 12$.
- (d) The Bartlett window (and many other windows) has the property that the leakage level depends on the distance between spectral components in the data, as seen in parts (a) and (b). For many practical applications it may be known what dynamic range the sinusoidal components in the data may have, and it is thus desirable to use a data window with a constant sidelobe level that can be chosen by the user. The Chebyshev window (or Taylor window) is a good choice for these applications, because the user can select the (constant) sidelobe level in the window design (see the MATLAB command `chebwin`).

Assume we know that the maximum dynamic range of sinusoidal components is 60 dB. Design a Chebyshev window $v(t)$ and corresponding Blackman–Tukey window $w(k)$ using (2.6.31) so that the two sinusoidal components of the data can be resolved for this dynamic range using (i) the Blackman–Tukey spectral estimator with window $w(k)$, and (ii) the windowed periodogram method with window $v(t)$. Plot the Fourier transform of the window and determine the spectral resolution of the window.

Test your window design by computing the Blackman–Tukey and windowed periodogram estimates for two sinusoids whose amplitudes differ by 50 dB in dynamic range, and whose frequency separation is the minimum value you predicted. Compare the resolution results with your predictions. Explain why the smaller amplitude sinusoid can be detected using one of the methods but not the other.

Exercise C3.17: Comparison of AR, ARMA and Periodogram Methods for ARMA Signals

In this exercise we examine the properties of parametric methods for PSD estimation. We will use two ARMA signals, one broadband and one narrowband, to illustrate the performance of these parametric methods.

Broadband ARMA Process: Generate realizations of the broadband ARMA process

$$y(t) = \frac{B_1(z)}{A_1(z)} e(t)$$

with $\sigma^2 = 1$ and

$$\begin{aligned} A_1(z) &= 1 - 1.3817z^{-1} + 1.5632z^{-2} - 0.8843z^{-3} + 0.4096z^{-4} \\ B_1(z) &= 1 + 0.3544z^{-1} + 0.3508z^{-2} + 0.1736z^{-3} + 0.2401z^{-4} \end{aligned}$$

Choose the number of samples as $N = 256$.

- (a) Estimate the PSD of the realizations by using the four AR and ARMA estimators described above. Use AR(4), AR(8), ARMA(4,4), and ARMA(8,8); for the MYW algorithm, use both $M = n$ and $M = 2n$; for the LS AR(MA) algorithms, use $K = 2n$. Illustrate the performance by plotting ten overlaid estimates of the PSD. Also, plot the true PSD on the same diagram.

In addition, plot pole or pole-zero estimates for the various methods. (For the MYW method, the zeroes can be found by spectral factorization of the numerator; comment on the difficulties you encounter, if any.)

- (b) Compare the two AR algorithms. How are they different in performance?
- (c) Compare the two ARMA algorithms. How does M impact performance of the MYW algorithm? How do the accuracies of the respective pole and zero estimates compare?
- (d) Use an ARMA(4,4) model for the LS ARMA algorithm, and estimate the PSD of the realizations for $K = 4, 8, 12,$ and 16 . How does K impact performance of the algorithm?
- (e) Compare the lower-order estimates with the higher-order estimates. In what way(s) does increasing the model order improve or degrade estimation performance?
- (f) Compare the AR to the ARMA estimates. How does the AR(8) model perform with respect to the ARMA(4,4) model and the ARMA(8,8) model?
- (g) Compare your results with those using the periodogram method on the same process (from Exercise C2.21 in Chapter 2). Comment on the difference between the methods with respect to variance, bias, and any other relevant properties of the estimators you notice.

Narrowband ARMA Process: Generate realizations of the narrowband ARMA process

$$y(t) = \frac{B_2(z)}{A_2(z)} e(t)$$

with $\sigma^2 = 1$ and

$$\begin{aligned} A_2(z) &= 1 - 1.6408z^{-1} + 2.2044z^{-2} - 1.4808z^{-3} + 0.8145z^{-4} \\ B_2(z) &= 1 + 1.5857z^{-1} + 0.9604z^{-2} \end{aligned}$$

- (a) Repeat the experiments and comparisons in the broadband example for the narrowband process; this time, use the following model orders: AR(4), AR(8), AR(12), AR(16), ARMA(4,2), ARMA(8,4), and ARMA(12,6).
- (b) Study qualitatively how the algorithm performances differ for narrowband and broadband data. Comment separately on performance near the spectral peaks and near the spectral valleys.

Exercise C4.12: Resolution Properties of Subspace Methods for Estimation of Line Spectra

In this exercise we test and compare the resolution properties of four subspace methods, Min–Norm, MUSIC, ESPRIT, and HOYW.

Generate realizations of the sinusoidal signal

$$y(t) = 10 \sin(0.24\pi t + \varphi_1) + 5 \sin(0.26\pi t + \varphi_2) + e(t), \quad t = 1, \dots, N$$

where $N = 64$, $e(t)$ is Gaussian white noise with variance σ^2 , and where φ_1, φ_2 are independent random variables each uniformly distributed on $[-\pi, \pi]$.

Generate 50 Monte–Carlo realizations of $y(t)$, and present the results from these experiments. The results of frequency estimation can be presented comparing the sample means and variances of the frequency estimates from the various estimators.

- (a) Find the exact ACS for $y(t)$. Compute the “true” frequency estimates from the four methods, for $n = 4$ and various choices of the order $m \geq 5$ (and corresponding choices of M and L for HOYW). Which method(s) are able to resolve the two sinusoids, and for what values of m (or M and L)?
- (b) Consider now $N = 64$, and set $\sigma^2 = 0$; this corresponds to the finite data length but infinite SNR case. Compute frequency estimates for the four techniques again using $n = 4$ and various choices of m, M and L . Which method(s) are reliably able to resolve the sinusoids? Explain why.
- (c) Obtain frequency estimates from the four methods when $N = 64$ and $\sigma^2 = 1$. Use $n = 4$, and experiment with different choices of m, M and L to see the effect on estimation accuracy (*e.g.*, try $m = 5, 8$, and 12 for MUSIC, Min–Norm and ESPRIT, and try $L = M = 4, 8$, and 12 for HOYW). Which method(s) give reliable “super–resolution” estimation of the sinusoids? Is it possible to resolve the two sinusoids in the signal? Discuss how the choices of m, M and L influence the resolution properties. Which method appears to have the best resolution?
You may want to experiment further by changing the SNR and the relative amplitudes of the sinusoids to gain a better understanding of the differences between the methods.
- (d) Compare the estimation results with the AR and ARMA results obtained in Exercise C3.18 in Chapter 3. What are the major differences between the techniques? Which method(s) do you prefer for this problem?

Exercise C5.13: The Capon Method

In this exercise we compare the Capon method to the RFB and AR methods. Consider the sinusoidal data sequence in equation (2.9.20) from Exercise C2.19, with $N = 64$.

- (a) We first compare the data filters corresponding to a RFB method (in which the filter is data independent) with the filter corresponding to the CM Version-1 method using both $m = N/4$ and $m = N/2 - 1$; we choose the Slepian RFB method with $K = 1$ and $K = 4$ for this comparison. For two estimation frequencies, $\omega = 0$ and $\omega = 2\pi \cdot 0.1$, plot the frequency response of the five filters (1 for $K = 1$ and 4 for $K = 4$) shown in the first block of Figure 5.1 for the two RFB methods, and also plot the response of the two Capon filters (one for each value of m ; see (5.4.5) and (5.4.8)). What are their characteristic features in relation to the data? Based on these plots, discuss how data dependence can improve spectral estimation performance.
- (b) Compare the two Capon estimators with the RFB estimator for both $K = 1$ and $K = 4$. Generate 50 Monte-Carlo realizations of the data and overlay plots of the 50 spectral estimates for each estimator. Discuss the similarities and differences between the RFB and Capon estimators.
- (c) Compare Capon and Least Squares AR spectral estimates, again by generating 50 Monte-Carlo realizations of the data and overlaying plots of the 50 spectral estimates. Use $m = 8, 16,$ and 30 for both the Capon method and the AR model order. How do the two methods compare in terms of resolution and variance? What are your main summarizing conclusions? Explain your results in terms of the data characteristics.