

EE512 - Applied Biomedical Signal Processing

Midterm Exam (“mock” exam which does NOT count for evaluation)

Remarks

- There are 4 exercises and 12 questions. The 4 exercises and 12 questions are worth 5/16 (approx. 0.3125) points each, hence a total of 5 points. There is an additional point for presence. [This information is for your own auto-evaluation.]
- The 4 exercises relate to the Lectures (Module 02 to Module 06).
- The 12 questions relate to the Modules 02 to 07 and corresponding Laboratories.
- All questions/problems are on the 3 pages.
- Focus on the essential ideas (like 5-10 lines for each question), but do not limit your answer to a couple of keywords.
- If you get stuck on one question, do not spend too much time on it.

Exercises

1. Let a filter be defined by the difference equation between the input x and the output y :

$$y(n) = x(n) + 2x(n-1) + x(n-2)$$

1.1. Demonstrate without any computation that this filter is a linear phase one.

The filter's coefficient has an axial symmetry axis ([1, 2, 1] and thus the filter is linear phase (see Module 02- Basics I, Slide 35)

1.2. What is the phase response of this filter?

The image shows a handwritten derivation of the phase response of the filter. The steps are as follows:

$$\begin{aligned} H(z) &= 1 + 2z^{-1} + z^{-2} = (1+z^{-1})^2 = (1+e^{-j\omega})^2 \\ &= (1 + \cos(\omega) - j\sin(\omega))^2 = (1 + \cos(\omega))^2 - 2j\sin(\omega)(1 + \cos(\omega)) - \sin^2(\omega) \\ &= 1 + 2\cos(\omega) + \cos^2(\omega) - \sin^2(\omega) - 2j\sin(\omega)(1 + \cos(\omega)) \\ &= 2\cos(\omega) + 2\cos^2(\omega) - 2j\sin(\omega)(1 + \cos(\omega)) \\ &= 2(1 + \cos(\omega)) (\cos(\omega) - j\sin(\omega)) \\ \angle H(z) &= \tan^{-1} \left(\frac{\text{Im}(H(z))}{\text{Re}(H(z))} \right) = \tan^{-1} \left(\frac{-\sin(\omega)}{\cos(\omega)} \right) = \underline{\underline{-\omega}} \end{aligned}$$

(Module 02 - Basics I, Slide 32)

1.3. If x is a white noise with variance 1, what is the autocorrelation function of y ?

$$y(n) = h(n) * x(n)$$

$$R_{yy}(k) = y(n) * y(n-k) = (h(n) * x(n)) * (x(n-k) h(n-k)) =$$

$$= h(n) * (x(n) * x(n-k)) h(n-k) =$$

$$= h(n) * \delta(n-k) * h(n-k) = h(n) * h(n-k)$$

$$R_{yy}(k) = \begin{bmatrix} 1 & 4 & 6 & 4 & 1 \end{bmatrix}$$

$k = -2 \quad k = -1 \quad k = 0 \quad k = 1 \quad k = 2$

(see Module 03 - Basics II, Slides 14 & 29)

2. A signal x with sampling frequency f_s is generated using the following AR model:

$$x(n) = -0.9 x(n-1) + \varepsilon(n)$$

with ε being white noise with variance σ_ε .

2.1. Write the expression for the power spectral density of x in terms of these parameters.

From the expression we can see this is an AR($p = 1$) model (refers to Module 05 - Linear models I, slide 10), with $a_1 = 0.9$. The general expression for the power spectral density (see Module 06 - Linear models II, slide 4) is:

$$\hat{P}_x(f) = \frac{\sigma_\varepsilon^2}{f_s \left| 1 + \sum_{l=1}^p a_l e^{-2\pi j \frac{f}{f_s}} \right|^2}$$

And so for this case the expression becomes:

$$\hat{P}_x(f) = \frac{\sigma_\varepsilon^2}{f_s \left| 1 + 0.9 e^{-2\pi j \frac{f}{f_s}} \right|^2}$$

2.2. Is the power spectral density higher in the low-frequency or high-frequency range?

We wish to understand how \hat{P}_x behaves when varying between 0 (the lowest frequency of interest) and $f_s/2$ (the highest). We present two valid (and partially related) approaches to look at the problem:

- 1st approach -

The denominator of the expression obtained in 2.1. can be developed as follows (including a replacement of $\frac{f}{f_s}$ by ω , just to simplify):

$$\begin{aligned} \left| 1 + 0.9e^{-2\pi j \frac{f}{f_s}} \right|^2 &= (1 + 0.9e^{-2\pi j \omega})^* (1 + 0.9e^{-2\pi j \omega}) \\ &= (1 + 0.9e^{+2\pi j \omega})(1 + 0.9e^{-2\pi j \omega}) \\ &= 1 + 0.9e^{-2\pi j \omega} + 0.9e^{+2\pi j \omega} + 0.81 \\ &= 1.81 + 1.8 \cos(2\pi \omega) \end{aligned}$$

We can see that this term varies because of the cosine: it is highest at $\omega = 0$, and then decreases monotonically to 1.81 at $\omega = \frac{1}{2}$. Since this term is the denominator of the power spectral density, we can then conclude, inversely, that the PSD has its peak at $f_s/2$ (where the denominator is lowest).

- 2nd approach -

Another strategy would be to draw insight from digital filter theory. Considering an extension of \hat{P}_x to the full complex domain:

$$\hat{Q}_x(z) = \frac{\sigma_\varepsilon^2}{f_s |1 + 0.9z^{-1}|^2}$$

where $z \in \mathbb{C}$, we can see that that it has a “pole” at $z = -0.9$, i.e., in the left half-plane on the real axis. Thus, for a $z = e^{j\omega}$, \hat{Q}_x will increase as ω increases towards $\pm\pi$, i.e., the higher frequencies – it can be seen as the “equivalent” to a single-pole high-pass filter. Consequently, we can conclude that \hat{P}_x is higher in the high-frequency range.

3. Find all the parameters of an AR model of order $p = 1$ for a signal x with autocorrelation function values:

$$R_{xx}(0) = 3 \quad R_{xx}(1) = 1$$

An AR(1) model is described by two parameters: a_1 and σ_1 . At such a low order these parameters can be easily calculated by applying the Yule-Walker equations (see general form in Module 05 - Linear models I, slide 18, and this specific case of $p = 1$ in slide 19):

$$\begin{aligned} a_1 &= -\frac{R_{xx}(1)}{R_{xx}(0)} = -\frac{1}{3} \\ \sigma_1^2 &= R_{xx}(0) + a_1 R_{xx}(1) = 3 - \frac{1}{3} = \frac{8}{3}. \end{aligned}$$

4. Six consecutive samples $x(0), x(1), \dots, x(5)$ of a blood pressure signal are available. One wishes to estimate the coefficients a_1 and a_2 of a linear predictor of order $p = 2$ for this signal.

4.1. Write the expression that defines this model, i.e. how a prediction for timepoint n , $\hat{x}(n)$, is obtained based on past instances of x and on the model parameters a_1, \dots, a_p .

A linear predictor of order 2 can be written as (see Module 05 - Linear models I, slide 28):

$$\hat{x}(n) = -a_1x(n-1) - a_2x(n-2)$$

4.2. An approach to estimate the model parameters, based on the six signal samples that we have available, is using the matrix formulation for the mean squared error (MSE) criterion, $\|\mathbf{e}\|^2$, with $\mathbf{e} = \mathbf{x} - X\mathbf{a}$ (**bold** denotes a vector, CAPITAL denotes a matrix). Following the “covariance” approach, write the vectors \mathbf{x} , \mathbf{a} , and the matrix X for this finite sample.

The “covariance” approach includes in the estimation only the prediction errors that are “fully supported”, i.e. the prediction equations for which the model does not extend beyond the samples that we have available (see Module 05 - Linear models I, slide 40 and also 36). This therefore includes only the prediction errors for timepoints $n = 2, 3, 4$ and 5.

For the matrix formulation that is requested, the objects will then be:

$$\mathbf{x} = \begin{bmatrix} x(2) \\ x(3) \\ x(4) \\ x(5) \end{bmatrix}, \quad \mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}, \quad X = \begin{bmatrix} x(1) & x(0) \\ x(2) & x(1) \\ x(3) & x(2) \\ x(4) & x(3) \end{bmatrix}$$

Questions

1. Describe the advantages / drawbacks of FIR and IIR filters.

FIR:

- + always stable
- + can exhibit linear phase
- + impulse response is finite (duration of transitions is independent of the input)
- more coefficients are needed to match IIR
- less sensitive to numerical approximations

IIR:

- + reduced number of coefficients
- stability is only ensured when all the poles have a radius smaller than one
- no linear phase
- impulse response is infinite (duration of transition is dependent of the input)

(see Module 02 - Basics, slide 37 + lab)

2.

a) What is a linear phase filter?

A filter with a phase that is linear which means that all the frequency components of the filtered signal are delayed by the same amount. (see Module 02 - Basics, slide 35)

b) What are the properties of the impulse response of a linear phase filter?

The filter is a FIR filter with an axial or central symmetry in its impulse response. (see Module 02 - Basics, slide 35)

c) What is a zero-phase filter?

A linear filter of any kind (FIR or IIR) that is firstly applied to the time-reversed signal and secondly applied to the time-reversed output signal. As the filter is applied in both time directions its phase is equal to zero for any frequency. Such a filter is suitable for comparing signals during analysis, but it is not suitable for real time implementations. (see Module 02 - Basics, slide 36)

3. Using the Welch algorithm to estimate the power spectral density:

a) What is the influence of the length of the block?

The length of the block impacts the frequency resolution that is f_s/N_{block} (see Module 03 - Basics II, slide 20)

b) What is the motivation of multiplying each block by a window before the DFT?

Without applying a window, the estimated spectrum will exhibit oscillations due to the sharp transitions at the boundaries of the block (Gibbs phenomenon). The application of a window before the computation of the DFT reduces these oscillations. Diverse kinds of windows can be selected, providing a compromise between oscillations and spectral resolution. (see Module 03 - Basics II, slide 16)

4. Using the measurement of mean blood pressure and cardiac inter-beat intervals:

a) What is a suitable sampling frequency to analyse the control of the autonomic nervous system?

The maximum frequency corresponding to the autonomic nervous system control is 0.5 Hz. Therefore, the sampling frequency must be larger than twice to avoid spectral aliasing. Practically sampling frequencies slightly higher (typ. 1.5 - 2.0 Hz) are used. (see Module 02 – Basics II, Lab. 2)

b) What is the advantage of using the auto-correlation method to estimate the power spectral density compared to discrete Fourier transform?

The Fourier transform is not a consistent spectral estimator and does not permit to obtain an averaged estimation of the power spectral density. The use of the central part of the autocorrelation function permits to increase the signal to noise ratio of the spectral estimate. (see Module 02 – Basics II, Slides 6 + 15)

5. Suppose a microphone is recording someone speaking for some time (“period A”); then the person stops talking and the microphone records only silence, which we can consider to be dominated by white noise, for some time (“period B”). An $AR(p)$ model,

with $p > 0$, is fit to period A and to period B separately. In which period is the ratio *excitation variance / signal variance* the highest? Why?

Assuming an adequate fit to the AR(p) model $x(n) = \varepsilon(n) - \sum_{l=1}^p a(l)x(n-l)$, where ε is the excitation and is assumed to be white noise, then we can expect that the ratio *excitation variance / signal variance* will be highest in period B, i.e., when the microphone is recording silence. This is because in the silent period the recording will be dominated by white noise, which will be attributed to the excitation term of the model. The deterministic part will have a negligible contribution, and therefore the excitation variance will approximate the full signal variance, leading to the highest ratio. (see Module 05 - Linear models I, slide 46, and lab Q2.2).

6. When an AR model is estimated with an order that is too low, what can we expect to find about the estimated excitation signal ε ? Why?

If the order is too low, the excitation term is likely to deviate significantly from white noise. This is because when the order assigned to the model is too low, the model will be unable to capture the structure presents fully and appropriately in the signal, i.e., its dependencies across time. As a result, the residual, non-modeled dependencies are expected to “leak” to the excitation term. (This effect is indeed explored in certain approaches for order estimation, wherein the estimated excitation signal is tested for whiteness to evaluate if the model order is high enough to capture all the signal structure or not – see Module 05 - Linear models I, slide 43).

7. Regarding the estimation of the power spectral density based on AR or ARMA models:

a) What are the potential advantages with respect to non-parametric estimation approaches (e.g., Welsh method)?

AR and ARMA models impose certain assumptions about the signal properties/structure, and in doing so they allow more robust estimates of the power spectral density, and more specifically, estimates that are less sensitive to noise in the available signal sample (see Module 06 - Linear models II, slide 5, and lab Q1.3). This can be especially useful when the sample is either relatively noisy or relatively short.

b) What are the risks when the model order p is too low?

If the order is too low, the signal structure will not be accurately and fully captured in the AR parameters a_i , and therefore the power spectral density estimate (which relies on these parameters to describe the signal spectrum) will be inaccurate. It is possible that certain meaningful spectral features (e.g., peaks) will be absent or smoothed out. (see Module 06 - Linear models II, lab Q1.2).

8. Regarding the Pisarenko harmonic retrieval method:

a) What type of signals is this method adequate for?

This method is appropriate for signals that are linear combinations of sinusoids with added white noise (see Module 06 - Linear models II, slide 32).

b) What model parameter(s) can the method estimate?

The method can estimate the power and frequency of each of the sinusoids, as well as the power of the white noise (see Module 06 - Linear models II, slide 41).

c) What model parameter(s) can the method not estimate?

The method is unable to estimate the phases of the sinusoids (see Module 06 - Linear models II, slide 37).

9. In a spectrogram, which window function should you choose if you wanted to maximize your chances to differentiate between the heart rate and a motion frequency component that is very close to it? Why?

The choice of window function is a trade-off between frequency resolution (ability to distinguish frequencies that are close to each other) and dynamic range (ability to distinguish frequencies of different strengths). In our case, the heart rate-related and the motion-related frequency components are very close to each other in frequency, and approximately of similar power. This should clearly guide the trade-off towards a good frequency resolution to help tell them apart, and therefore choosing a window such as the rectangular window. (see Module 04 - Time frequency, slide 13).

10. When evaluating the sympathovagal balance through time-frequency analysis, why would you probably have a lower time resolution on the low frequency component if you used the wavelet transform instead of the short-term Fourier transform?

In real signals, high frequencies often span over a shorter time than long frequencies. This observation makes wavelets naturally more suited for time-frequency analysis as wavelets have a time support (the equivalent of the window length) that depends on their frequency. Thus, low frequencies will be projected onto longer wavelets, at the cost of a low time resolution. With the STFT, the window length is the same for all frequencies. Because of the trade-off between frequency and time resolution, the window in the STFT will often be longer than desired in the high frequencies and shorter than desired in the low frequencies. As a result, the time resolution at low frequencies is often better with the STFT compared to the wavelet transform. (see Module 04 - Time frequency, slide 37).

11. One wants to estimate the instantaneous frequency (IF) of atrial fibrillation from the surface electrocardiogram (ECG). Following the subtraction of the ventricular activity (the atrial and ventricular activities overlap on the surface ECG), the remaining atrial signals are usually bandpass filtered to remove the spectral frequency components outside of the frequency range of atrial fibrillation.

a) From a theoretical perspective, why should the signals be bandpass filtered before estimating the instantaneous frequency?

The underlying assumption for the estimation of instantaneous frequency is that the signal is locally mono-component (or at least narrow-band). This is rarely the case in practice and thus some band-pass filtering should be performed. (see Module 07 – Instantaneous Frequency Estimation, slides 9-12, and Lab Exercise 1)

b) What would be the benefit of bandpass filtering in terms of accuracy of IF estimation?

The IF estimate obtained on the filtered signals would be smoother (with less distortions) and centered on the correct frequency value, whereas the IF estimate from the raw atrial signals would present large distortions due to the large signal spectral variations as a function of time. (see Module 07 - Instantaneous Frequency Estimation, slide 9-12, and Lab Exercise 1)

12. An oscillatory signal with slowly-time varying amplitude and frequency can be expressed as $x(t) = A(t) \cos(\Phi(t))$. How can instantaneous amplitude and frequency be defined through the analytic representation of $x(t)$?

An analytic signal is a complex signal without negative frequency components. It is defined as $x_a(t) = x(t) + jx_h(t)$, where $x_h(t)$ is the Hilbert transform of $x(t)$. (see Module 07 - Instantaneous Frequency tracking, slides 12, 13, and 14)

For a cosine signal, the analytic signal is $x_a(t) = A(t) \cos(\Phi(t)) + jA(t) \sin(\Phi(t)) = A(t) \exp(j\Phi(t))$. The instantaneous amplitude is defined as $|x_a(t)| = A(t)$ and the instantaneous frequency is defined as $f(t) = \frac{1}{2\pi} \frac{d\Phi(t)}{dt}$, where the instantaneous phase is $\angle x_a(t) = \phi(t)$. (see Module 07 - Instantaneous Frequency tracking, slides 12, 13, and 14)