Signal Analysis and Estimation

Paul Van den Hof
Arjan den Dekker

Lecture Notes Stochastic Signal Analysis, sc3011tn

Course 2010/2011

DCSC
Delft Center for Systems and Control

Delft Center for Systems and Control
Delft University of Technology

Address of the authors:

Paul M.J. Van den Hof
Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
Tel. +31-15-2784509; Fax: +31-15-2786679
Email: p.m.j.vandenhof@tudelft.nl
Web: http://www.dcsc.tudelft.nl/~pvandenhof

Arnold J. den Dekker
Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
Tel. +31-15-2781823; Fax: +31-15-2786679
Email: a.j.dendekker@tudelft.nl
Web: http://www.dcsc.tudelft.nl/~adendekker

Contents

[1] H. Amandus Schwarz (1843-1921) was a German mathematician.

List of Symbols

$\mathbb{C}$	Set of complex numbers
$C_{\mathbf{x}}(t_1, t_2)$	Autocovariance function of stochastic process $\mathbf{x}$
$C_{\mathbf{yx}}(t_1, t_2)$	Crosscovariance function of stochastic processes $\mathbf{y}$ and $\mathbf{x}$
$C_{\mathbf{x}}(\tau)$	Autocovariance function of stationary stochastic process $\mathbf{x}$
$C_{\mathbf{yx}}(\tau)$	Crosscovariance function of jointly stationary stochastic processes $\mathbf{y}$ and $\mathbf{x}$
$\delta(t)$	Discrete pulse function: $\delta(0) = 1; \delta(t) = 0,\ t \neq 0.$
$det(\cdot)$	Determinant of a matrix
ε	One-step-ahead prediction error
$\mathbf{e}$	White noise stochastic process
$\mathbb{E}$	Expectation operator
$\mathbb{E}_{\mathbf{x}}$	Expectation over the p.d.f. of random variable $\mathbf{x}$
$\mathbb{E}_{\mathbf{x}\|\mathbf{y}}$	Conditional expectation over the p.d.f. $f_{\mathbf{x}\|\mathbf{y}}$
$\mathbb{E}_{\mathbf{x},\mathbf{y}}$	Expectation over the joint p.d.f. of $\mathbf{x}$ and $\mathbf{y}$
$f_{\mathbf{x}}$	Probability density function of random variable $\mathbf{x}$
$f_{\mathbf{x},\mathbf{y}}$	Joint probability density function of r.v.'s $\mathbf{x}$ and $\mathbf{y}$
$f_{\mathbf{x}\|\mathbf{y}}$	Conditional probability density function
$F_{\mathbf{x}}(x)$	Distribution function of r.v. $\mathbf{x}$, i.e. $P(\mathbf{x} \leq x)$
$F_{\mathbf{x},\mathbf{y}}(x, y)$	Joint distribution function of r.v.'s $\mathbf{x}$ and $\mathbf{y}$, i.e. $P[(\mathbf{x} \leq x) \wedge (\mathbf{y} \leq y)]$
$\varphi(t)$	Regression vector
$\Phi_{\mathbf{x}}$	(Auto-)spectral density of stochastic process $\mathbf{x}$
$\Phi_{\mathbf{xy}}$	Cross-spectral density of stochastic processes $\mathbf{x}$ and $\mathbf{y}$
i	Imaginary number
$\mu_{\mathbf{x}}$	Expectation of random variable $\mathbf{x}$, i.e. $\mathbb{E}[\mathbf{x}]$
$\mathbb{N}$	Set of natural numbers, 1, 2, $\cdots$
$\mathcal{N}$	Gaussian or normal probability density function
N	Number of data
ω	Radial frequency
ω_s	Sampling (radial) frequency
q^{-1}	Delay operator, $q^{-1}\mathbf{x}(t) = \mathbf{x}(t-1)$
$\mathbb{R}$	Set of real nunbers
$r_{\mathbf{xy}}$	Correlation between scalar random variables $\mathbf{x}$ and $\mathbf{y}$
$r_{\mathbf{xx}}$, $r_{\mathbf{x}}^2$	Correlation of scalar random variable $\mathbf{x}$
$R_{\mathbf{x}}$	Correlation matrix of vector random variable $\mathbf{x}$
$R_{\mathbf{xy}}$	Correlation matrix of two vector random variables $\mathbf{x}$ and $\mathbf{y}$
$R_{\mathbf{x}}(t_1, t_2)$	Autocorrelation function of stochastic process $\mathbf{x}$
$R_{\mathbf{xy}}(t_1, t_2)$	Crosscorrelation function of stochastic processes $\mathbf{x}$ and $\mathbf{y}$
$R_{\mathbf{x}}(\tau)$	Autocorrelation function of stationary stochastic process $\mathbf{x}$
$R_{\mathbf{xy}}(\tau)$	Crosscorrelation function of jointly stationary stochastic processes $\mathbf{x}$ and $\mathbf{y}$
$\rho_{\mathbf{xy}}$	Correlation coefficient between random variables $\mathbf{x}$ and $\mathbf{y}$
$\sigma_{\mathbf{x}}^2$, $\sigma_{\mathbf{xx}}$	Variance of scalar random variable $\mathbf{x}$, i.e. $\mathbb{E}[\mathbf{x} - \mu_{\mathbf{x}}]^2$
$\sigma_{\mathbf{xy}}$	Covariance of scalar random variables $\mathbf{y}$ and $\mathbf{x}$, i.e. $\mathbb{E}[\mathbf{x} - \mu_{\mathbf{x}}][\mathbf{y} - \mu_{\mathbf{y}}]^*$
$\Sigma_{\mathbf{x}}$	Covariance matrix of a vector random variable: $\mathbb{E}[\mathbf{x} - \mu_{\mathbf{x}}][\mathbf{x} - \mu_{\mathbf{x}}]^*$
$\Sigma_{\mathbf{xy}}$	Covariance matrix of two vector random variables: $\mathbb{E}[\mathbf{x} - \mu_{\mathbf{x}}][\mathbf{y} - \mu_{\mathbf{y}}]^*$
θ	Parameter vector
$\hat{\boldsymbol{\theta}}_N$	Estimated parameter based on N data

Θ	Domain of parameter vector
T_s	Sampling interval
$tr(A)$	Trace of a matrix
$\mathbf{x}$	Random variable
$\mathbf{x}(t)$	Stochastic process $\mathbf{x}$ as a function of t
x	One realization of random variable $\mathbf{x}$
$x(t)$	One realization of stochastic process $\mathbf{x}$ at time t
$\mathbb{Z}$	Set of integer numbers (positive and negative)
$\star$	convolution
$(\cdot)^*$	Complex conjugate transpose
$(\cdot)^T$	Transpose
$\|A\|_F$	Frobenius-norm of a matrix
$\|A\|_2$	2-Norm of a matrix

Chapter 1

Introduction

In this chapter an introduction is given to the role and use of (measurement) signals in science and engineering. The choice for probabilistic or deterministic signal models is discussed and an overview is given of different signal analysis and estimation problems.

1.1 Introduction

Signal analysis and estimation are two important subjects that are closely tied to many problems of measurement that occur in almost all areas of science and engineering. Measurements of physical quantities lead to observations or signals and in most cases it is our aim to extract particular information from the measurement data. In this respect signals are considered to be observed quantities that are sequenced most often in time, but possible also in other domains as e.g. the spatial domain. The characterization and analysis of signals is the first step in retrieving the necessary physical information that is present (or hidden) in the measurements.

- Consider for instance the radial position of the laser spot on the track of a Compact Disc in a CD-player, as schematically depicted in figure 1.1. An example of a time sequenced signal of the distance of the center of the spot to the track is given in figure 1.2. It is indirectly measured by measuring the intensity of the reflected spot in the so-called OPU (optical pick-up unit). It is important to characterize the principal properties of this signal, e.g. in order to control the actuator that positions the laser spot on the Compact Disc so as to attenuate the track following error as much as possible.

- In microscopy the several pixels of a still image can be considered a signal that is indexed over space rather than over time. To every pixel in the image a (real-valued) light intensity is connected, which is ordered either row- or column-wise into a sequenced signal. In figure 1.3 an image is shown from transmission electron microscopy (TEM) where the atomic structure of some material is investigated. In other image processing applications, sequences of timed images will occur, introducing an important time dependency also. In the example of TEM data, the image contains information on the atomic structure of the material, but this information is not exact, but rather “noise disturbed”. The image (signal) needs to be processed in such a

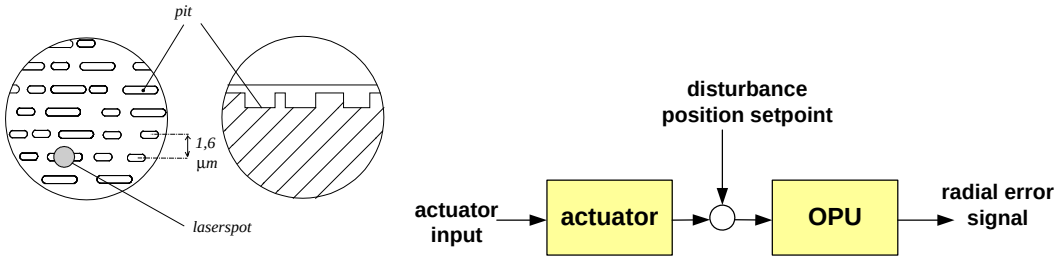

Figure 1.1: A CD player track (left) and a block diagram of the radial error (right).

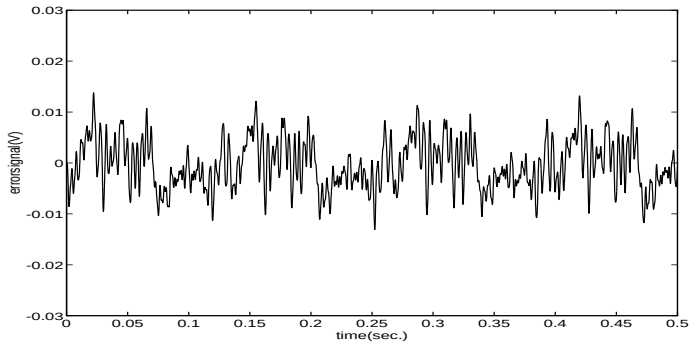

Figure 1.2: Radial position error signal of the laser spot in a CD-player, as a function of time.

way that the underlying physical information is retrieved with optimal accuracy and precision.

In order to quantify physical variables one generally has to measure. To find out the temperature in a particular reactor, the heartbeat of a phoetus, the position of a star in space, or the position of a laser beam when it "writes" an electronic circuit map on a silicon wafer, measurement systems are developed and installed to measure the respective quantities. However in most of these situations the measurements that the equipment deliver are not "exact". There are several underlying reasons for this:

- The measurement can be contaminated by noise. Due to all kinds of disturbing phe-

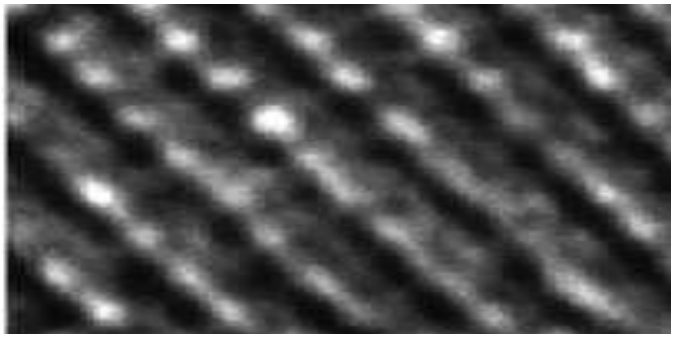

Figure 1.3: Image from transmission electron microscopy.

nomena, measurements are affected by perturbations that are considered not relevant for the problem at hand. The instruments will have finite precision and so they will introduce measurement noise. However there are also effects of other noisy perturbations. For a communication system for either speech, images or general type of data, the several noise sources that occur in coding and decoding the signals are schematically depicted in figure 1.4.

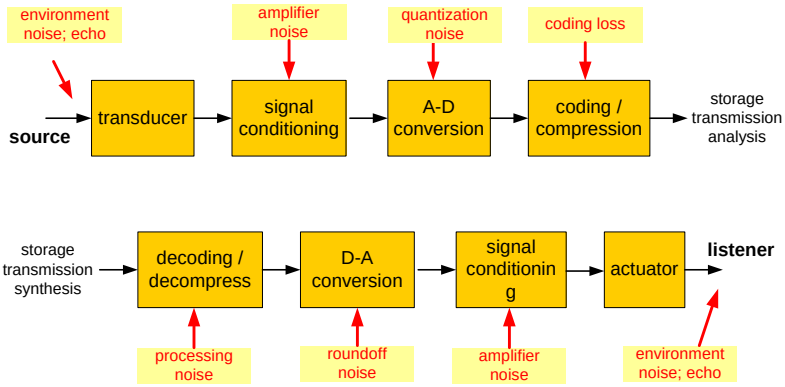

Figure 1.4: Sequence of coding and decoding operations in a communication channel for speech, images or data in general.

- Quantum noise, i.e. noise attributable to the intrinsic discrete and probabilistic nature of physical phenomena and their interactions. One can distinguish photon noise[1] in optical systems and shot noise [2] in electrical (semi)conductors.

Besides signals that are disturbed by noise, there are also many signals that show "noisy" behaviour by nature. When measuring the voltage between two locations on the human skull, a noisy-like signal will be measured. However this signal includes important information on the functionality of the human brain.

Indirect measurements

It is not only disturbance effects and noises that limit the possibility of making exact measurements and that complicate the handling of measurement data. In many measurement systems it is not possible to measure the exact quantity that one is interested in, but *indirect measurements* are used, i.e. physical variables are measured that are related to the variable of interest (in a particular way), but they are not the same. E.g. one measures the number of photons that hit a detector plate, but one is interested in an image, and maybe even in structural information on the object that is studied, e.g. in microscopy. Or: one is interested in the temperature, and one measures the volume of a certain amount of mercury. In these situations a model of the measurement system will have to determine the relation between the measured variables and the variables to be determined or estimated.

[1] attributable to the statistical nature of optical quanta.
[2] Noise caused by random fluctuations in the motion of charge carriers in a conductor.

1.2 Signals as source of information

Example 1.1 (EEG) In analysis of electroencephalographic (EEG) signals voltages are measured between well-defined locations on the human skull. Typical measurements are indicated in the right upper part of figure 1.5. EEG is a practically useful tool for studying the functional states of the brain (e.g. sleep-stage analysis) and for diagnosing functional brain malfunctions as e.g. epilepsy. For an assessment of the EEG, the on-average power spectrum of these signals, depicted in the lower part of figure 1.5, is essential. This power spectrum is a important tool in diagnosing particular diseases, e.g. by evaluating the relative importance of the several frequency regions.

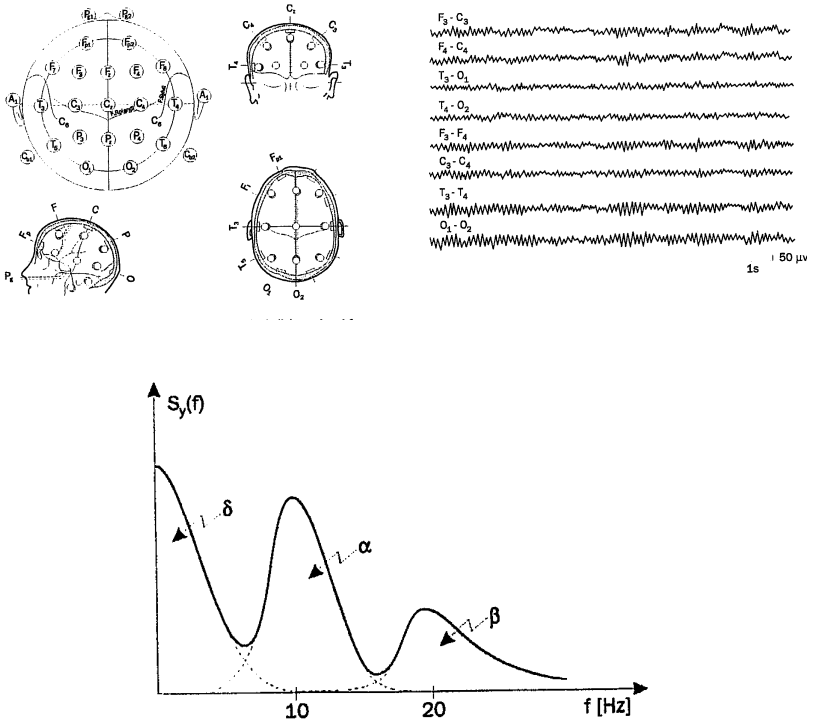

Figure 1.22 Spectrum of a typical 'normal' EEG recording.

Figure 1.5: EEG signal recordings.

Example 1.2 (Radar) In radar a transmit pulse is transmitted and reflected by an object to be detected. The received waveform contains a reflection of the transmitted signal, contaminated with noise, due to all kind of disturbance effects e.g. from the atmosphere, see Figure 1.6. A detection of the time delay between transmission and reception has to be done in order to quantify the distance to the target.

Example 1.3 (Adaptive Optics) In adaptive optics, the light that is received from a star on a (large scale) telescope will be contaminated with all kind of disturbance effects, e.g. due to atmospheric turbulence. The optical signals are filtered and corrected by way of a deformable mirror that changes the orientation of each separate element of the deformable

mirror, so as to compensate for the disturbance effects on the optical signals. The control system is steered by a measured signal that is deflected by the wavefront sensor. In order to control the system appropriately, the control system will need to use (prior) knowledge of the type of disturbances that affect the signals. Therefore a characterization of these disturbance signals is required. A typical example of the effect of an adaptive control system on the image received from a triple star is given in figure 1.7.

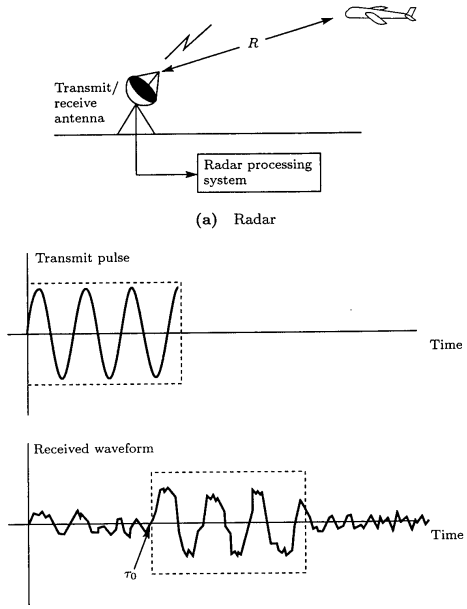

Figure 1.6: Radar signal recordings.

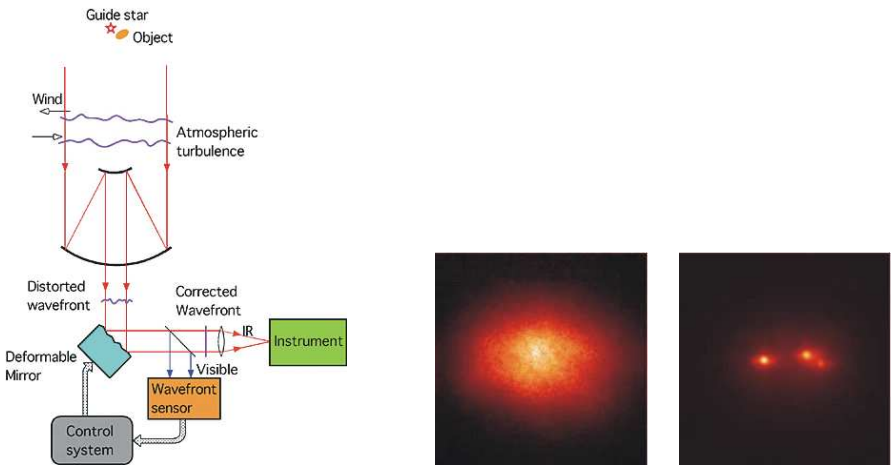

Figure 1.7: Adaptive Optics system (left), and images of a triple star received by a telescope without (middle) and with (right) a correction by an adaptive optics control system.

Example 1.4 (Production Technology) Acoustic emission (AE) is a sound wave or, more properly, a stress wave that travels through a material as the result of some sudden

release of strain energy. When measured such waves can be employed for quality control in manufacturing operations such as turning, grinding and milling.
In these applications, AE signals have a highly random nature, which is related to the complexity of the underlying cutting process. There are multiple AE sources that can be identified, such as deformation and fracture of the workpiece material, chip breakage, rubbing between tool and the material and between the tool and the chips, vibrations of the machine, etc. (see Figure 1.8). All these sources generate sound waves that when measured at some position in the machine result in a signal that understandably has a highly random nature. Nonetheless, the properties of this AE signal appear to change when the cutting conditions change, in particular when the cutting tool wears. Tool condition monitoring systems have been developed that use these changing characteristics of the AE signal to estimate the tool wear. This improves the overall quality of the production, as in general worn tools adversely affect the surface finish of the workpiece. By means of a tool wear condition monitoring system, the operator of the machine can be alerted to the state of tool, thereby avoiding undesirable consequences.

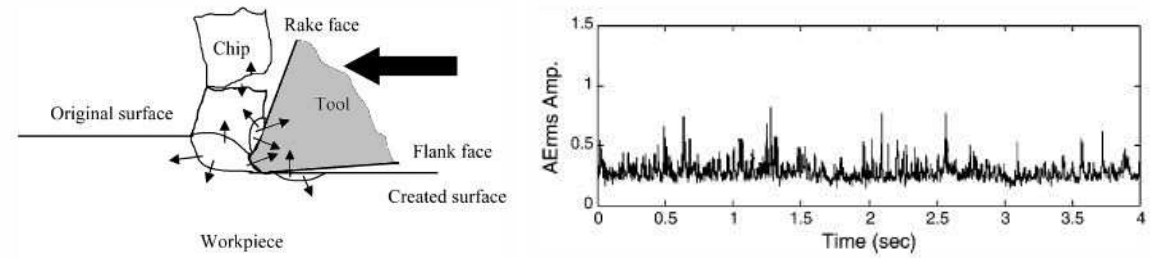

Figure 1.8: Cutting process (left), and example of a generated sound wave signal (right).

1.3 Deterministic or probabilistic signals models

The common feature of the signals presented in the previous section, is that they do not appear to very regular, i.e. they include components that seem to vary when experiments are repeated, giving the signals an "erratic" type of behaviour.
For smooth analytical signals, Fourier analysis is available to analyze the signals' properties, such as power/energy density over frequency, and/or decomposition in terms of frequency components (sine/cosine). However for "erratic" behaving signals, this analysis doe not seem to be the most appropriate, particularly in situations where experiment repetition is leading to different signals.
Consider for instance three measurements of the vertical acceleration of a car when it is driving over a particular road track. In Figure 1.9 three (repeated) experimental signals are shown of the same track.
The measurements are taken from the same physical variable, but due to all kind of irregularities the time realizations are different. Nevertheless the mechanical and physical mechanism behind the measured signals is considered to be the same.
When analyzing the three different time realizations, Fourier analysis will lead to different results for the several signals. The purpose of random signal analysis is to determine and to characterize the underlying principles and characteristic properties of the collection of signals.

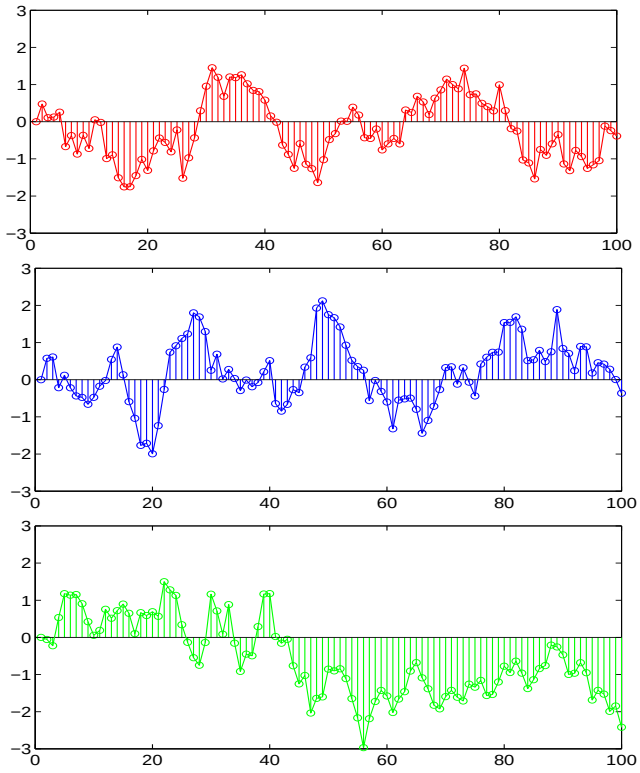

Figure 1.9: Three measurement signals obtained after experiment repetition.

As remarked above, for some physical phenomena we know that there is an intrinsic stochastic or probabilistic behaviour, as e.g. witnessed by Heisenberg's relation. From a macroscopic point of view, many physical phenomena can be modelled in different ways.

What is the probabilistic principle behind tossing a die?

If we toss a die, and we would be able to exactly specify the position, orientation and velocity of the die when it leaves our hand, and we would know exactly its material properties, the material properties of the surface on which we toss it, possibly also determine air movements that the die faces, then one could construct a first principles model that would exactly model the number of eyes that would be thrown in a particular nature. In this process there is nothing intrinsically stochastic or probabilistic! However, the detailed modelling of this situation would be extremely complex. And there is an alternative. By simply repeating the experiment a large number of times a probabilistic law can be empirically determined that also predicts the outcome of an experiment, just as the first principles model. However this prediction will not be a 100% sure prediction, but will have a probabilistic nature.

The example of tossing a die, can be replaced by any complex physical process or experiment, where one realizes that at a certain level one limits the modelling of the observed variables to a particular level of abstraction, and consequently some phenomena are discarded. Operational conditions are simplified, objects (as e.g. lenses) are considered to be ideal, temperatures are assumed to be constant, flow regimes are supposed to be known and fixed, etcetera. In all situations that the real physical world deviates from our model assumptions, perturbations will be involved, and will be measured, and these perturbations will generally have a non-structured format. I.e. they will not simply be periodic signals

or signals of any other reproducible form. They will typically be non-reproducible, in the sense that if the experiment or measurement is repeated, a different signal will result, from which it is nontrivial to determine in which sense it is related to the first measurement.
One of the examples to be mentioned here is the radial error signal of the CD optical unit, as depicted in figure 1.2. If the measurement is repeated some time later, under the same experimental conditions (in the eyes of the experimenter), then a different signal will appear.

For the modelling of non-reproducible signals a signals framework and modelling tool is needed that goes beyond the signal analysis tools treated in a basic (deterministic) signals and systems course, such as Fourier analysis (valid for periodic and finite-energy signals). A probabilistic model for signals will appear attractive, leading to stochastic processes. In this model exact representations of signals are replaced by "on average" properties, leading to a more rough description.
It has to be noted, that in many instances the deterministic or probabilistic character of signals is simply in the eye of the beholder. It is the modeler that determines how to interpret a particular measured signal. From the signal in figure 1.2 alone, there is no objective conclusion that the signal has a probabilistic nature. If the experiment is repeated, exactly the same signal could be possible, or something completely different could happen. However, in most cases a probabilistic model will appear attractive in modelling signals that deviate essentially when repeating experiments under the same conditions.

1.4 A global framework for signal analysis and estimation

1.4.1 Introduction

In order to put all different aspects of signal analysis and estimation in one picture, we consider a framework as schematically depicted in Figure 1.10. In this scheme the following notions are considered:

- $\mathbf{y}$ is the signal that is available from direct measurement with a sensor, resulting from some kind of physical/chemical/electrical process; it may be a voltage, a current, a temperature, a pressure, a number of photons, etcetera.

- $\mathbf{v}$ is a disturbance signal that affects $\mathbf{y}$ in some way. It causes $\mathbf{y}$ to be not exactly equivalent to the underlying phenomenon that one is interested in; it may be a measurement error, external disturbances that influence the measurement, but also (as is the case with the EEG signals) some unknown (non-controllable) source signal that generates the measurements $\mathbf{y}$.

- $\mathbf{u}$ is the experiment design variable; it is a possible excitation signal that is under control of an experimenter (it can be designed and applied to the real-world set-up). It may be the acoustic source signal that is used in acoustic imaging problems (the pulse signal or sinewave that is transmitted), or the electric current that steers the position of the actuator moving the position of the laser spot on a track of a DVD player.

- $M(\theta)$ is a set of (differential) equations that describes how the measurements $\mathbf{y}$ result from a particular excitation signal $\mathbf{u}$, under influence of disturbances $\mathbf{v}$. It is, what

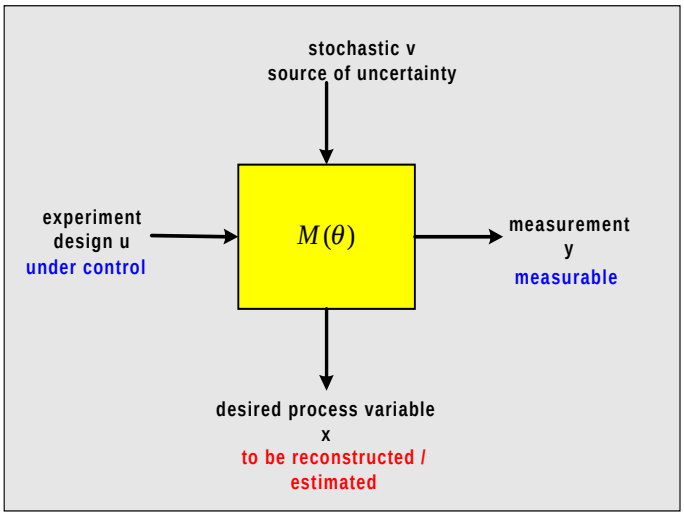

Figure 1.10: Framework for signal analysis and estimation.

is often called the "physical model" that underlies the observations $\mathbf{y}$. The variable θ reflects possible physical quantities (variables or parameters) that govern the physical model. E.g. in an acoustic imaging problem $M(\theta)$ reflects the differential equations that describe the reflection wave behavior of the object under test. $M(\theta)$ can be either known completely, or known in structure, while some variables (parameters) θ that govern M are unknown.

- $\mathbf{x}$ is the desired process variable; it is the signal that the user actually is interested in. It can be measured indirectly through the observation of $\mathbf{y}$, which is related to $\mathbf{x}$ through the physical model M.

For this general setup, a number of particular cases can be distinguished that will be discussed in more detail in particular parts of this course.

1.4.2 Situation 1: Detection of (known) signals in noise

In this case $\mathbf{x} = \mathbf{s}$, and the goal is to reconstruct a known signal $\mathbf{s}$ from noisy observations $\mathbf{y}$, see figure 1.11.

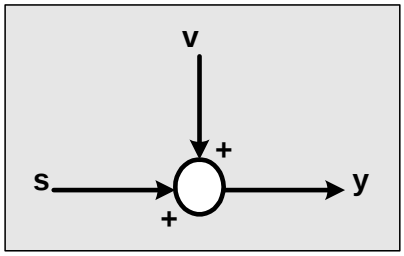

Figure 1.11: Detection of known signal in noise.

Example: transmitted and received radar signal.

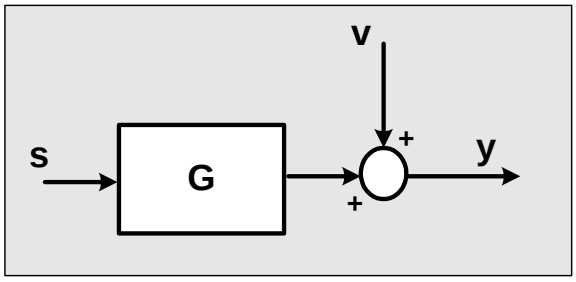

Figure 1.12: Signal reconstruction / filter problem.

1.4.3 Situation 2: Signal reconstruction / filter problem

In this situation, which is sketched in Figure 1.12, an information carrying signal $\mathbf{x} = \mathbf{s}$ that cannot directly be observed, is measured through a measurement variable $\mathbf{y}$ and the objective is to reconstruct/estimate of $\mathbf{s}$ on the basis of $\mathbf{y}$.

- Model G determines a (dynamical) relation between the (unknown) information carrying signal $\mathbf{s}$ and the noisy measurement variable $\mathbf{y}$.
- Objective is to determine an optimal estimation/reconstruction of $\mathbf{s}$ on the basis of measurement $\mathbf{y}$ (and knowledge of G)
- Required information may include a priori knowledge of the relation between $\mathbf{v}$ and $\mathbf{s}$.

Example: Noise cancellation in the cockpit of an airplane, as depicted in figure 1.13. The speech signal $\mathbf{s}$ of the captain is transmitted to the control tower under the influence of noise disturbance $\mathbf{v}_1$. By measuring, with a second sensor, a signal $\mathbf{v}_2$ that is related to the disturbance $\mathbf{v}_1$, an optimal filter can estimate $\mathbf{v}_1$ and correct the original signal $\mathbf{y}$ so as to reduce the noise level.

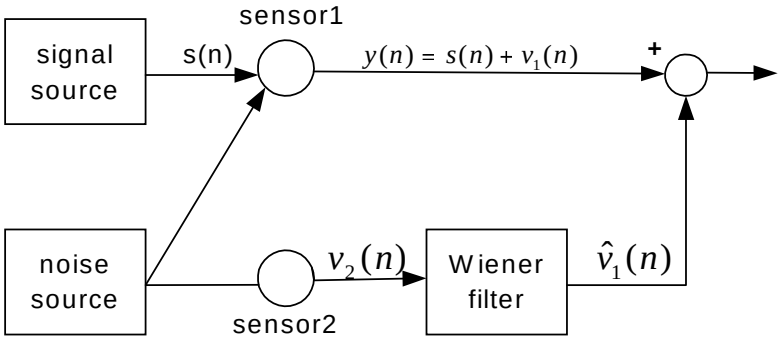

Figure 1.13: Noise cancellation.

The problem is most simple if $\mathbf{s}$ and $\mathbf{v}$ have their essential contributions in different frequency regions, because in that case a simple high-pass or low-pass filter could be applied.

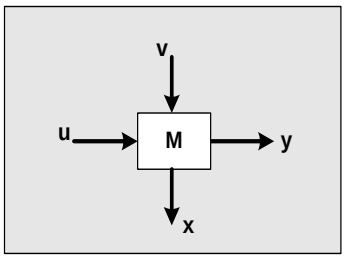

Figure 1.14: State reconstruction / filter problem

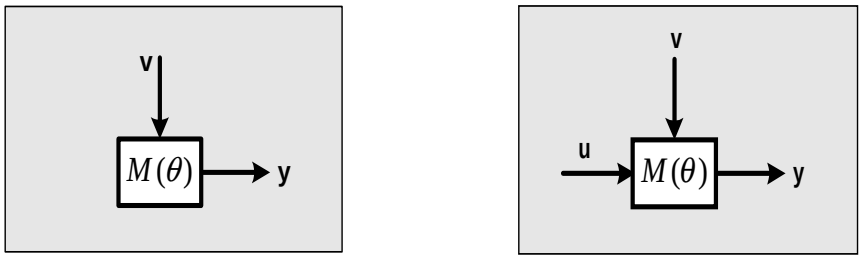

Figure 1.15: Signal modelling (left) as e.g. in EEG signals and in measurement problems with multiple (redundant) equipment; System modelling (right) as e.g. in dynamical systems as e.g. servo systems.

1.4.4 Situation 3: State reconstruction / filter problem

See figure 1.14.

- A known model M determines the relation between observation $\mathbf{y}$ and desired process variable $\mathbf{x}$;
- Determine an optimal estimate/reconstruction of $\mathbf{x}$ on the basis of observations $\mathbf{y}$ (and possibly $\mathbf{u}$)
- Required: knowledge of basic properties of disturbance $\mathbf{v}$.

Example: Construct a weather chart on the basis of pressure/temperature measurements at a restricted number of stations in Europe.

1.4.5 Situation 4: Parameter estimation / Identification

See figure 1.15.

- $\mathbf{y}$ is the result of a physical process that is not exactly known, i.e. there are unknown parameters
- Estimate θ on the basis of measurements $\mathbf{y}$ possibly in dedicated experiments introduced by exciting $\mathbf{u}$

1.4.6 Situation 5: Control design

In a feedback control system (figure 1.16) the influence of $\mathbf{v}$ on $\mathbf{x}$ needs to be reduced by way of feedback control, i.e. the actuation of the system through an input $\mathbf{u}$ that is constructed on the basis of measurements $\mathbf{y}$. In control system design this is called disturbance rejection.

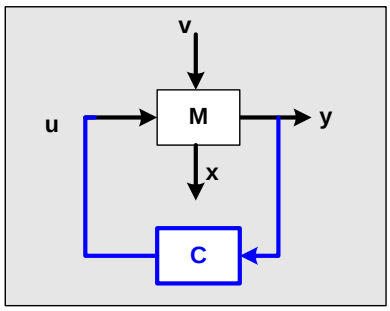

Figure 1.16: Control system design

In all these problems it may be evident that knowledge on the characteristic properties of the disturbance signal $\mathbf{v}$ plays a crucial role in the solutions. In this course the disturbance signal $\mathbf{v}$ will be modelled as a stochastic noise source. Note also that there is likely to be a clear relation between the problem of estimating $\mathbf{x}$ and that of estimating θ. In other words: there is some redundancy whether one wants to consider an unknown quantity as a process variables $\mathbf{x}$, or as a physical parameter θ.

1.5 Summary

Signals occur in very many engineering domains. Deterministic signal analysis tools as Fourier analysis are attractive to apply to deterministic signals, i.e. signals that typically are invariant under experiment repetition. When signals vary under experiment repetition, an alternative signal analysis framework may be more attractive, addressing the random nature of the measurements.
Several signal estimation, parameter estimation and reconstruction problems have been sketched in brief; they will be further explored and analyzed in subsequent chapters.

Chapter 2

Modeling of deterministic signals and systems

This chapter contains a brief review of discrete-time signals and systems analysis, including the formulation of Discrete Time Fourier Transforms, DFT and the representation of discrete-time linear systems/filters.

2.1 Introduction

In almost all situations of applied signal analysis and signal processing one will be dealing with sampled signals. These signals are easily treated by modern microprocessor-based equipment. The analysis of signals will particularly address questions as: "what is the frequency content of a signal?", and "how is the energy or power of a signal distributed over frequency?". (Frequency) transforms of signals play an important role in these questions. They will provide the necessary insight, but even more. For particular type of signals the frequency (Fourier) transform has a particular interpretation. The Fourier transform of the pulse response of a linear dynamical system is known to be equal to the frequency response of the system, and plays a crucial role in systems analysis and filtering.

The representation of signal properties, and the analysis of corresponding frequency transforms is the basic content of this chapter.

The treatment of the material will be done in a summarizing style rather than on an introductory level. It is assumed that the reader has a basic knowledge of signals and systems theory.

2.2 Discrete-time signals

For discrete-time signals the notation $u_d(k) := u(kT_s)$ will be adopted. In this expression u is the continuous-time signal that possibly underlies the discrete-time signal, and T_s is the sampling period. k is the (discrete) time variable that is an integer running index: $k = 1, 2, \cdots$. The sample frequency ω_s is defined by $\omega_s = \frac{2\pi}{T_s}$.

Deterministic signals appear in several different forms. As an illustration of this difference three different types of signals are depicted in Figure 2.1.

In specifying relevant signal properties we denote the *energy* of the signal u_d by

$$\mathcal{E}_u := \sum_{k=-\infty}^{\infty} u_d^2(k)$$

and the *power* of the signal by

$$\mathcal{P}_u := \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N-1} u_d^2(k).$$

Signals that have finite energy are referred to as energy signals , while signals with finite power are called power signals . In this way it can be verified that the signal in Figure 2.1(a) is an energy-signal, while the signals in (b) and (c) are power-signals.

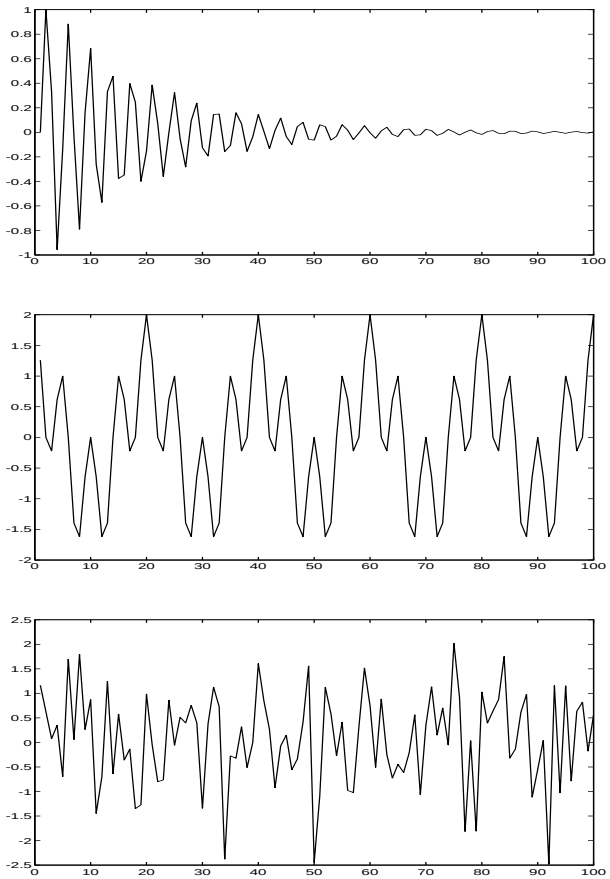

Figure 2.1: Three different types of signals: (a) finite-energy signal, (b) periodic finite-power signal, and (c) realization of a stationary stochastic process.

The basic tool for analyzing the frequency content of a signal is the Fourier analysis, i.e. the Fourier series and the Fourier transform. The Fourier series refers to periodic signals, showing that any periodic signal can be written as a summation of harmonic functions (sinusoids). The Fourier transform is a generalization that can also handle non-periodic signals.

Periodic signals
For a periodic signal with period N_0, which means that $u_d(k+N_0) = u_d(k)$ for all $k \in \mathbb{Z}$, the *Discrete-Time Fourier Series* of the signal u_d is given by[1]:

$$u_d(k) = \sum_{\ell=0}^{N_0-1} a_\ell e^{i\frac{2\pi}{N_0}\ell k} \tag{2.1}$$

where the Fourier coefficients are given by

$$a_\ell = \frac{1}{N_0} \sum_{k=0}^{N_0-1} u_d(k) e^{-i\frac{2\pi}{N_0}\ell k}. \tag{2.2}$$

The power of periodic signals can again be written directly as a function of the Fourier coefficients:

$$\mathcal{P}_u = \frac{1}{N_0} \sum_{k=0}^{N_0-1} u_d^2(k) = \sum_{\ell=0}^{N_0-1} |a_\ell|^2.$$

This shows that every exponential function in u has an independent contribution to the power of the signal, which is simply a summation of the contributions of each separate frequency. As the Fourier coefficients a_ℓ are periodic with period N_0, the sum on the right hand side can be taken over any N_0 consecutive values of ℓ.

2.3 Discrete-time Fourier Transform

The *Discrete-Time Fourier Transform* for sampled (discrete-time) signals is given by the transform pair:

$$U_s(\omega) := \sum_{k=-\infty}^{\infty} u_d(k) e^{-i\omega k T_s} \tag{2.3}$$

$$u_d(k) = \frac{T_s}{2\pi} \int_{2\pi/T_s} U_s(\omega) e^{i\omega k T_s} d\omega. \tag{2.4}$$

Note that the discrete-time Fourier transform (DTFT) transforms a discrete sequence of time-domain samples, into a function $U_s(\omega)$ that takes its values continuously over ω's. By construction (since k is integer valued) the transform $U_s(\omega)$ is a periodic function with period $2\pi/T_s = \omega_s$. Corresponding to this, the integral in (2.4) is taken over any range of ω with length $2\pi/T_s$, being the period length of the integrand.

Finite-time signals
When considering discrete time signals over a finite time, the corresponding Fourier transform is denoted by:

$$U_N(\omega) := \sum_{k=0}^{N-1} u_d(k) e^{-i\omega k T_s}. \tag{2.5}$$

[1]The given expression for $u_d(k)$ actually has resulted from $u(kT_s) = \sum_{\ell=0}^{N_0-1} a_\ell e^{i\frac{\omega_s}{N_0}\cdot\ell\cdot kT_s}$, which shows that the effect of the sampling interval T_s is cancelled out in the exponent.

Periodic signals
For a periodic signal with period N_0 the coefficients of the Fourier series can be directly related to a finite-time Fourier transform taken over one period of the periodic signal. Directly from (2.2) it follows that

$$a_\ell = \frac{1}{N_0} U_{N_0}(\ell\omega_0). \tag{2.6}$$

Additionally the expressions for the Fourier transform can be shown to be directly related to the Fourier series coefficients. Equating the Fourier series (2.1) with the inverse Fourier transform (2.4) it follows that for this periodic signal u the Fourier transform satisfies

$$U_s(\omega) = \frac{2\pi}{T_s} \sum_{k=-\infty}^{\infty} a_k \delta_c(\omega - k\omega_0) \tag{2.7}$$

with $\omega_0 = \frac{2\pi}{N_0 T_s}$. In this expression the δ-functions serve to replace the integral expression in (2.4).

Example 2.1 (Discrete-time Fourier transform of a sinusoid) Consider the signal

$$u(k) = A \cdot cos(\omega_0(kT_s))$$

with $\omega_0 = 2\pi/(N_0 T_s)$, i.e. there are N_0 samples in a single period of the signal. We consider N to be a multiple of N_0, $N = rN_0$, with $r \in \mathbf{N}$. Then

$$U_N(\omega) = \sum_{k=0}^{N-1} \frac{A}{2} \left[e^{i(\omega_0 - \omega)kT_s} + e^{-i(\omega_0 + \omega)kT_s} \right].$$

Using lemma 2A.1 it follows that

$$U_N(\omega) = \begin{cases} N \cdot \frac{A}{2} & \text{for } \omega = \pm\omega_0 = \pm\frac{2\pi}{N_0 T_s}, \\ 0 & \text{for } \omega = \frac{2\pi\ell}{NT_s}, \quad \ell \in \mathbb{Z}, \ell \neq r. \end{cases} \tag{2.8}$$

□

For a more extensive explanation see also example 2A.2.

Spectral densities of energy and power signals
Again, similar to the situation of continuous-time signals we can consider the distribution of energy and/or power of a signal over frequency.

Proposition 2.2 (Energy Spectral Density Function.) Let $u_d(k)$ be a finite-energy sampled-data signal, sampled with a sampling interval T_s. Then

$$\mathcal{E}_u = \frac{T_s}{2\pi} \int_{2\pi/T_s} \Psi_u(\omega) d\omega$$

where the Energy Spectral Density $\Psi_u(\omega)$ is given by

$$\Psi_u(\omega) = |U_s(\omega)|^2.$$

Proposition 2.3 (Power Spectral Density Function.) Let $u_d(k)$ be a finite-power sampled-data signal, i.e. $\frac{1}{N}\sum_{k=0}^{N-1}|u_d(k)|^2 = \mathcal{P}_u < \infty$, sampled with a sampling interval T_s. Then

$$\mathcal{P}_u = \frac{T_s}{2\pi}\int_{2\pi/T_s} \Phi_u(\omega)d\omega \tag{2.9}$$

where the Power Spectral Density $\Phi_u(\omega)$ is given by

$$\Phi_u(\omega) = \frac{1}{N}|U_N(\omega)|^2$$

The proof is added in the appendix.
As in the case of continuous-time signals, the (discrete-time) Fourier transform of sampled-data signals constitutes a way to characterize the distribution of energy and/or power of the corresponding signals over the different frequencies.
For finite power signals the quantity $\frac{1}{N}|U_N(\omega)|^2$ is referred to as the *periodogram* of the (finite-time) discrete-time signal. This periodogram determines the distribution of power over frequencies.

For periodic signals the power spectral density can again be computed directly on the basis of the discrete-time Fourier coefficients of the signals. Since in this case $\mathcal{P}_u = \sum_{\ell=0}^{N_0-1}|a_\ell|^2$ it follows from combination of (2.6) and (2.9) that

$$\Phi_u(\omega) = \frac{2\pi}{T_s}\sum_{k=-\infty}^{\infty}|a_k|^2\delta_c(\omega - k\omega_0).$$

Sampled-data signals and discrete-time signals
It is quite important to remark that discrete-time signals can be considered to be either originated from sampling continuous-time signals or simply as just a discrete-time sequence of numbers. In this section the first interpretation is followed, leading to the use of the sampling interval T_s in all expressions for Fourier transforms and the like. As a result, DTFT and spectral densities are formulated as functions of the continuous-time radial frequency ω. This allows a direct analysis of frequency properties of signals stated in terms that relate to their continuous-time equivalents.
However, in most situations of discrete-time signal analysis, the connection with a sampling mechanism is simply discarded, and indeed this is very well possible without loosing any information. When interpreting discrete-time signals just as discrete sequences of numbers, all results and notions introduced here still apply. This situation is best dealt with by introducing the variable transformation

$$\omega T_s \to \omega$$

which is actually equivalent to inserting $T_s = 1$.
Under this variable transformation all expressions for Fourier transforms and spectral densities are given in terms of the "discrete" radial frequency and the considered notions become periodic with a period length of 2π.

2.4 Discrete Fourier Transform

We now restrict attention for a moment to the situation of finite-time signals. For finite-time signals we have already mentioned the notation $U_N(\omega)$ as given in (2.5).
Actually this finite-time DTFT concerns the following transform pair:

$$U_N(\omega) = \sum_{k=0}^{N-1} u_d(k) e^{-i\omega k T_s}. \tag{2.10}$$

$$u_d(k) = \frac{1}{N} \sum_{\ell=0}^{N-1} U_N(\frac{\ell}{N}\omega_s) e^{i\frac{2\pi\ell}{N}k}. \tag{2.11}$$

A verification of the validity of this transform pair is added in the appendix. Considering this transform pair, a few remarks have to be made.

- Note that while $U_N(\omega)$ takes its values on a continuous domain of ω, only N discrete values of U_N are necessary for reconstructing the original signal u_d. These N discrete values are N points within one period of the periodic function $U_N(\omega)$.

- The DTFT is periodic with period $2\pi/T_s$.

- The sequence $\{U_N(\omega),\ \omega = \frac{\ell}{N}\omega_s,\ \ell = 0, \cdots N-1\}$ is defined as the *Discrete Fourier Transform* (DFT) of the signal $u_d(k)$, $k = 0, \cdots N-1$. It is given by

$$U_N(\frac{\ell\omega_s}{N}) = \sum_{k=0}^{N-1} u_d(k) e^{-i\frac{2\pi\ell}{N}k}, \quad \ell = 0, \cdots N-1.$$

- The DFT constitutes a one-to-one mapping from an N-length sequence of time-domain samples to an N-length sequence of frequency-domain samples, where the frequency domain samples are taken equidistantly in the frequency.

- The inverse DFT, defined by (2.11), also defines a time-domain sequence outside the interval $[0, N-1]$. Actually it induces a periodic extension of the original time-sequence $u_d(k)$, as the reconstructed signal (2.11) is periodic with period N.

- Because of reasons of symmetry, the DTFT (and also the DFT) satisfies

$$U_N(-\omega) = U_N(\omega)^*.$$

  As a result the DTFT is completely determined by $U_N(\omega)$ for ω in the interval $\omega \in [0, \pi/T_s]$. For the DFT this implies that the one-to-one mapping between time- and frequency domain actually takes place between N real-valued time-domain samples, and $N/2$ complex-valued frequency domain samples.

In very many situations discrete-time signals are being analyzed without taking account of the fact that they originate from sampled continuous-time signals. Similar to the situation of the previous section, this implies that in that case the expressions for the DTFT are used

for $T_s = 1$:

$$U_N(\omega) = \sum_{k=0}^{N-1} u_d(k) e^{-i\omega k}. \tag{2.12}$$

$$u_d(k) = \frac{1}{N} \sum_{\ell=0}^{N-1} U_N(\frac{2\pi\ell}{N}) e^{i\frac{2\pi\ell}{N}k}. \tag{2.13}$$

In many books on discrete-time signal processing this is the only situation that is considered. Discrete-time Fourier transforms, spectral densities, periodograms will then be considered generally over the frequency interval $\omega \in [0, \pi]$, being half of a single period of the corresponding periodic function in the frequency domain. Whenever we connect a sampling time to the discrete-time signal, then $\omega = \pi$ gets the interpretation of being equal to half of the (radial) sampling frequency.

Spectral properties of finite-time sampled signals
Similar to the situation of infinite-time signals, we can exploit Parsseval's relation for quantifying the energy and power of finite-time (deterministic) sampled signals.
Consider the Discrete Fourier Transform as discussed above. Then

$$\sum_{k=0}^{N-1} u_d(k)^2 = \frac{1}{N} \sum_{k=0}^{N-1} |U_N(\frac{k\omega_s}{N})|^2 \tag{2.14}$$

$$\frac{1}{N} \sum_{k=0}^{N-1} u_d(k)^2 = \frac{1}{N} \sum_{k=0}^{N-1} |\frac{1}{\sqrt{N}} U_N(\frac{k\omega_s}{N})|^2. \tag{2.15}$$

It may be clear that the first expression is used for signals having the character of having finite energy, while the second expression is specially used for finite power signals. Note that over a finite time-interval this distinction is not really relevant as the operation of dividing by a finite N is just a matter of scaling. The main difference has to be found in the corresponding asymptotic analysis, when $N \to \infty$. Note that the expressions above actually are alternatives for the integral expressions for signal power as presented in Proposition 2.3. For finite time signals, there is no need to take the integral over the power spectral density as in (2.9); the power also results from summing the squared magnitude of the DFT over an equidistant frequency grid.

2.5 Discrete-time systems analysis

General notation

Linear time-invariant (LTI) discrete-time dynamical systems specify (dynamical) relations between input and output signals. They can be represented in several forms, one of which is the convolutional representation:

$$y(k) = \sum_{\ell=-\infty}^{\infty} g(\ell) u(k-\ell), \tag{2.16}$$

where the sequence $\{g(k)\}_{\ell=-\infty,\cdots,\infty}$ is the *pulse response* of the system, and $y(k)$ and $u(k)$ are the discrete-time output and input signal of the system, i.e. $k \in \mathbb{Z}$.

To facilitate this description, two shift operators will be used; the *forward shift operator* q:

$$qu(k) = u(k+1)$$

and the *backward shift operator* q^{-1}:

$$q^{-1}u(k) = u(k-1)$$

Using these operators it follows that

$$y(k) = \sum_{\ell=-\infty}^{\infty} g(\ell)(q^{-\ell}u(k)) = G(q)u(k)$$

with

$$G(q) = \sum_{\ell=-\infty}^{\infty} g(\ell)q^{-\ell}.$$

With slight abuse of notation, we will also refer to $G(q)$ as the *transfer function* of the system. Strictly speaking, however, the transfer function is defined by the complex function $G(z)$:

$$G(z) = \sum_{\ell=-\infty}^{\infty} g(\ell)z^{-\ell} \tag{2.17}$$

where z is a complex indeterminate.
A system is called bounded-input bounded-output stable (BIBO-stable) if every input signal with bounded amplitude generates an output signal with bounded amplitude. This property is reflected by the condition that the pulse response is absolutely summable, that is

$$\sum_{k=-\infty}^{\infty} |g(k)| < \infty. \tag{2.18}$$

It implies that the series expansion (2.17) is convergent for $|z| = 1$ (on the unit circle).
In this book, attention will be restricted to linear time-invariant systems having a *rational* transfer function. That is, $G(z)$ can be written as a rational of polynomials in z. We note in particular that the transfer function of a system satisfying a linear constant-coefficient difference equation is always rational. As a result $G(z)$ can be written as a fraction of polynomials in z: $G(z) = \frac{b(z)}{a(z)}$ and the transfer function can be characterized by the *zeros* z_i, determined by $b(z_i) = 0$, and the *poles* p_i, determined by $a(p_i) = 0$.

Causal systems
Usually attention will be restricted to causal systems, i.e. systems for which holds that $g(\ell) = 0,\ \ell < 0$, so that the transfer function satisfies

$$G(z) = \sum_{\ell=0}^{\infty} g(\ell)z^{-\ell}. \tag{2.19}$$

Causal systems are stable if the poles of the related transfer function are located strictly inside the unit circle. This implies that the series expansion (2.19) converges on the domain

$|z| \geq 1$, i.e. on and outside the unit circle in the complex plane. Causality of a stable system can therefore also be checked by verifying that $\lim_{|z|\to\infty} G(z)$ exists and is bounded. Additionally the transfer function of a causal system will be called *monic* if $g(0) = 1$.

Noncausal systems
Although it seems unnatural in a real engineering world, we sometimes have to deal with noncausal systems also, as well as with their related transfer functions. Their representation is as in (2.16) with a possible two-sided pulse response $\{g(k)\}_{\ell=-\infty,\cdots,\infty}$ and a transfer function

$$G(z) = \sum_{\ell=-\infty}^{\infty} g(\ell) z^{-\ell}. \tag{2.20}$$

For any rational transfer function $G(z)$ there exists a unique expansion (2.20) that is convergent on the unit circle, i.e. $\sum_{\ell=-\infty}^{\infty} |g(\ell)| < \infty$, and therefore is representing a stable (noncausal) system.
We define the transfer function operation $[\cdot]_+$ as taking the causal part of this unique expansion of the transfer function:

$$G_+(z) := \sum_{\ell=0}^{\infty} g(\ell) z^{-\ell} \tag{2.21}$$

and the operation $[\cdot]_-$ as taking the noncausal part:

$$G_-(z) := \sum_{\ell=-\infty}^{-1} g(\ell) z^{-\ell} \tag{2.22}$$

so that $G(z) = G_+(z) + G_-(z)$.
The system represented by $G_+(z)$ is causal and stable by construction, and will have all its poles strictly inside the unit circle. The expansion (2.21) is convergent in $|z| \geq 1$.
The system represented by $G_-(z)$ is anti-causal and stable. It will have all poles strictly outside the unit circle, while the expansion (2.22) is convergent in $|z| \leq 1$.
For more details see also Appendix B.5 and Oppenheim and Willskey (1997).

Frequency response and Bode plots The frequency response of a discrete-time system is determined by the output of the system when excited with a sinusoidal input signal. Consider the signal:

$$u(k) = cos(\omega k) = Re\{e^{i\omega k}\}.$$

The output of the (causal) discrete-time system is given by

$$\begin{aligned} y(k) &= \sum_{\ell=0}^{\infty} g(\ell) u(k-\ell) = Re\{e^{i\omega k} \sum_{\ell=0}^{\infty} g(\ell) e^{-i\omega\ell}\} && (2.23) \\ &= Re\{e^{i\omega k} G(e^{i\omega})\}. && (2.24) \end{aligned}$$

Consequently

$$y(k) = |G(e^{i\omega})| \cdot cos(\omega k + \phi) \tag{2.25}$$

with $\phi = \arg[G(e^{i\omega})]$.

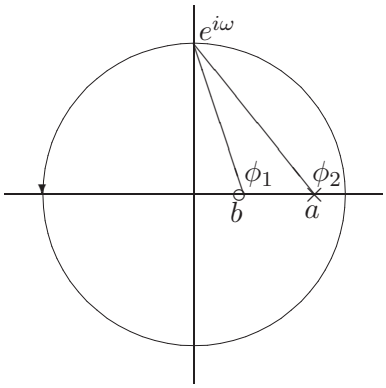

Figure 2.2: Zero/pole location and evaluation of frequency response of first order discrete-time system

The complex-valued function $G(e^{i\omega})$ is referred to as the *frequency response* of the discrete-time system. It evaluates the transfer function in the complex plain over the unit circle $z = e^{i\omega}$. Note the difference with a continuous-time system, where the frequency response is reflected by the evaluation of $G(s)$ over the imaginary axis $s = i\omega$.
For discrete-time systems with a real-valued pulse response (and thus real-valued coefficients) it holds that $G(e^{-i\omega}) = G(e^{i\omega})^*$ and so for reasons of symmetry, full information on the frequency response of the system is obtained by $G(e^{i\omega})$ for $\omega \in [0, \pi]$. In Figure 2.2 this is illustrated for a first order system, given by

$$G(z) = \frac{z-b}{z-a}. \tag{2.26}$$

By writing the expression

$$G(e^{i\omega}) = \frac{e^{i\omega}-b}{e^{i\omega}-a} \tag{2.27}$$

it follows that

$$|G(e^{i\omega})| = \frac{|e^{i\omega}-b|}{|e^{i\omega}-a|} \quad \text{and} \tag{2.28}$$

$$\arg G(e^{i\omega}) = \arg(e^{i\omega}-b) - \arg(e^{i\omega}-a). \tag{2.29}$$

The first equation generates the amplitude Bode plot, whereas the second defines the phase Bode plot. For the considered first order system these Bode plots are given in Figure 2.3, where the values $b = 0.3$ and $a = 0.8$ are chosen. Note that the frequency function is given for frequencies up to $\omega = \pi$.
In this discrete-time case, unlike the situation in the continuous-time case, there is no asymptotic point in frequency where ω tends to. Note that in the discrete-time case the phase contribution of every (real) zero varies between 0^0 and $+180^0$, while the contribution of each (real) pole varies between 0^0 and -180^0. For a complex conjugate pair of zeros/poles it can simply be verified that the contribution to the phase in $\omega = \pi$ is given by respectively $+360^0$ and -360^0.

Signal properties processed by discrete-time systems.

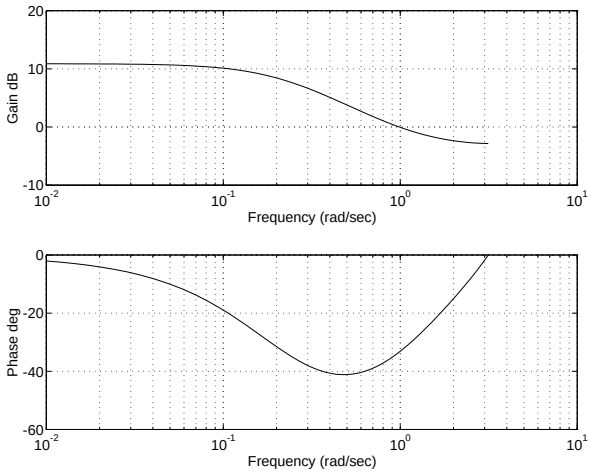

Figure 2.3: Bode-amplitude and Bode-phase plot of first order discrete-time system

Similar to the situation of continuous-time systems, it can be verified that when given an input/output system $y(k) = G(q)u(k)$, and u and y quasi-stationary, the following relations hold:

- $\Phi_{yu}(\omega) = G(e^{i\omega})\Phi_u(\omega)$;
- $\Phi_y(\omega) = |G(e^{i\omega})|^2 \cdot \Phi_u(\omega)$;

If u is a deterministic sequence for which the DTFT exists, then additionally, under the assumption of zero initial conditions (i.e. $u(t) = 0$, $t < 0$):

- $Y(\omega) = G(e^{i\omega})U(\omega)$.

2.6 Overview of Fourier Transforms

	Transform	Inverse transform
CTFT	$U(\omega) = \int_{-\infty}^{\infty} u(t)e^{-i\omega t}dt$	$u(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} U(\omega)e^{i\omega t}d\omega$
DTFT	$U_s(\omega) := \sum_{k=-\infty}^{\infty} u_d(k)e^{-i\omega kT_s}$	$u_d(k) = \frac{T_s}{2\pi}\int_{2\pi/T_s} U_s(\omega)e^{i\omega kT_s}d\omega$
DTFT, $T_s = 1$	$U_s(\omega) := \sum_{k=-\infty}^{\infty} u_d(k)e^{-i\omega k}$	$u_d(k) = \frac{1}{2\pi}\int_{2\pi} U_s(\omega)e^{i\omega k}d\omega$
DFT	$U_N(\frac{\ell\omega_s}{N}) = \sum_{k=0}^{N-1} u_d(k)e^{-i\frac{2\pi\ell}{N}k}$	$u_d(k) = \frac{1}{N}\sum_{\ell=0}^{N-1} U_N(\frac{\ell\omega_s}{N})e^{i\frac{2\pi\ell}{N}k}$

Table 2.1: Summary of Fourier transforms; CTFT = Continuous-Time Fourier Transform; DTFT = Discrete-Time Fourier Transform; DFT = Discrete Fourier-Transform

2.7 Summary

In this section a brief review and summary of the basic concepts in deterministic discrete-time signals and systems analysis has been presented. For a detailed treatment of the fundamentals the reader is referred to more specialized textbooks, as e.g. Oppenheim and Willsky (1997).

Appendix

Lemma 2A.1 Let $k \in \mathbb{Z}$ and $N \in \mathbf{N}$. Then

$$\begin{aligned} \sum_{\ell=0}^{N-1} e^{i\frac{2\pi\ell}{N}k} &= N \quad \text{for } k = 0, \\ &= 0 \quad \text{for } k \neq 0. \end{aligned}$$

Proof: The proof follows simply by applying the finite sum formula: $\sum_{\ell=0}^{N-1} a^\ell = \frac{1-a^N}{1-a}$ for $a \in \mathbb{C}$, $a \neq 1$. □

Example 2A.2 (DTFT of a periodic signal) Let $u(t)$ be a periodic signal with length $N = rN_0$ and basic period N_0. Then

$$\begin{aligned} U_N(\omega) &= \sum_{t=0}^{rN_0-1} u(t)e^{-i\omega t} \\ &= \sum_{\ell=1}^{r}\sum_{m=0}^{N_0-1} u(m)e^{-i\omega[(\ell-1)N_0+m]} \\ &= \sum_{\ell=1}^{r} e^{-i\omega(\ell-1)N_0} \sum_{m=0}^{N_0-1} u(m)e^{-i\omega m}. \end{aligned} \tag{2A.1}$$

Since with Lemma 2A.1

$$\begin{aligned} \sum_{\ell=1}^{r} e^{-i\omega(\ell-1)N_0} &= r \quad \text{for } \omega = \frac{2\pi k}{N_0},\ k = 1, \cdots N_0, \\ &= 0 \quad \text{for } \omega = \frac{2\pi j}{N},\ j = 1, \cdots, N; j \neq r, 2r, \cdots N_0 r, \end{aligned}$$

it follows that

$$U_N(\omega) = r \cdot U_{N_0}(\omega), \tag{2A.2}$$

for $\omega = \frac{2\pi k}{N_0},\ k = 1, \cdots N_0$, where $U_{N_0}(\omega)$ is the DTFT of u over one period of the signal, and $U_N(\omega)$ will be 0 at frequencies outside this grid but being a member of the grid $\{\omega = \frac{2\pi k}{N}, k = 1, \cdots N\}$.
For $N \to \infty$ and N_0 finite, it follows that $U_N(\omega) = 0$ almost everywhere, except in the frequencies $\omega = \frac{2\pi k}{N_0}$ where the DTFT grows with r.

Lemma 2A.3 Let u be a signal defined on the time interval $[0, N-1]$. Consider the sample covariance,

$$\begin{aligned} \hat{R}_u^N(\tau) &:= \frac{1}{N}\sum_{t=0}^{N-1} u(t)u(t-\tau), \qquad |\tau| \leq N-1 \\ &:= 0 \qquad\qquad\qquad\qquad\qquad\quad |\tau| \geq N. \end{aligned} \tag{2A.3}$$

Then the Discrete-Time Fourier Transform of this sample covariance satisfies:

$$\sum_{\tau=-\infty}^{\infty} \hat{R}_u^N(\tau)e^{-i\omega\tau} = \frac{1}{N}|U_N(\omega)|^2, \tag{2A.4}$$

with $U_N(\omega)$ the DTFT of the signal, $U_N(\omega) = \sum_{t=0}^{N-1} u(t)e^{-i\omega t}$.

Proof: Substituting the definitions it follows that

$$\sum_{\tau=-\infty}^{\infty} \hat{R}_u^N(\tau)e^{-i\omega\tau} = \frac{1}{N}\sum_{t=-\infty}^{\infty}\sum_{\tau=-\infty}^{\infty} u(t)u(t-\tau)e^{-i\omega\tau}$$

with $u(t) := 0$ for $t < 0$ or $t \geq N$. By variable substitution $\ell = t - \tau$ we arrive at

$$\frac{1}{N}\sum_{t=-\infty}^{\infty} u(t)e^{-i\omega t} \cdot \sum_{\ell=-\infty}^{\infty} u(\ell)e^{i\omega\ell}$$

which equals $\frac{1}{N}U_N(\omega)U_N(\omega)^*$. □

Proof of Proposition 2.3.

Starting from $\frac{1}{N}\sum_{k=0}^{N-1} u_d(k)^2 = \hat{R}_u^N(0)$ and using the result of Lemma 2A.3 that $\frac{1}{N}|U_N(\omega)|^2$ is the DTFT of $\hat{R}_u^N(\tau)$, it follows that

$$\hat{R}_u^N(\tau) = \frac{T_s}{2\pi}\int_{-\pi/T_s}^{\pi/T_s} \frac{1}{N}|U_N(\omega)|^2 e^{i\omega\tau}d\omega$$

showing that

$$\hat{R}_u^N(0) = \frac{T_s}{2\pi}\int_{-\pi/T_s}^{\pi/T_s} \frac{1}{N}|U_N(\omega)|^2 d\omega.$$

Proof of the DFT-pair (2.10)-(2.11).

Substituting the expression (2.10) into (2.11) shows that

$$u_d(k) = \frac{1}{N}\sum_{\ell=0}^{N-1}\sum_{m=0}^{N-1} u_d(m)e^{-i\frac{\ell\omega_S}{N}mT_s}e^{i\frac{2\pi\ell}{N}k} \tag{2A.5}$$

$$= \frac{1}{N}\sum_{m=0}^{N-1} u_d(m)\cdot\sum_{\ell=0}^{N-1} e^{i\frac{2\pi\ell}{N}(k-m)}. \tag{2A.6}$$

With Lemma 2A.1 the sum of exponentials will equal $N\delta(k-m)$ [2], which proves the validity of the transform pair.

[2] $\delta(\cdot)$ is the discrete pulse function, i.e. $\delta(k) = 1$ for $k = 0$ and $\delta(k) = 0$ elsewhere.

Chapter 3

Modelling of Stochastic Signals

Signals that have an inherent random nature are described by using random variables and stochastic processes. The necessary tools for describing and analyzing these signals are presented. After defining the basic concepts of stochastic processes, attention will be given to correlation properties (in time domain) as well as to spectral properties (in frequency domain).

3.1 Introduction

For many signals it is not appropriate to describe their properties in terms of a fixed and given amplitude at every time instant, as e.g. a sinusoid or an exponential function. This is due to the fact that many signals that result from measurements are not exactly reproducible: if we measure the same object at a different moment, a signal results that is not exactly the same as the signal measured in first instance.

When studying the suspension properties of a particular car, consider that we drive it over a particular highway track and measure the vertical acceleration with an acceleration sensor. When repeating the experiment over the same track a different acceleration signal will generally result, although some underlying characteristics of the signal may be invariant. See e.g. examples of three measurements given in Figure 3.1. The difference in the signals is caused by several sources of "uncertainty" (the exact horizontal position on the track, effects of wind, initial driving conditions of the car, etcetera).

There are two options for describing (modelling) the difference between the several measured signals. One can try to model all the effects mentioned above by using (physical) knowledge of the related phenomena, and thus removing all of the "uncertainty" from the signal model. In many cases this will not be very realistic. Alternatively one can apply a less "specific" description of the signals that are measured, by considering them to be generated by random variables. As a result a statistical way of describing the several signals is introduced.

Random signals are not described by a fixed and given amplitude at every time instant but only by a probability density function that characterizes the probability that a certain signal value will result.

In this chapter random or stochastic processes are introduced to describe signals, and related modeling tools for these processes are discussed, such as correlation and covariance functions and spectral density functions.

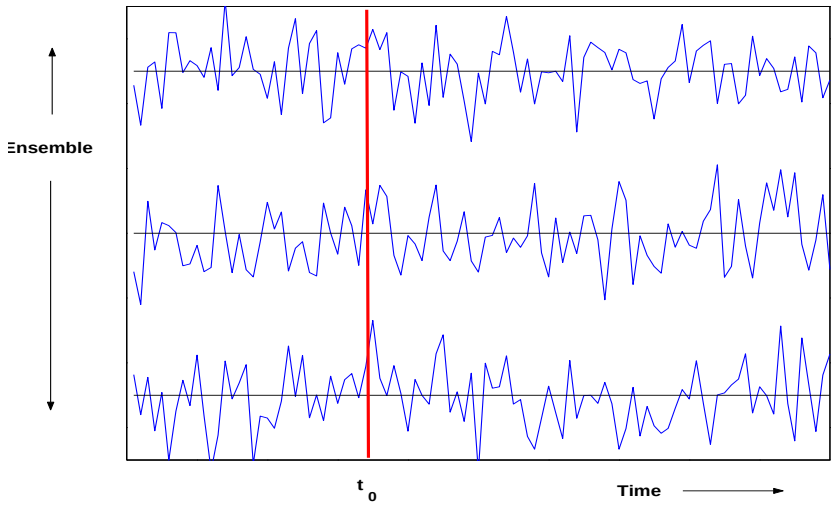

Figure 3.1: Three realizations of a random process.

3.2 From random variables to stochastic processes

Random Processes
A random process generates time sequences (signals) $\{x(t)\}$ as result of a random experiment. Every outcome in the sample space (also referred to as a realization) from the ensemble is a single time sequence.
The random process will formally be denoted by $\mathbf{x}(t, \Lambda)$, and is actually described in two dimensions, where t represents time, and Λ reflects the outcomes in the sample space (the ensemble). In this way one can distinguish several quantities:

- $\mathbf{x}(t, \Lambda)$ is a collection of time signals $\{x_1(t)\}$, $\{x_2(t)\}$, etc.
- $\mathbf{x}(t, \lambda_i)$ is a particular time signal $\{x_i(t)\}$ (one realization)
- $\mathbf{x}(t_0, \Lambda)$ is a collection of signal values at time t_0, $\{x_1(t_0), x_2(t_0), \cdots\}$, i.e., a random variable;
- $\mathbf{x}(t_0, \lambda_i)$ $x_i(t_0)$, a particular numerical value of x_i at time t_0.

As a general notation $\mathbf{x}(t)$ will be used to denote the random (stochastic) process, presuming that there exists an associate set Λ that underlies the generation of realizations from an ensemble.

Example 3.1 Consider an ensemble of 6 signals being given by

$$x_1(t) = -4; \quad x_2(t) = -2; \quad x_3(t) = +2; \quad x_4(t) = +4; \quad x_5(t) = -t/2; \quad x_6(t) = t/2;$$

being the result of a random experiment, where $x_i(t)$ is the result of $\mathbf{x}(t, \lambda_i) := \mathbf{x}(t, \Lambda = i)$. The sample space is given by $\{1, 2, 3, 4, 5, 6\}$. The six signals are sketched in Figure 3.2. Note that $x_4(t)$ and $x_6(t)$ have clearly different properties in time.

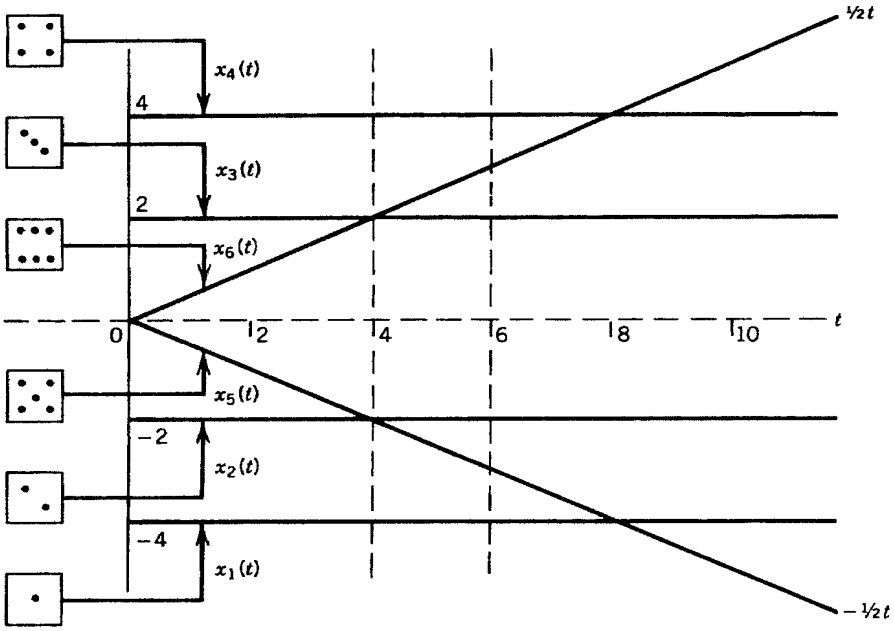

Figure 3.2 Example of a random process.

Figure 3.2: Example of a random process with a finite number of 6 signals in the ensemble, that each have equal probability.

In this example:
$\mathbf{x}(t = 6, \Lambda = 5) = -3$ is a constant;
$\mathbf{x}(6, \Lambda) = \mathbf{x}(6)$ is a random variable with values taken from the set $\{-4, -3, -2, 2, 3, 4\}$.

Problem 3.2 Assume that the probability of each of the 6 signals in the previous example is 1/6. Verify that for all values of t: $\mathbb{E}[\mathbf{x}(t)] = 0$.

Processes where t is *discrete* are sometimes referred to as random *sequences* ;
Processes where $\mathbf{x}(t)$ takes on values in a continuous domain (e.g. the real line) are called *continuous*, whereas processes where $\mathbf{x}(t)$ takes values in a finite (discrete) set are called *discrete*. Additionally $\mathbf{x}(t)$ can be either real- or complex valued.
In this course main attention will be given to real-valued continuous random sequences.

Description of a Random Process

Consider a random process $\mathbf{x}(t, \Lambda)$ with $t \in \Gamma \subset \mathbb{R}$. For every value of t, $\mathbf{x}(t, \Lambda)$ is a random variable that is characterized by a particular probability density function (pdf) $f_{\mathbf{x}}(x)$, that determines the distribution function:

$$F_{\mathbf{x}}(a) := P(\mathbf{x} \leq a) = \int_{-\infty}^{a} f_{\mathbf{x}}(x)dx,$$

and therefore completely describes the probabilistic properties of the random variable $\mathbf{x}$ at time instant t_1. This mechanism is illustrated in Figure 3.3, where the pdf of $\mathbf{x}$ at one particular time instant t_1 is sketched in the left part of the figure.

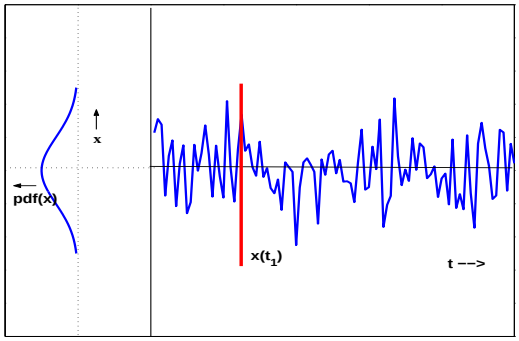

Figure 3.3: Realization of a random process, and indication of the pdf of $\mathbf{x}$ at one particular time instant t_1, indicated by the red line.

In first instance it may look sufficient to fully characterize the random process by its distribution function at every time instant t. However such a description does not take into account the possible interrelation between $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$. When considering the random process in Figure 3.4, one of the questions to be addressed is how the random variables $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ are statistically related to each other.

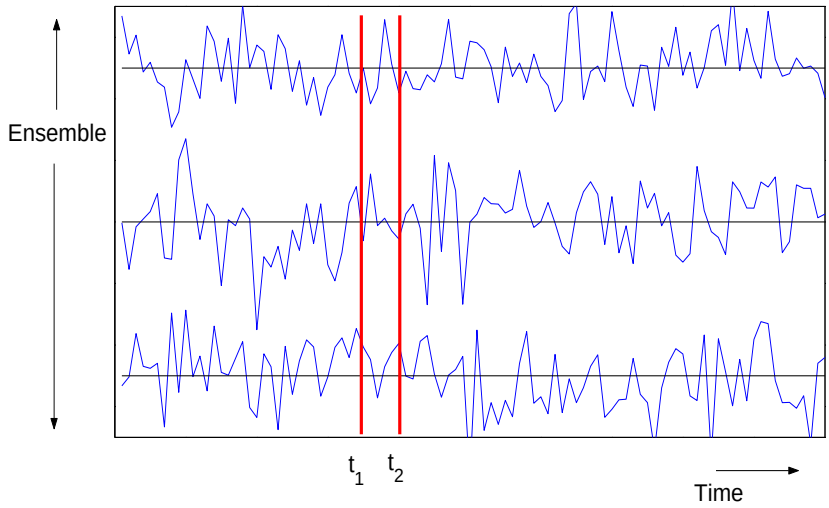

Figure 3.4: For properly characterizing the properties of a random process, it is necessary to specify also the statistical relationship between the random variables $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ for different values of t_1, t_2.

In order to illustrate that this statistical relation can contain important characteristics of the process, consider Figure 3.5 where single realizations are sketched of two different stochastic processes.
One of the apparent differences between the stochastic processes in this figure is that in the left picture the signal at a time t seems to be dependent on the signal value in the neighborhood of t (there is correlation in time), whereas the right picture shows a signal that seems to be more "noisy" and therefore in its time behaviour less correlated with neighboring time samples.
A full characterization of the stochastic process is obtained when considering $\mathbf{x}(t)$ to be

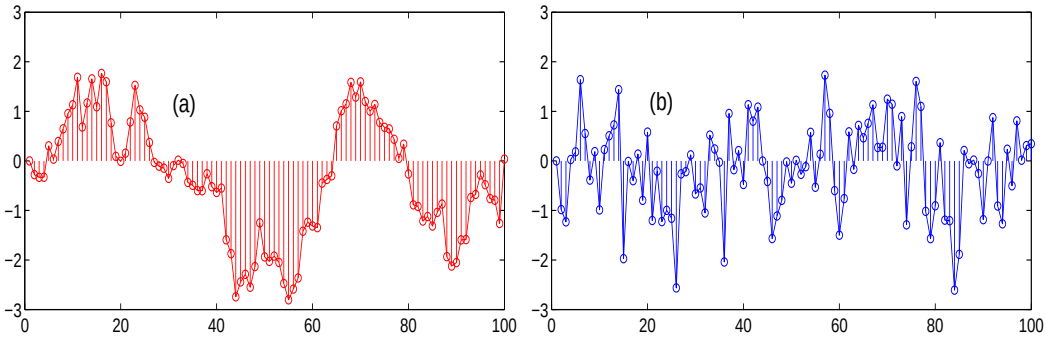

Figure 3.5: Realizations of two different stochastic processes

a random variable at every time instant t, while the full statistical behaviour of $\mathbf{x}(t,\Lambda)$ is described by all joint distribution functions of any selection of n time moments $t_1,..,t_n$ and any value of n. As a result, the stochastic process will be fully characterized by *all* n^{th} order distribution functions

$$F_{\mathbf{x}(t_1),\mathbf{x}(t_2),\cdots,\mathbf{x}(t_n)}(x_1,x_2,\cdots,x_n) = \\ P\left[\mathbf{x}(t_1)\leq x_1,\mathbf{x}(t_2)\leq x_2,\cdots\mathbf{x}(t_n)\leq x_n\right] \quad \forall n \text{ and } t_1,\cdots t_n\in\Gamma \tag{3.1}$$

In the continuous-time situation this is only of theoretical interest, as all possible combinations of n time instants have to be considered for all values of n. This is truly an impossible task. For the discrete-time case, however, there is one time vector: $[t_1,t_2,\cdots t_n]$, and the single distribution function $F_{\mathbf{x}(t_1),\mathbf{x}(t_2),\cdots,\mathbf{x}(t_n)}(x_1,x_2,\cdots,x_n)$ fully describes all statistical properties of the random process.
For many purposes it will suffice to consider the first and second order distribution functions, i.e.,

$$P\left[\mathbf{x}(t_i)\leq x_i\right] \qquad \text{and} \tag{3.2}$$

$$P\left[\mathbf{x}(t_i)\leq x_i,\mathbf{x}(t_j)\leq x_j\right]. \tag{3.3}$$

The statistical behaviour of the process is then considered to be characterized by these first two distribution functions, that describe the interrelation between any two given time instants. And an ever more crude characterization, but still a practical one, is when one only considers the expected values related to the first and second order distribution functions, i.e., the mean and the correlation. This situation is considered in the next section.

3.3 Autocorrelation and autocovariance functions

Basic properties of a random variable are reflected in their mean value and in their variance. A simple extension of these quantities to the first and second order distributions of random processes, leads to the notions of mean value and autocorrelation/autocovariance functions, as formally defined next.

Mean.
The mean value of $\mathbf{x}(t)$ is the expected value of the random variable $\mathbf{x}(t)$:

$$\mu_{\mathbf{x}}(t) := \mathbb{E}[\mathbf{x}(t)]. \tag{3.4}$$

By construction, the mean value is dependent on n.

Autocorrelation:

$$R_{\mathbf{x}}(t_1, t_2) := \mathbb{E}[\mathbf{x}(t_1)\mathbf{x}^*(t_2)] \tag{3.5}$$

where $\mathbf{x}^*$ is the complex conjugate of $\mathbf{x}$, for the situation that $\mathbf{x}$ is complex-valued.

Autocovariance:

$$C_{\mathbf{x}}(t_1, t_2) := \mathbb{E}\{(\mathbf{x}(t_1) - \mu_{\mathbf{x}}(t_1))(\mathbf{x}^*(t_2) - \mu^*_{\mathbf{x}}(t_2))\} \tag{3.6}$$

The autocorrelation is basically the correlation (A.6) between the random variables $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$. With varying t_1 and t_2 the autocorrelation function becomes a two-dimensional function. Similarly the autocovariance function is simply the covariance (A.7) between the random variables $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$. For processes that have zero mean, i.e., $\mu_{\mathbf{x}}(t) = 0$ for all t, autocorrelation and autocovariance functions are the same.
The correlation and covariance functions partially describe the time domain behaviour of the underlying stochastic process. They will be used later to derive the spectral (frequency domain) properties of the process.
Note that for $t_1 = t_2 = t$, $C_{\mathbf{x}}(t, t) = var[\mathbf{x}(t)] = \sigma^2_{\mathbf{x}(t)}$.

Random signals do not necessarily look "noisy". This is illustrated in the following example.

Example 3.3 In many applications in communication systems carrier signals are defined by sinusoidal signals of a particular (carrier) frequency and with random amplitude and phase. As an example consider the random process

$$\mathbf{x}(t) = \mathbf{a} \cdot cos(100t + \mathbf{b}) \tag{3.7}$$

with $\mathbf{a}$ a normally (Gaussian) distributed random variable with mean value 0 and variance 1, $\mathbf{b}$ uniformly distributed in $[-\pi, \pi]$, and $\mathbf{a}, \mathbf{b}$ independent random variables.
This implies that every realization of this process has the character of a cosine-signal, however with an amplitude and phase that are random. Five different realizations of this random process are sketched in Figure 3.6.
Here we address the question to analyze $\mu_{\mathbf{x}}(t)$ and $R_{\mathbf{x}}(t, t - \tau)$.

Solution.

$$\mathbb{E}[\mathbf{x}(t)] = \int_{-\infty}^{\infty} \int_{-\pi}^{\pi} \mathbf{a} \cdot cos(100t + \mathbf{b}) f_{\mathbf{a},\mathbf{b}}(a, b) da db$$

Since $\mathbf{a}$ and $\mathbf{b}$ are statistically independent, their joint pdf is equal to the product of the separate pdf's, i.e., $f_{\mathbf{a},\mathbf{b}}(a, b) = f_{\mathbf{a}}(a) \cdot f_{\mathbf{b}}(b)$, and as a result

$$\mathbb{E}[\mathbf{x}(t)] = \mathbb{E}[\mathbf{a}] \cdot \mathbb{E}\left[(cos(100t + \mathbf{b}))\right] = 0.$$

$$\begin{aligned} R_{\mathbf{x}}(t, t - \tau) &= \mathbb{E}\left[\mathbf{a}^2 cos(100t + \mathbf{b}) cos(100t - 100\tau + \mathbf{b}\right] && (3.8) \\ &= \mathbb{E}\left[\frac{\mathbf{a}^2}{2}\left\{cos(100\tau) + cos(200t - 100\tau + 2\mathbf{b})\right\}\right] && (3.9) \end{aligned}$$

following from the trigonometric rule that $2 \cdot cos\alpha \cdot cos\beta = cos(\alpha + \beta) + cos(\alpha - \beta)$.

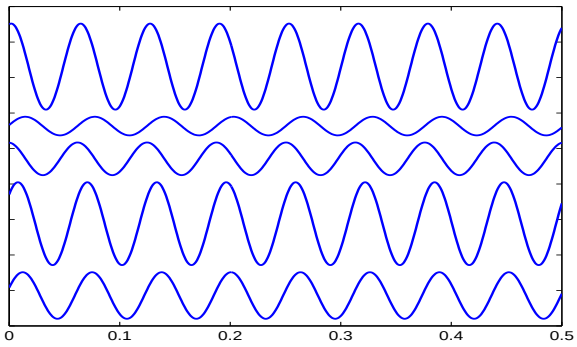

Figure 3.6: Five different realizations of the random process (3.7).

Since $\mathbb{E}[cos(200t - 100\tau + 2\mathbf{b})] = \int_{-\pi}^{\pi} cos(200t - 100\tau + 2b)\frac{1}{2\pi}db = 0$, it follows that

$$R_{\mathbf{x}}(t, t-\tau) = \frac{1}{2}cos(100\tau).$$

Note that the autocorrelation function is independent of the absolute time t; it is only dependent on the time difference τ. It also shows that the autocorrelation function is periodic in τ.
As every realization of the stochastic process is a cosine function with frequency 100 rad/sec, it is not surprising that the correlation between time samples that are one period apart is high. One period of the cosine refers to $\tau = 2\pi/100$, which implies that $cos(100\tau) = cos(2\pi) = 1$, which is its maximum value.
The value of the autocorrelation function can be interpreted by considering a scaled version of $R_{\mathbf{x}}(t, t-\tau)$, which is denoted as the correlation coefficient $\rho_{\mathbf{x}(t)\mathbf{x}(t-\tau)}$ between the (zero mean) random variables $\mathbf{x}(t)$ and $\mathbf{x}(t-\tau)$ (see Appendix A1):

$$\rho_{\mathbf{x}(t)\mathbf{x}(t-\tau)} = \frac{R_{\mathbf{x}}(t, t-\tau)}{\sqrt{var\{\mathbf{x}(t)\}var\{\mathbf{x}(t-\tau)\}}} = \frac{R_{\mathbf{x}}(t, t-\tau)}{\sqrt{R_{\mathbf{x}}(0)R_{\mathbf{x}}(0)}} = \cos(100\tau).$$

The correlation coefficient directly relates the statistical relation between the random variables $\mathbf{x}(t)$ and $\mathbf{x}(t-\tau)$. Because of its scaling it takes on real values between -1 and 1. If $|\rho_{\mathbf{x}(t)\mathbf{x}(t-\tau)}| = 1$, $\mathbf{x}(t)$ and $\mathbf{x}(t-\tau)$ are linearly related.

As mentioned, the previous example clearly shows that a random process does not necessarily generate signals that are "noisy". In the example, all signals that are generated are "nicely-behaving" sinusoids. In this case it is the non-deterministic character of the amplitude and phase that leads to the random nature of this process.

3.4 Particular examples of stochastic processes

3.4.1 White noise process

Consider a real-valued stochastic process $\mathbf{x}(t)$ defined by the following properties:

- $\mathbf{x}(t)$ has a pdf $f_{\mathbf{x}}(x)$ that is fixed for all t, with $\mu_{\mathbf{x}} = 0$;

- $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ are independent random variables for all $t_1 \neq t_2$.

This is also denoted as a sequence of zero mean identically distributed independent random variables.
It follows that

$$R_{\mathbf{x}}(t_1, t_2) = \mathbb{E}[\mathbf{x}(t_1)\mathbf{x}(t_2)] = \mathbb{E}[\mathbf{x}(t_1)] \cdot \mathbb{E}[\mathbf{x}(t_2)] = 0 \ \forall t_1 \neq t_2$$

For $t_1 = t_2$:

$$R_{\mathbf{x}}(t_1, t_1) = \mathbb{E}[\mathbf{x}(t_1)^2] = var(\mathbf{x}(t_1)).$$

The autocorrelation (and also the autocovariance) function is zero for all t_1, t_2, except in those values where $t_1 = t_2$.
This process has the (spectral) properties of a so-called white noise process, as explained later in this chapter.
For processes (sequences) where the time is discrete, there are several ways to generate a white noise process. The most straightforward example is to draw independent samples from a random variable with a fixed zero-mean pdf (e.g., normal distribution or uniform distribution). Since for all time instants the samples are drawn independently, there will be no correlation in time and consequently the autocorrelation function $R_{\mathbf{x}}(t_1, t_2)$ will be 0 for $t_1 \neq t_2$. Typical time realizations are given in Figure 3.7. Note that a white noise process generated by a uniform distribution has an amplitude that is bounded by the pdf-interval. The Gaussian white noise can have any amplitude, although high amplitudes (larger than 3σ) have very small probability.

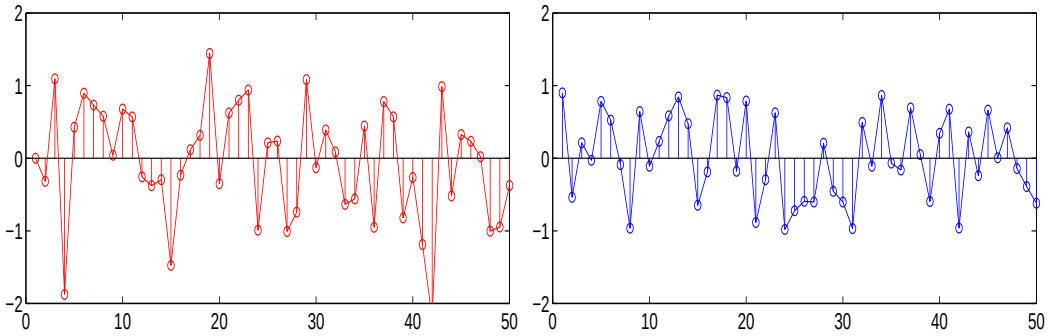

Figure 3.7: Realizations of two different white noise processes based on a Gaussian (normal) pdf $\mathcal{N}(0, 1)$ (left), and on a uniform pdf over the interval $[-1, 1]$ (right).

An alternative way to generate a white noise process in discrete time is given in section 3.4.5.

3.4.2 Random walk - Wiener process

In 1923, Norbert Wiener[1] derived a random process model for Brownian motion, i.e., the random movement of a small particle when immersed in a medium, due to a bombardment

[1] Norbert Wiener (1894-1964) was an American mathematician who is considered to be the founder of modern (statistical) signal analysis.

of the molecules of the medium.[2] The principle of Brownian motion was explained by Albert Einstein in 1905. The Wiener process can be derived as a limiting behaviour of a so-called random walk.

Consider a stochastic process defined by

$$\mathbf{x}(n) = \sum_{k=1}^{n} \mathbf{a}_k \qquad \mathbf{a}_i = \pm d$$

with d a real-valued number.
A possible realization of $\mathbf{x}(n)$ is indicated in Figure 3.8. At every (discrete) time step, a coin is tossed and dependent on its outcome (head or tail) a positive or a negative step ($\pm d$) is added to the previous signal value.

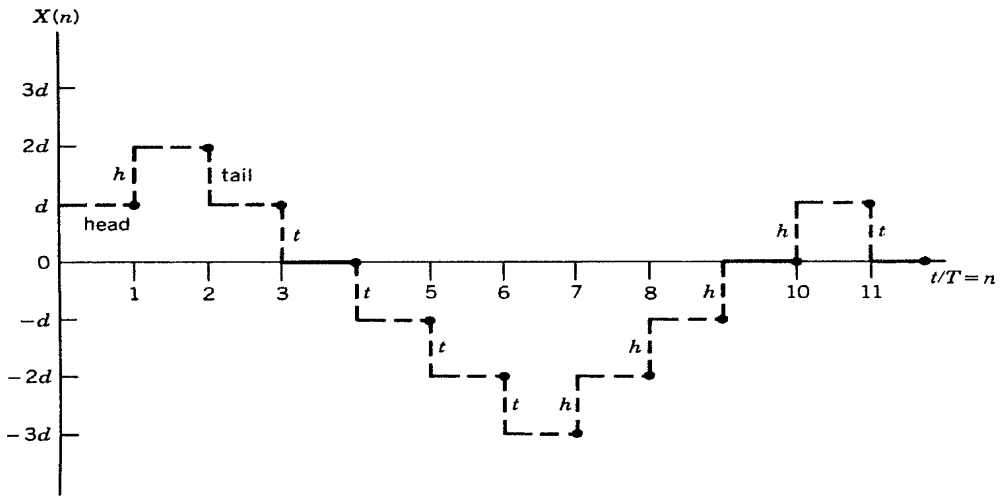

Figure 3.8: Realization of random walk

$$\mathbb{E}[\mathbf{x}(n)] = \sum_{k=1}^{n} \mathbb{E}[\mathbf{a}_k] = 0 \tag{3.10}$$

$$\mathbb{E}[\mathbf{x}(n)^2] = \mathbb{E}\left[(\mathbf{a}_1 + \mathbf{a}_2 + \cdots + \mathbf{a}_n)^2\right] = nd^2 \tag{3.11}$$

On the basis of this one can construct a continuous process:

$$\mathbf{y}(t) = \begin{cases} 0, & t = 0 \\ \mathbf{x}(n), & (n-1)T < t \leq nT, \quad n = 1, 2, 3, \cdots \end{cases}$$

Then at $t = nT$:

$$\mathbb{E}[\mathbf{y}(t)] = 0 \quad t = nT \tag{3.12}$$

$$\mathbb{E}[\mathbf{y}^2(t)] = nd^2 = \frac{t}{T}d^2 \tag{3.13}$$

[2]For an illustrative example showing Brownian motion see
http://www.phys.virginia.edu/classes/109N/more_stuff/Applets/brownian/applet.html.

If we take the limit for $T \to 0$, $d \to 0$, with $\frac{d^2}{T} = \alpha$ constant, a so-called Wiener process $\mathbf{w}(t)$ results, having the following properties:

- $\mathbf{w}(t)$ is continuous in time and amplitude;
- $\mathbf{w}(t)$ has independent increments, i.e., $\forall t_1 < t_2 < t_k :\ \mathbf{w}(t_k) - \mathbf{w}(t_2),\ \mathbf{w}(t_2) - \mathbf{w}(t_1)$ are mutually independent;
- $\mathbb{E}[\mathbf{w}(t)] = 0,\ \ \mathbb{E}[\mathbf{w}^2(t)] = \alpha t$
- Because $\mathbf{w}(t)$ is constructed as a summation of n independent, identically distributed random variables, the Central Limit Theorem (CLT) (see section A.5) shows that $\mathbf{w}(t)$ has a Gaussian probability density function for $n \to \infty$.

$\mathbf{w}(t)$ has mean value 0 and a variance that grows linearly with t. An example of a realization of a Wiener process is given in figure 3.9.

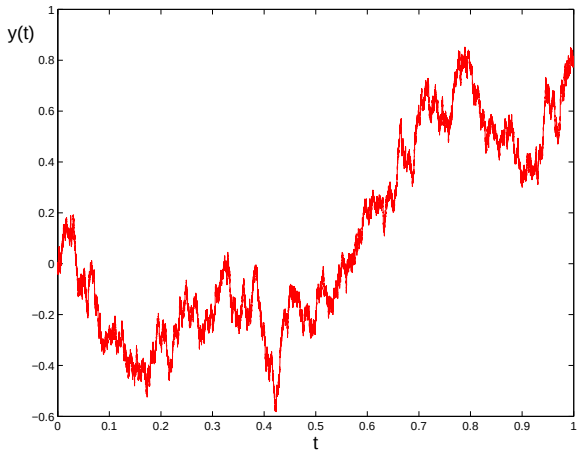

Figure 3.9: Realization of a Wiener process

3.4.3 Poisson random process

The Poisson process is a continuous-time, discrete-amplitude random process that is used to model random phenomena that are the result of a counting process, as, e.g., the emission of photons from a light-emitting diode, or the number of α-particles from a radio-active source arriving at a particular point in an image over a fixed observation time, or electrons arriving at a positively charged plate.
It is based on the positioning of n points randomly in the interval $(0, T)$.
When randomly positioning 1 point in the time interval $(0, T)$ the probability that it is positioned in (t_1, t_2) equals $\frac{t_2 - t_1}{T}$. In other words: the probability that it is positioned in an interval of length t equals t/T.
Now as a generalization of this, we randomly position n points, and the counting function $\mathbf{y}(t)$ denotes the number of events in an interval of length t. Then

$$P[\mathbf{y}(t) = k] = \left(\begin{array}{c} n \\ k \end{array} \right) p^k q^{n-k}$$

with $p = t/T$ and $q = 1 - p$.
For $n >> 1$, $p << 1$, but $np \approx 1$ the Poisson theorem states that the above mentioned probability can be approximated by (see e.g. Papoulis, 1991):

$$\left(\begin{array}{c} n \\ k \end{array} \right) p^k q^{n-k} \simeq \frac{(np)^k}{k!} e^{-np}.$$

For $n \to \infty$, $p \to 0$ and $np = nt/T \to \lambda t$ with $\lambda = n/T$ the average number of events per time unit, it follows that

$$P[\mathbf{y}(t) = k] = \frac{(\lambda t)^k}{k!} e^{-\lambda t}.$$

A random process with this property is called a Poisson process, under the additional restriction that the number of events in nonoverlapping time intervals are independent. A typical realization $\mathbf{y}(t)$ is sketched in Figure 3.10.

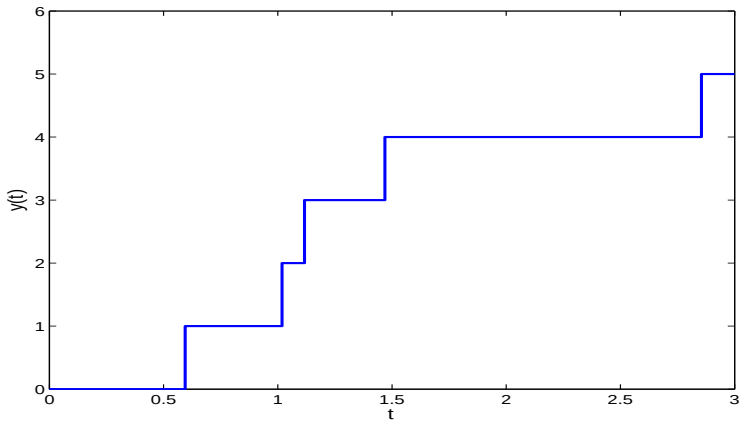

Figure 3.10: One realization $y(t)$ of the Poisson random process.

Mean value

$$\mathbb{E}[\mathbf{y}(t)] = \sum_k k \cdot P[\mathbf{y}(t) = k] = \sum_k k \frac{(\lambda t)^k}{k!} e^{-\lambda t} = \lambda t.$$

This can be understood by realizing that $\sum_{k=1}^{\infty} k \frac{a^k}{k!} = a e^a$ which is a result of differentiating the standard series expansion $e^a = 1 + \sum_{k=1}^{\infty} \frac{a^k}{k!}$ with respect to a.

Variance

$$\begin{aligned} var\{\mathbf{y}(t)\} &= \mathbb{E}[\mathbf{y}(t) - \mathbb{E}(\mathbf{y}(t))]^2 = \mathbb{E}[\mathbf{y}^2(t)] + (\mathbb{E}[\mathbf{y}(t)])^2 - 2(\mathbb{E}[\mathbf{y}(t)])^2 \\ &= \mathbb{E}[\mathbf{y}^2(t)] - (\mathbb{E}[\mathbf{y}(t)])^2 \end{aligned} \tag{3.14}$$

The first term in this expression is given by

$$\mathbb{E}[\mathbf{y}(t)^2] = \sum_k k^2 \frac{(\lambda t)^k}{k!} e^{-\lambda t}.$$

Differentiating $ae^a = \sum_k k\frac{a^k}{k!}$ with respect to a delivers:

$$\begin{aligned} e^a(1+a) &= \sum_k k^2 \frac{a^{k-1}}{k!} \\ ae^a(1+a) &= \sum_k k^2 \frac{a^k}{k!} \\ a(1+a) &= \sum_k k^2 \frac{a^k}{k!} e^{-a} \\ &= \mathbb{E}[\mathbf{y}(t)^2] \quad \text{for } a = \lambda t. \end{aligned} \tag{3.15}$$

As a result $\mathbb{E}[\mathbf{y}^2(t)] = \lambda t(1+\lambda t)$ and (3.14) shows that

$$var\{\mathbf{y}(t)\} = \lambda t(1+\lambda t) - (\lambda t)^2 = \lambda t.$$

For a more thorough description of the Poisson process the reader is referred to Dekking et al. (2005).

3.4.4 Continuous-time random binary waveform

$$\mathbf{x}(t) = \sum_{k=-\infty}^{\infty} \mathbf{a}_k p(t - kT - \mathbf{b})$$

with $p(\cdot)$ a unit pulse of duration T, see figure 3.11.

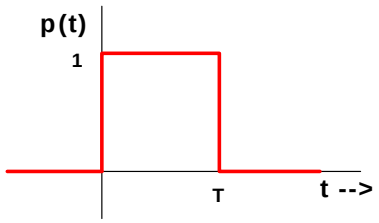

Figure 3.11: Unit pulse function $p(t)$

$\mathbf{a}_k$ is a binary random variable:

$$P[\mathbf{a}_k = +1] = P[\mathbf{a}_k = -1] = \frac{1}{2},$$

and $\mathbf{b}$ is a continuous random variable with a uniform distribution in $[0, T]$. All pulse amplitudes $\mathbf{a}_k$ are independent.

For any value of t, $\mathbf{x}(t)$ has one of two possible values ± 1, with equal probability, because it is determined by a single variable $\mathbf{a}_k$. Hence the mean and variance of $\mathbf{x}(t)$ are

$$\mathbb{E}[\mathbf{x}(t)] = 0; \qquad \mathbb{E}[\mathbf{x}^2(t)] = 1.$$

The autocorrelation $R_{\mathbf{x}}(t_1, t_2)$ is given by

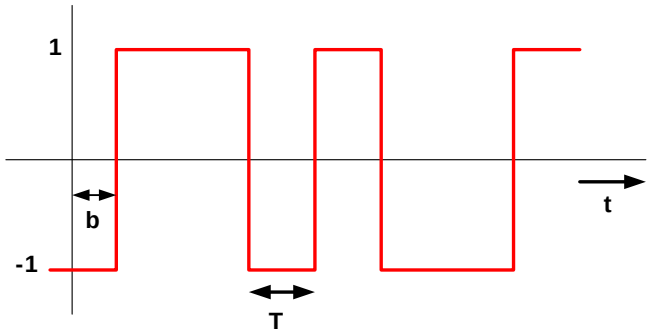

Figure 3.12: Random binary waveform

$$\begin{aligned} R_{\mathbf{x}}(t_1,t_2) &= 1-\frac{|t_2-t_1|}{T}, \qquad |t_2-t_1|<T \\ &= 0 \qquad\qquad\qquad\quad \text{elsewhere.} \end{aligned} \tag{3.16}$$

The derivation of this is as follows:
For $0<t_1<t_2<T$:

$$\begin{aligned} \mathbf{x}(t_1)\mathbf{x}(t_2) &= 1 \quad \text{if } 0<\mathbf{b}<t_1 \text{ or } t_2<\mathbf{b}<T \\ &= \pm 1 \quad \text{if } t_1\le \mathbf{b}\le t_2. \end{aligned}$$

If $t_1<\mathbf{b}<t_2$ then the probability that $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ have different sign is $1/2$.

$$\begin{aligned} \mathbb{E}[\mathbf{x}(t_1)\mathbf{x}(t_2)] &= \mathbb{E}\{\mathbf{x}(t_1)\mathbf{x}(t_2)|0<\mathbf{b}<t_1, \text{or } t_2<\mathbf{b}<T\}\cdot P[0<\mathbf{b}<t_1 \text{ or } t_2<\mathbf{b}<T]+ \\ &\quad +\mathbb{E}\left[\mathbf{x}(t_1)\mathbf{x}(t_2)|t_1<\mathbf{b}<t_2\right]\cdot P[t_1<\mathbf{b}<t_2] \\ &= 1\cdot\left[1-\frac{(t_2-t_1)}{T}\right]+1\cdot\frac{1}{2}\cdot\frac{(t_2-t_1)}{T}-1\cdot\frac{1}{2}\cdot\frac{(t_2-t_1)}{T} \\ &= 1-\frac{(t_2-t_1)}{T}. \end{aligned}$$

Interchanging the role of t_1 and t_2 shows the validity of the expression (3.16).
For $|t_2-t_1|>T$ there does not exist a statistical relation anymore between $\mathbf{x}(t_1)$ and $\mathbf{x}(t_2)$ because of the randomly occurring sign-changes. As a result $R_{\mathbf{x}}(t_1,t_2)=0$ for $|t_2-t_1|>T$.

3.4.5 Discrete-time random binary process

In the discrete-time situation, i.e., when the time axis is a discrete sequence of time instants, the following random binary waveform can be considered:

$$\mathbf{x}(t)=\sum_{k=-\infty}^{\infty}\mathbf{a}_k\delta(t-k), \quad t=\cdots,-2,-1,0,1,2,\cdots$$

with $\delta(t)$ a discrete pulse function, as indicated in Figure 3.13.

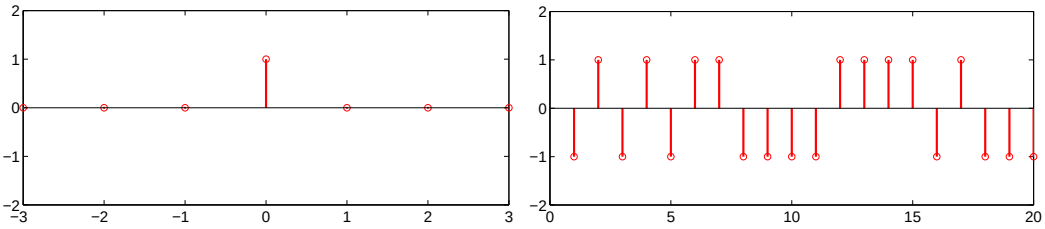

Figure 3.13: Discrete pulse function $\delta(t)$ (left) and a realization of random binary process (right)

$\mathbf{a}_k$ is a binary random variable:

$$P[\mathbf{a}_k = +1] = P[\mathbf{a}_k = -1] = \frac{1}{2},$$

while all pulse amplitudes $\mathbf{a}_k$ are independent, which implies that $\mathbb{E}[\mathbf{a}_k \mathbf{a}_\ell] = 0$ for $k \neq \ell$.

For any value of t, $\mathbf{x}(t)$ has one of two possible values ± 1, with equal probability, because it is determined by a single variable $\mathbf{a}_k$. As a result the mean and variance of $\mathbf{x}(t)$ are determined by

$$\begin{aligned} \mathbb{E}[\mathbf{x}(t)] &= \sum_{k=-\infty}^{\infty} \delta(t-k) \cdot \mathbb{E}(a_k) = 0; \\ \mathbb{E}[\mathbf{x}^2(t)] &= 1. \end{aligned}$$

The autocorrelation $R_{\mathbf{x}}(t_1, t_2)$ is given by

$$\begin{aligned} R_{\mathbf{x}}(t_1, t_2) &= \mathbb{E}\left[\sum_{k=-\infty}^{\infty} \mathbf{a}_k \delta(t_1 - k) \sum_{\ell=-\infty}^{\infty} \mathbf{a}_\ell \delta(t_2 - \ell) \right] && (3.17) \\ &= \mathbb{E}[\mathbf{a}_{t_1} \mathbf{a}_{t_2}] && (3.18) \\ &= \begin{cases} 0 & t_1 \neq t_2 \\ 1 & t_1 = t_2 \end{cases} && (3.19) \end{aligned}$$

This shows that the random binary process is indeed a white noise process. The autocorrelation function of this process can be further influenced by adapting the probabilistic properties of the amplitudes $\mathbf{a}_k$, e.g. by

$$P[\mathbf{a}_k = a_{k-1}] = \alpha \qquad P[\mathbf{a}_k = a_{k-1}] = 1 - \alpha.$$

A white noise process will only result if $\alpha = 0.5$.

3.4.6 Gaussian process

A random process is called Gaussian if *all* n-th order distribution functions $F_{\mathbf{x}_1, \mathbf{x}_2, \cdots \mathbf{x}_n}(x_1, x_2, \cdots x_n)$ are n-variate Gaussian distributions, i.e.,

$$F_{\mathbf{x}_1, \cdots \mathbf{x}_n}(x_1, \cdots x_n) = \int_{-\infty}^{x_1} \cdots \int_{-\infty}^{x_n} \frac{1}{(2\pi)^{n/2} \sqrt{det(\Sigma_{\mathbf{x}})}} \, exp\left(-\frac{1}{2} (x - \mu_{\mathbf{x}})^T \Sigma_{\mathbf{x}}^{-1} (x - \mu_{\mathbf{x}}) \right) dx_1 \cdots dx_n.$$

See also Appendix A.

3.4.7 Autoregressive process

A special class of random processes is formed by the so-called autoregressive (AR) processes. A discrete-time autoregressive process is one represented by a difference equation of the form:

$$\mathbf{x}(t) = \alpha_1 \mathbf{x}(t-1) + \alpha_2 \mathbf{x}(t-2) + \cdots + \alpha_p \mathbf{x}(t-p) + \mathbf{e}(t) \tag{3.20}$$

where $\mathbf{x}(t)$ is a real random sequence, $\alpha_i, i = 1, \ldots, p$ are real-valued constant coefficients, and $\mathbf{e}$ is a discrete-time white noise process, that is,

$$\mathbb{E}[\mathbf{e}(t)] = 0, \quad \forall t \tag{3.21}$$

$$\mathbb{E}[\mathbf{e}(t_1)\mathbf{e}(t_2)] = \begin{cases} \sigma^2_{\mathbf{e}}, & t_1 = t_2 \\ 0, & t_1 \neq t_2 \end{cases} \tag{3.22}$$

If the difference equation is of the order p (i.e., $\alpha_p \neq 0$), then the process is called a pth order autoregressive process (denoted by AR(p)).

First order autoregressive process

Let us now consider the AR(1) process in more detail. An AR(1) process is described by the following difference equation:

$$\mathbf{x}(t) = \alpha \mathbf{x}(t-1) + \mathbf{e}(t) \tag{3.23}$$

where $\mathbf{e}$ is a white noise process, and α is a real-valued constant. To study the properties of this model, we start by solving the difference equation (3.23) so as to obtain an expression for $\mathbf{x}(t)$ in terms of $\mathbf{e}(t), \mathbf{e}(t-1), \mathbf{e}(t-2), \ldots$. The precise solution of this equation depends on the "initial condition", so let us suppose that initially $\mathbf{x}(0) = 0$, so that $\mathbf{x}(1) = \mathbf{e}(1)$. Then solving (3.23) by repeated substitution we obtain

$$\mathbf{x}(t) = \sum_{i=0}^{t-1} \alpha^i \mathbf{e}(t-i). \tag{3.24}$$

Since $\mathbf{e}$ is a white noise process, it follows that $\mathbb{E}[\mathbf{x}(t)] = 0$ and

$$\mathbb{E}[\mathbf{x}^2(t)] = \mathbb{E}\left[\left(\mathbf{e}(t) + \alpha\mathbf{e}(t-1) + \alpha^2\mathbf{e}(t-2) + \cdots + \alpha^{t-1}\mathbf{e}(1)\right)^2\right] \tag{3.25}$$

$$= \sigma^2_{\mathbf{e}}\left(1 + \alpha^2 + \alpha^4 + \ldots + \alpha^{2(t-1)}\right) \tag{3.26}$$

$$= \begin{cases} \sigma^2_{\mathbf{e}}\left(\frac{1-\alpha^{2t}}{1-\alpha^2}\right), & |\alpha| \neq 1 \\ \sigma^2_{\mathbf{e}} t, & |\alpha| = 1 \end{cases} \tag{3.27}$$

Similarly, we find for the autocorrelation function of $\mathbf{x}$ (if $\tau \geq 0$):

$$\begin{aligned} R_{\mathbf{x}}(t+\tau, t) &:= \mathbb{E}[\mathbf{x}(t+\tau)\mathbf{x}(t)] \\ &= \mathbb{E}\left[\left(\mathbf{e}(t+\tau) + \alpha\mathbf{e}(t-1+\tau) + \cdots + \alpha^\tau \mathbf{e}(t) + \cdots + \alpha^{t-1+\tau}\mathbf{e}(1)\right)\right. \\ &\quad \left.\left(\mathbf{e}(t) + \alpha\mathbf{e}(t-1) + \cdots + \alpha^{t-1}\mathbf{e}(1)\right)\right] \\ &= \sigma^2_{\mathbf{e}}\left(\alpha^\tau + \alpha^{\tau+2} + \ldots + \alpha^{\tau+2(t-1)}\right) \\ &= \begin{cases} \sigma^2_{\mathbf{e}} \cdot \alpha^\tau \left(\frac{1-\alpha^{2t}}{1-\alpha^2}\right), & |\alpha| \neq 1 \\ \sigma^2_{\mathbf{e}} t, & |\alpha| = 1 \end{cases} \end{aligned} \tag{3.28}$$

Similarly we find for $\tau \leq 0$ (if $t > |\tau|$):

$$R_{\mathbf{x}}(t+\tau,t) = \begin{cases} \sigma_{\mathbf{e}}^2 \cdot \alpha^{-\tau} \left(\frac{1-\alpha^{2t}}{1-\alpha^2}\right), & |\alpha| \neq 1 \\ \sigma_{\mathbf{e}}^2 t, & |\alpha| = 1 \end{cases} \tag{3.29}$$

The proof is left as an exercise to the reader. Generally, both $\mathbb{E}[\mathbf{x}^2(t)]$ and $R_{\mathbf{x}}(t+\tau,t)$ are functions of the absolute time t. However, if $|\alpha| < 1$ we may argue that for t sufficiently large,

$$\begin{aligned} \mathbb{E}[\mathbf{x}^2(t)] &\approx \frac{\sigma_{\mathbf{e}}^2}{1-\alpha^2}, \\ R_{\mathbf{x}}(t,t+\tau) &\approx \alpha^{|\tau|} \frac{\sigma_{\mathbf{e}}^2}{1-\alpha^2}. \end{aligned} \tag{3.30}$$

The right hand side of (3.30) is now a function of the time difference τ only. AR processes will be discussed more extensively in section 4.3

3.5 Cross-correlation and cross-covariance functions

When considering two different stochastic processes, the following notions are instrumental in describing and quantifying statistical relationships between the two processes.

Cross-correlation function:

$$R_{\mathbf{yx}}(t_1,t_2) := \mathbb{E}[(\mathbf{y}(t_1)\mathbf{x}^*(t_2)] \tag{3.31}$$

Cross-covariance function:

$$C_{\mathbf{yx}}(t_1,t_2) := \mathbb{E}[(\mathbf{y}(t_1) - \mu_{\mathbf{y}}(t_1))(\mathbf{x}^*(t_2) - \mu_{\mathbf{x}}^*(t_2))] \tag{3.32}$$

These functions reflect the statistical relationship between two processes.

Two stochastic processes $\mathbf{x}$ and $\mathbf{y}$ are called:

- *uncorrelated* if $C_{\mathbf{yx}}(t_1,t_2) = 0, \ \forall t_1,t_2 \in \Gamma$
- *orthogonal* if $R_{\mathbf{yx}}(t_1,t_2) = 0, \ \forall t_1,t_2 \in \Gamma$
- *independent* if *all* joint distribution functions are equal to the product of the corresponding single-variable distribution functions; this implies e.g. that

$$P\left[\mathbf{x}(t_1) \leq x_1, y(t_2) \leq y_2\right] = P\left[\mathbf{x}(t_1) \leq x_1\right] \cdot P\left[\mathbf{y}(t_2) \leq y_2\right].$$

Example 3.4 (Time-delayed white noise process) Suppose an acoustic noise signal is observed by a sensor at a particular location, i.e., $\mathbf{x}(t) = \mathbf{v}(t)$ with $\mathbf{v}(t)$ a stochastic process with zero mean and variance $\sigma_{\mathbf{v}}^2$ for all t. The same acoustic noise is also observed at a different location, with a sensor $\mathbf{y}(t)$ and will appear as a time-delayed version of the signal that is measured at sensor $\mathbf{x}$, i.e., $\mathbf{y}(t) = \mathbf{v}(t-\tau)$.

The relation between $\mathbf{x}$ and $\mathbf{y}$ is reflected in the cross-correlation function:

$$R_{\mathbf{yx}}(t_1, t_2) = \mathbb{E}[\mathbf{v}(t_1 - \tau)\mathbf{v}(t_2)].$$

In the particular situation that $\mathbf{v}$ is a white noise process, it follows that

$$\begin{aligned} R_{\mathbf{yx}}(t_1, t_2) &= \sigma_{\mathbf{v}}^2 \quad \text{for } t_2 = t_1 - \tau \\ &= 0 \qquad\quad \text{elsewhere} \end{aligned}$$

or

$$\begin{aligned} R_{\mathbf{yx}}(t, t-k) &= \sigma_{\mathbf{v}}^2 \quad \text{for } k = \tau \\ &= 0 \qquad\quad \text{elsewhere} \end{aligned} \tag{3.33}$$

The cross-correlation function is independent of t, and only dependent on the time difference between t_2 and t_1, as reflected in the latter expression.
Note that in this case the relation between the stochastic processes $\mathbf{x}$ and $\mathbf{y}$ is given by $\mathbf{y}(t) = q^{-\tau}\mathbf{x}(t)$. This time delay of τ time steps, appears similarly in the cross-correlation function as a (discrete) delta function shifted over τ time steps: $R_{\mathbf{yx}}(t, t-k) = \sigma_{\mathbf{v}}^2\delta(k - \tau)$. This is depicted in figure 3.14.

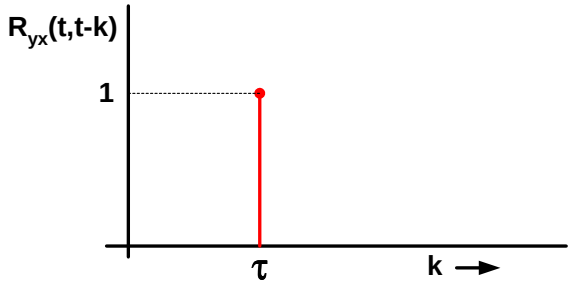

Figure 3.14: $R_{\mathbf{yx}}(t, t-k) = \sigma_{\mathbf{v}}^2\delta(k - \tau)$ (example 3.4).

Example 3.5 (Detection of a phase shift in two noise disturbed sinusoids) Suppose that measurements are available of two stochastic processes $\mathbf{u}$ and $\mathbf{y}$, being sinusoids of given amplitude and fixed frequency, but disturbed by additive noise:

$$\mathbf{u}(t) = sin(\omega t) + \mathbf{v}_1(t) \tag{3.34}$$
$$\mathbf{y}(t) = sin(\omega t + \varphi) + \mathbf{v}_2(t) \tag{3.35}$$

with ω fixed but possibly unknown, φ fixed but unknown, and $\mathbf{v}_1, \mathbf{v}_2$ uncorrelated zero-mean stochastic processes. Determine φ from the cross-correlation function $R_{\mathbf{yu}}(t, t-\tau)$.

Solution

$$R_{\mathbf{yu}}(t, t-\tau) = \mathbb{E}[sin(\omega t + \varphi) + \mathbf{v}_2(t)][sin(\omega(t-\tau)) + \mathbf{v}_1(t-\tau)] \tag{3.36}$$
$$= 0.5[cos(\varphi + \omega\tau) - cos(2\omega t + \varphi - \omega\tau)] + \mathbb{E}\mathbf{v}_2(t)\mathbf{v}_1(t-\tau) \tag{3.37}$$
$$= 0.5[cos(\varphi + \omega\tau) - cos(2\omega t + \varphi - \omega\tau)]. \tag{3.38}$$

The cross-correlation function is sketched in figure 3.15(left) as a function of t and τ, for $\omega = 5rad/sec$ and $\varphi = \pi/4$.
For a fixed value $\tau = \tau_f$ the cross-correlation function is a function of t only. In that case ω can be determined from the period length of the periodic signal $R_{\mathbf{yu}}(t, t-\tau_f)$, while φ can be determined with the following procedure:
Suppose the function $R_{\mathbf{yu}}(t, t-\tau_f)$ reaches a maximum value of r_{max}, then

$$0.5[cos(\varphi + \omega\tau_f) + 1] = r_{max}.$$

As a result:

$$\varphi = -\omega\tau_f \pm arccos(2r_{max} - 1). \tag{3.39}$$

The curve $R_{\mathbf{yu}}(t, t-\tau_f)$ is sketched in figure 3.15(right) for $\tau_f = 1$. Substituting the

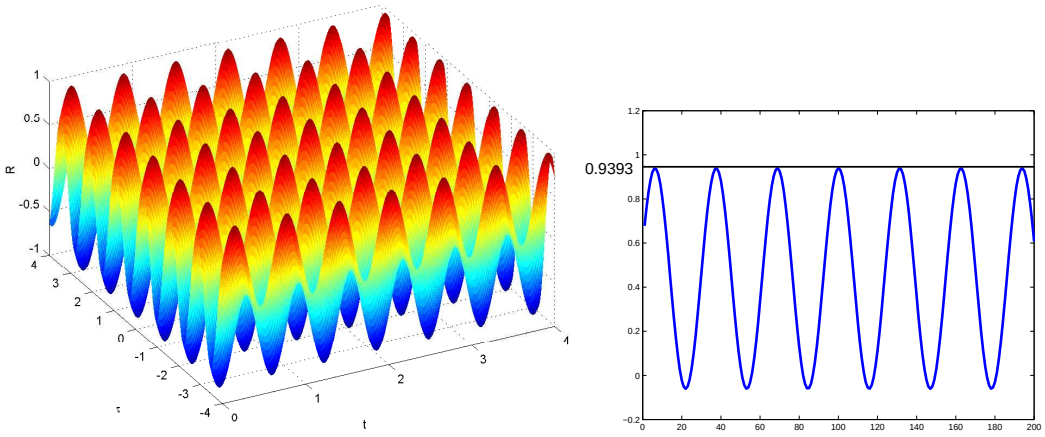

Figure 3.15: Left: $R_{\mathbf{yu}}(t, t-\tau)$ for $\omega = 5rad/sec$ and $\varphi = \pi/4$; right: $R_{\mathbf{yu}}(t, t-\tau_f)$ for $\tau_f = 1$.

numerical values, leads to:

$$\varphi = -5 \pm arccos(2 \cdot 0.9393 - 1) = -5 \pm 0.4978rad.$$

As a result: $\varphi = -1.75\pi$ or -1.43π. The first solution corresponds to $\pi/4$ rad. The duality in the expression (3.39) can be removed, by either incorporating a second curve $R_{\mathbf{yu}}(t, t-\tau_f)$ for another value of τ_f, or by determining the phase shift $\varphi - \omega\tau_f$ in the cross-correlation function $R_{\mathbf{yu}}(t, t-\tau_f)$.

From the definitions the following two expressions are immediate:

- If two stochastic processes are independent they are also uncorrelated.
- If two processes $\mathbf{x}$, $\mathbf{y}$ are uncorrelated then $R_{\mathbf{yx}}(t_1, t_2) = \mu_{\mathbf{y}}(t_1)\mu_{\mathbf{x}}(t_2)$.

3.6 Stationarity

Several properties of a stochastic process can be invariant over time. While separate signals taken from the ensemble clearly exhibit time-dependent behavior, the statistical properties

of the process can be time-invariant. To this end there are several notions of stationarity of a stochastic process, depending on whether all or only part of the distribution functions and moments thereof are time-invariant. Basically a stochastic process is called *stationary* if its statistical properties are independent of time. We distinguish two different concepts:

- *Strict sense stationary (SSS)*: all joint distribution functions are invariant under a shift in time. This means that for all values of $t_1, \cdots t_n$, and for all values of τ and n:

$$P\left[\mathbf{x}(t_1) \leq x_1, \mathbf{x}(t_2) \leq x_2, \cdots \mathbf{x}(t_n) \leq x_n\right] = \tag{3.40}$$
$$P\left[\mathbf{x}(t_1+\tau) \leq x_1, \mathbf{x}(t_2+\tau) \leq x_2, \cdots \mathbf{x}(t_n+\tau) \leq x_n\right] \tag{3.41}$$

- *Wide sense stationary (WSS) or weakly stationary*: the first two moments are independent of time t, i.e., the mean value and the autocorrelation function are independent of time:

$$\mathbb{E}[\mathbf{x}(t)] = \mu_{\mathbf{x}}, \quad \text{constant} \tag{3.42}$$
$$\mathbb{E}[\mathbf{x}(t)\mathbf{x}^*(t-\tau)] = R_{\mathbf{x}}(\tau) \tag{3.43}$$

  for all values of t and τ.

It is straightforward to show that strict-sense stationarity (SSS) implies wide-sense stationarity (WSS). However the converse is in general not true. There is an important exception: for Gaussian stochastic processes, wide-sense stationarity implies strict-sense stationarity. This follows from the fact all statistical properties of a Gaussian stochastic process are determined by its mean and autocorrelation function, which is a direct result of the description of the multivariate Gaussian probability density function (A.10).

Two stochastic processes are "jointly" stationary (in SSS or in WSS sense) if each of the two is stationary, and if additionally all joint distribution functions (for SSS-sense) or their first two moments (for WSS-sense) are time-invariant.

An example of a non-stationary stochastic process is the noise disturbed sinusoidal

$$\mathbf{u}(t) = sin(\omega t) + \mathbf{v}_1(t)$$

of Example 3.5.
While $\mathbb{E}\mathbf{v}_1(t) = 0$ for all t, $\mathbb{E}\mathbf{u}(t) = sin(\omega t)$ is time-varying, and therefore the conditions for wide-sense stationarity are not met.
A typical example of such a situation is found in speech signals, of which an example is given in figure 3.16. The several vowels and consonants appear in particular time frames and lead to time-dependent correlation properties.
Except for the Wiener and the Poisson process, the white noise stochastic processes described in sections 3.4 are all wide sense stationary, as can easily be verified from the fact that mean value and autocorrelation function are independent of absolute time.

3.7 First and second order moments of WSS-processes

3.7.1 Introduction

The role of first and second order moments (mean value and correlation function) of WSS processes is twofold:

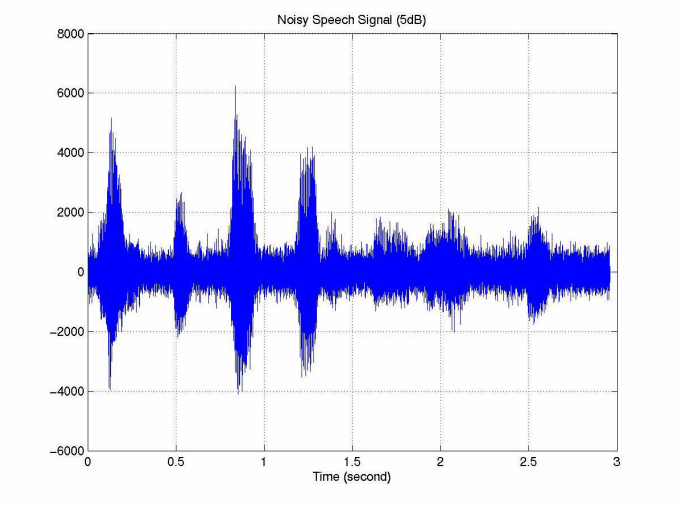

Figure 3.16: Speech signal

- To specify properties, i.e., to model stochastic processes; and in particular
- to find expressions for the distribution of power of the processes over frequency.

For deterministic types of signals the power or energy distribution of a signal over frequency can directly be obtained from Fourier analysis. For stochastic processes a different approach has to be followed, since the transformation of a single time sequence or signal does not only reflect the structural properties of the underlying stochastic process, but also the specific realization that is taken from the process.

3.7.2 Autocorrelation and autocovariance function

First we consider the autocorrelation function of a real-valued wide-sense stationary stochastic process.

$$\boxed{R_{\mathbf{x}}(\tau) = \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-\tau)]} \tag{3.44}$$

The variable τ is called the *lag*. It represents the shift between the two processes that are considered. It may take on both positive and negative values. The following basic properties of an autocorrelation function can be observed:

1. The *mean power* of the process:

$$\boxed{R_{\mathbf{x}}(0) = \mathbb{E}[\mathbf{x}^2(t)] \geq 0.}$$

2. $R_{\mathbf{x}}(\tau)$ is an even function

$$R_{\mathbf{x}}(\tau) = R_{\mathbf{x}}(-\tau).$$

This is a direct result of the definition.

3. $R_{\mathbf{x}}(\tau)$ is bounded by $R_{\mathbf{x}}(0)$

$$|R_{\mathbf{x}}(\tau)| \leq R_{\mathbf{x}}(0).$$

This can be shown as follows:

$$\begin{cases} \mathbb{E}[(\mathbf{x}(t+\tau) - \mathbf{x}(t))^2] \geq 0 \\ \mathbb{E}[(\mathbf{x}(t+\tau) + \mathbf{x}(t))^2] \geq 0 \end{cases} \Rightarrow$$

$$\begin{cases} 2R_{\mathbf{x}}(0) - 2R_{\mathbf{x}}(\tau) \geq 0 \\ 2R_{\mathbf{x}}(0) + 2R_{\mathbf{x}}(\tau) \geq 0 \end{cases} \Rightarrow$$

$$-R_{\mathbf{x}}(0) \leq R_{\mathbf{x}}(\tau) \leq R_{\mathbf{x}}(0).$$

4. If $\mathbf{x}(t)$ is periodic, then $R_{\mathbf{x}}(\tau)$ will have the same periodic component, i.e., if $\mathbf{x}(t+T_0) = \mathbf{x}(t)$ for all t, then

$$R_{\mathbf{x}}(\tau) = \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-\tau)] = \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-\tau-T_0)] = R_{\mathbf{x}}(\tau+T_0).$$

5. If $R_{\mathbf{x}}(T_0) = R_{\mathbf{x}}(0)$ for some $T_0 \neq 0$ then $R_{\mathbf{x}}(\tau)$ is periodic with period T_0, i.e.,

$$R_{\mathbf{x}}(\tau+T_0) = R_{\mathbf{x}}(\tau).$$

The proof of this follows from the inequality

$$\{\mathbb{E}[(\mathbf{x}(t+\tau+T_0) - \mathbf{x}(t+\tau))x(t)]\}^2 \leq \mathbb{E}[(\mathbf{x}(t+\tau+T_0) - \mathbf{x}(t+\tau))^2] \cdot \mathbb{E}[\mathbf{x}^2(t)]$$

which originates from the Schwartz inequality (A.12).

Hence

$$[R_{\mathbf{x}}(\tau+T_0) - R_{\mathbf{x}}(\tau)]^2 \leq \underbrace{[2R_{\mathbf{x}}(0) - 2R_{\mathbf{x}}(T_0)]}_{=0} \cdot R_{\mathbf{x}}(0)$$

showing that $R_{\mathbf{x}}(\tau+T_0) = R_{\mathbf{x}}(\tau)$ for all τ.

These five properties of autocorrelation functions show that not every function can be an autocorrelation function.
For discrete-time stochastic processes the allowable autocorrelation functions are exactly characterized by the following two restrictions:

- $R_{\mathbf{x}}(\tau) = R_{\mathbf{x}}(-\tau)$, and
- $R_{\mathbf{x}}(\tau)$ constitutes a correlation matrix

$$\mathbf{R}_{\mathbf{x}} := \begin{bmatrix} R_{\mathbf{x}}(0) & R_{\mathbf{x}}(1) & \cdots & R_{\mathbf{x}}(p-1) \\ R_{\mathbf{x}}(1) & R_{\mathbf{x}}(0) & \cdots & R_{\mathbf{x}}(p-2) \\ \vdots & \vdots & \ddots & \vdots \\ R_{\mathbf{x}}(p-1) & R_{\mathbf{x}}(p-2) & \cdots & R_{\mathbf{x}}(0) \end{bmatrix}$$

that is positive semidefinite for any finite value of p.

The latter condition is based on the property that if we consider any linear map

$$\mathbf{y} := a_1\mathbf{x}(1) + a_2\mathbf{x}(2) + \cdots + a_p\mathbf{x}(p)$$

with a_i any real-valued coefficients, and $\mathbf{y}$ a resulting random variable, then $\mathbb{E}[\mathbf{y}^2]$ which is determined by

$$\mathbb{E}[\mathbf{y}^2] = \mathbb{E}[a^T\mathbf{x}\mathbf{x}^T a] = a^T\mathbf{R}_\mathbf{x} a \geq 0$$

for all values of vector $a = [a_1 \cdots a_p]^T$. Consequently $\mathbf{R}_\mathbf{x} \geq 0$ for all values of p.

Example 3.6 Consider the discrete-time stochastic processes AR(1) $\mathbf{x}$ determined by the following time-recursive relation

$$\mathbf{x}(t) = \alpha\mathbf{x}(t-1) + \mathbf{e}(t) \tag{3.45}$$

where $\mathbf{e}$ is a white noise stochastic processes, and α is a real-valued constant, $|\alpha| < 1$. It follows from subsection 3.4.7 that this process is asymptotically WSS. That is, $\mathbf{x}(t)$ is WSS for t sufficiently large.
Consider the autocorrelation function of $\mathbf{x}$:

$$R_\mathbf{x}(\tau) = \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-\tau)] = \alpha\mathbb{E}[\mathbf{x}(t-1)\mathbf{x}(t-\tau)] + \mathbb{E}[\mathbf{e}(t)\mathbf{x}(t-\tau)]. \tag{3.46}$$

For positive values of τ the last term of this expression will be 0 due to the fact that $\mathbf{x}(t)$ will only be correlated with past and present values of $\mathbf{e}$ but not with future ones. In that case it follows that

$$R_\mathbf{x}(\tau) = \alpha R_\mathbf{x}(\tau - 1), \quad \tau \geq 1 \tag{3.47}$$

and consequently

$$R_\mathbf{x}(\tau) = \alpha^\tau R_\mathbf{x}(0), \quad \tau \geq 1. \tag{3.48}$$

For negative values of τ, $R_\mathbf{x}(\tau)$ is simply obtained by following the property that $R_\mathbf{x}(\tau)$ is an even function, i.e., $R_\mathbf{x}(\tau) = R_\mathbf{x}(-\tau)$.
In figure 3.17 the autocorrelation of this process is sketched for two different values of α, while the signals are scaled so as to obtain two processes with the same mean power $R_\mathbf{x}(0)$.

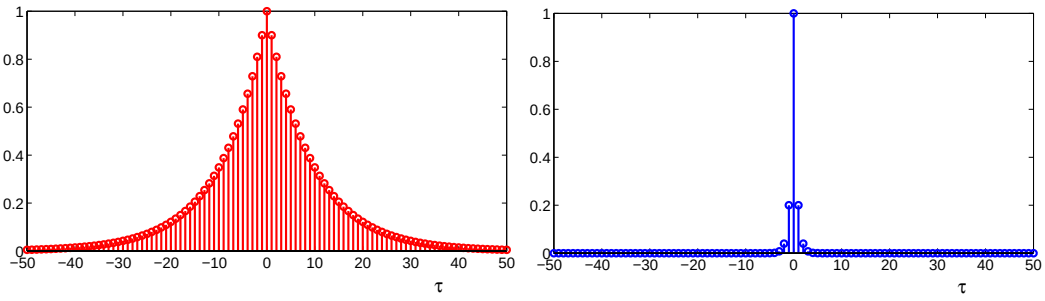

Figure 3.17: Autocorrelation function $R_\mathbf{x}(\tau)$ of stochastic process $\mathbf{x}$ (3.45) for $\alpha = 0.9$ (left) and $\alpha = 0.2$ (right) while the mean power of $\mathbf{e}$ is adjusted to obtain two processes with equal mean power $R_\mathbf{x}(0)$.

The interpretation of the two curves in figure 3.17 is that the correlation of figure (a) ranges over a longer horizon than the correlation of figure (b). One could state that the process

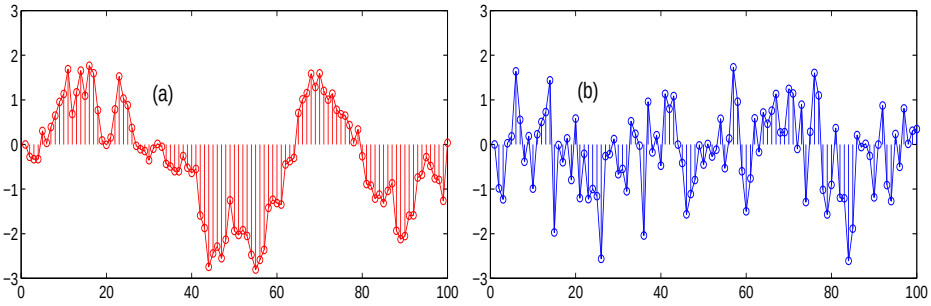

Figure 3.18: One realization (time sequence) of the stochastic process $\mathbf{x}$ (3.45) for $\alpha = 0.9$ (left) and $\alpha = 0.2$ (right)

for $\alpha = 0.9$ has a longer memory than for $\alpha = 0.2$. As a result, time sequences from process (a) will show a more smooth behaviour, while time sequences from process (b) will be more erratic. In figure 3.18 two corresponding time sequences are plotted.
As can be observed from these figures, $R_{\mathbf{x}}(\tau)$ will highly determine the frequency-content of the stochastic process.

Example 3.7 (White noise process) A discrete-time white noise stochastic process $\mathbf{x}(t)$ is determined by the property that it is a sequence of equally distributed zero-mean uncorrelated random variables, implying that

$$\mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-\tau)] = 0, \qquad \tau \neq 0.$$

There is no correlation between time samples at different time instants.

Autocovariance functions
The expressions for autocorrelation functions presented in the previous section can simply be extended to the autocovariance functions:

$$\boxed{C_{\mathbf{x}}(\tau) := \mathbb{E}[(\mathbf{x}(t) - \mu_{\mathbf{x}})(\mathbf{x}(t-\tau) - \mu_{\mathbf{x}})].} \tag{3.49}$$

Whereas $R_{\mathbf{x}}(0)$ denotes the mean power of the process,

$$\boxed{C_{\mathbf{x}}(0) = \mathbb{E}[(\mathbf{x}(t) - \mu_{\mathbf{x}})^2]} \tag{3.50}$$

is the (mean) *variance*. It can be observed that

$$R_{\mathbf{x}}(\tau) = C_{\mathbf{x}}(\tau) + \mu_{\mathbf{x}}^2. \tag{3.51}$$

Since the autocovariance function *is* an autocorrelation function for processes with zero mean, it follows that any property that holds for an autocorrelation function also holds for an autocovariance function. This applies to all properties considered before:

1. $C_{\mathbf{x}}(\tau) = C_{\mathbf{x}}(-\tau)$.

2. $|C_{\mathbf{x}}(\tau)| \leq C_{\mathbf{x}}(0)$.

3. If $\mathbf{x}(t+T_0) = \mathbf{x}(t)$ for all t, then $C_{\mathbf{x}}(\tau) = C_{\mathbf{x}}(\tau+T_0)$.

4. If $C_{\mathbf{x}}(T_0) = C_{\mathbf{x}}(0)$ for some $T_0 \neq 0$ then

$$C_{\mathbf{x}}(\tau+T_0) = C_{\mathbf{x}}(\tau).$$

Along the same line of reasoning as applied to autocorrelation functions, the characterizing properties of an autocovariance function are

- $C_{\mathbf{x}}(\tau) = C_{\mathbf{x}}(-\tau)$, and
- $C_{\mathbf{x}}(\tau)$ constitutes a covariance matrix

$$\Sigma_{\mathbf{x}} := \begin{bmatrix} C_{\mathbf{x}}(0) & C_{\mathbf{x}}(1) & \cdots & C_{\mathbf{x}}(p-1) \\ C_{\mathbf{x}}(1) & C_{\mathbf{x}}(0) & \cdots & C_{\mathbf{x}}(p-2) \\ \vdots & \vdots & \ddots & \vdots \\ C_{\mathbf{x}}(p-1) & C_{\mathbf{x}}(p-2) & \cdots & C_{\mathbf{x}}(0) \end{bmatrix}$$

  that is positive semidefinite, i.e., $\Sigma_{\mathbf{x}} \geq 0$ for any finite value of p.

3.7.3 Correlation and covariance matrices

When considering a discrete-time stochastic process $\mathbf{x}$ over a range of N time instants, the resulting N random variables can be collected in a random vector:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}(1) \\ \mathbf{x}(1) \\ \vdots \\ \mathbf{x}(N) \end{bmatrix}.$$

The correlation matrix of this vector random variable is denoted as

$$\mathbf{R}_{\mathbf{x}} := \mathbb{E} \begin{bmatrix} \mathbf{x}(1) \\ \mathbf{x}(2) \\ \vdots \\ \mathbf{x}(N) \end{bmatrix} [\mathbf{x}(1)\ \mathbf{x}(2)\ \cdots\ \mathbf{x}(N)] \tag{3.52}$$

which for wide-sense stationary processes satisfies

$$\mathbf{R}_{\mathbf{x}} = \begin{bmatrix} R_{\mathbf{x}}(0) & R_{\mathbf{x}}(1) & \cdots & R_{\mathbf{x}}(N-1) \\ R_{\mathbf{x}}(1) & R_{\mathbf{x}}(0) & \cdots & R_{\mathbf{x}}(N-2) \\ \vdots & \vdots & \ddots & \vdots \\ R_{\mathbf{x}}(N-1) & R_{\mathbf{x}}(N-2) & \cdots & R_{\mathbf{x}}(0) \end{bmatrix} \tag{3.53}$$

and by construction is a symmetric *Toeplitz matrix*. A Toeplitz matrix is a matrix with equal elements on all diagonals and subdiagonals.

Similarly, the covariance matrix of the vector random variable is denoted as

$$\Sigma_{\mathbf{x}} := \mathbb{E}\left[\begin{array}{c} \mathbf{x}(1)-\mu_{\mathbf{x}} \\ \mathbf{x}(2)-\mu_{\mathbf{x}} \\ \vdots \\ \mathbf{x}(N)-\mu_{\mathbf{x}} \end{array}\right][\mathbf{x}(1)-\mu_{\mathbf{x}}\ \mathbf{x}(2)-\mu_{\mathbf{x}}\ \cdots\ \mathbf{x}(N)-\mu_{\mathbf{x}}] \tag{3.54}$$

leading to

$$\Sigma_{\mathbf{x}} = \left[\begin{array}{cccc} C_{\mathbf{x}}(0) & C_{\mathbf{x}}(1) & \cdots & C_{\mathbf{x}}(N-1) \\ C_{\mathbf{x}}(1) & C_{\mathbf{x}}(0) & \cdots & C_{\mathbf{x}}(N-2) \\ \vdots & \vdots & \ddots & \vdots \\ C_{\mathbf{x}}(N-1) & C_{\mathbf{x}}(N-2) & \cdots & C_{\mathbf{x}}(0) \end{array}\right] \tag{3.55}$$

which also is a symmetric Toeplitz matrix.
For a WSS stochastic process the correlation and covariance matrices do not change if the N samples that are considered are taken over a shifted time period. Due to the stationarity property the correlation and covariance functions are independent of the (absolute) time and so a time-shift does not matter.
As formulated in section 3.7.2, the positive-semidefinite property of $\mathbf{R}_{\mathbf{x}}$ and $\Sigma_{\mathbf{x}}$ is essential in its characterization. It can simply be checked by calculating its eigenvalues. Since for a symmetric matrix all eigenvalues are real-valued, for $\mathbf{R}_{\mathbf{x}}$ and $\Sigma_{\mathbf{x}}$ to be positive semi-definite, their eigenvalues should be ≥ 0.

Example 3.8 The autocorrelation function of a stochastic process is given by

$$R_{\mathbf{x}}(\tau) = 4 \cdot (-0.5)^{|\tau|}.$$

The correlation matrix for $N = 3$ is

$$\left[\begin{array}{ccc} 4 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 4 \end{array}\right].$$

This symmetric Toeplitz matrix has eigenvalues 7.4, 3.0 and 1.6. Due to the positivity of all eigenvalues, the correlation matrix is positive definite.
For $N = 4$ the correlation matrix becomes

$$\left[\begin{array}{cccc} 4 & -2 & 1 & -0.5 \\ -2 & 4 & -2 & 1 \\ 1 & -2 & 4 & -2 \\ -0.5 & 1 & -2 & 4 \end{array}\right].$$

The eigenvalues of this matrix are: 8.3, 4.0, 2.2 and 1.5, and again the correlation matrix is positive definite.
It is shown in example 3.6 that the considered function can indeed be the autocorrelation of a WSS stochastic process; therefore the positive semidefinite property will hold for all values of N.

3.7.4 Power spectral density function

For a real-valued wide-sense stationary stochastic process, we define the
Mean power as

$$\mathbb{E}[\mathbf{x}^2(t)] = R_{\mathbf{x}}(0). \tag{3.56}$$

Because of the WSS-property of the process, this notion is independent of the absolute time t.
One is generally interested in determining how the mean power of a process is distributed over different frequency components. This can be obtained by writing the mean power as an integral over frequency.

Continuous-time case
We first will consider the continuous-time situation.
In that case one is looking for a density function $\Phi_{\mathbf{x}}(\omega)$ that satisfies:

$$R_{\mathbf{x}}(0) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \Phi_{\mathbf{x}}(\omega)d\omega. \tag{3.57}$$

It can simply be verified that this expression holds true for a particular choice of function $\Phi_{\mathbf{x}}(\omega)$, namely the
Power spectral density function:

$$\Phi_{\mathbf{x}}(\omega) := \int_{-\infty}^{\infty} R_{\mathbf{x}}(\tau)e^{-i\omega\tau}d\tau \tag{3.58}$$

which of course is simply related to the autocorrelation function by Fourier transform: $R_{\mathbf{x}}(\tau) \overset{Fourier}{\Leftrightarrow} \Phi_{\mathbf{x}}(\omega)$. Expression (3.58) is known as the Wiener-Khinchine or the Wiener-Khinchine-Einstein relation.

Note that the validity of (3.57) can simply be verified by writing the inverse Fourier transform:

$$R_{\mathbf{x}}(\tau) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \Phi_{\mathbf{x}}(\omega)e^{i\omega\tau}d\omega, \tag{3.59}$$

which reduces to (3.57) by taking $\tau = 0$.
The simple fact that the definition (3.58) satisfies (3.57) does not warrant a sensible definition of spectral density; it only puts a limitation on integral properties of $\Phi_{\mathbf{x}}(\omega)$ when integrated over frequency.
However the expression (3.57) holds for *all* WSS stochastic processes, and therefore also for narrow-band processes that build up their power only in a very small frequency range. In order to examine this, consider e.g. the stochastic process

$$\mathbf{x}(t) = \sum_{i=1}^{n} a_i sin(\omega_i t + \boldsymbol{\theta}_i) \tag{3.60}$$

with $a_i \in \mathbb{R}$ given constants and $\boldsymbol{\theta}_i$ independent random variables with a uniform distribution over $[-\pi, \pi]$.

Problem 3.9 Verify that for stochastic process $\mathbf{x}$ it follows that

$$R_{\mathbf{x}}(\tau) = \sum_{i=1}^{n} \frac{1}{2} a_i^2 cos(\omega_i \tau) \tag{3.61}$$

and consequently

$$\Phi_{\mathbf{x}}(\omega) = \frac{\pi}{2} \sum_{i=1}^{n} a_i^2 [\delta_c(\omega - \omega_i) + \delta_c(\omega + \omega_i)] \tag{3.62}$$

where δ_c is the continuous-time Dirac function.

Equation (3.62) shows that the spectral density function not only leads to a total power (integrated over frequency) that is correct, but that it also has appropriate properties that are localized in frequency: a sinusoidal signal of one particular frequency contributes to the power spectral density function in terms of Dirac functions at that same frequency.

Properties of power spectral density functions

- $\Phi_{\mathbf{x}}(\omega) \geq 0$ and real-valued for all ω
- When $\mathbf{x}(t)$ is real-valued, $R_{\mathbf{x}}(\tau)$ is an even function and so $\Phi_{\mathbf{x}}(\omega)$ is even too:

$$\Phi_{\mathbf{x}}(\omega) = \Phi_{\mathbf{x}}(-\omega)$$

- If $\mathbf{x}(t)$ has periodic components then $\Phi_{\mathbf{x}}(\omega)$ contains impulse functions (Dirac functions).

Discrete-time case

In case of a discrete-time stochastic process, time is not a continuous variable but it is discretized. Following the same reasoning as above we now get the expressions

$$R_{\mathbf{x}}(0) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{\mathbf{x}}(\omega) d\omega \tag{3.63}$$

which is being realized by the spectral density function [3]

$$\boxed{\Phi_{\mathbf{x}}(\omega) := \sum_{k=-\infty}^{\infty} R_{\mathbf{x}}(k) e^{-i\omega k}} \tag{3.64}$$

being a function that is periodic in frequency with period 2π.

Again the correctness of (3.63) can be verified by considering the inverse (Discrete-Time) Fourier Transform

$$R_{\mathbf{x}}(k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{\mathbf{x}}(\omega) e^{i\omega k} d\omega \tag{3.65}$$

and substituting $k = 0$.

For an overview of the considered Fourier Transforms, see section 2.6.

The power spectral density of the process considered in example 3.6 is sketched in figure 3.19. As can be clearly observed the process with $a = 0.9$ has higher power density in the lower frequency range, whereas for $a = 0.2$ the power spectral density becomes more flat, and the higher frequencies are more dominantly present. Since both processes have the same power $\mathbb{E}[\mathbf{x}^2(t)] = R_{\mathbf{x}}(0)$ (see figure 3.17) the integral of both spectral density functions is equal. In the "visual" inspection of this equality one has to take account of the fact that in the right side plot of the figure logarithmic scales are used.

[3] For some future analysis of spectral density functions a notation $\Phi_{\mathbf{x}}(e^{i\omega})$ will appear to be more appropriate. However for simplicity of notation the expression $\Phi_{\mathbf{x}}(\omega)$ will be used here.

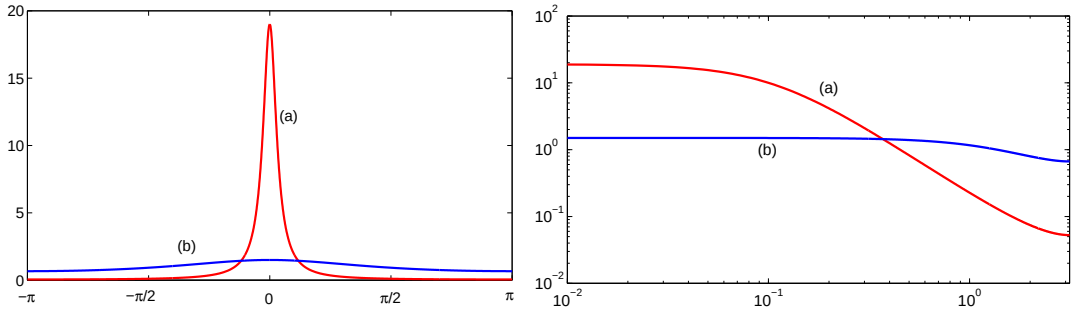

Figure 3.19: Power spectral density function $\Phi_{\mathbf{x}}(\omega)$ of process $\mathbf{x}$ (3.45) for $a = 0.9$ (a) and $a = 0.2$ (b), with linear (left) and logarithmic (right) frequency scale.

Example 3.10 (Power spectral density of a white noise process) A discrete-time white noise stochastic process $\mathbf{x}$ with variance $\sigma^2_{\mathbf{x}}$ (example 3.7) has a power spectral density function that is given by

$$\Phi_{\mathbf{x}}(\omega) = \sigma^2_{\mathbf{x}} \quad \text{for all } \omega. \tag{3.66}$$

The property of having a constant (flat) power spectral density function, is exactly the reason why this process is called "white".

3.7.5 Cross-correlation and cross-covariance functions

For two real-valued stochastic processes that are jointly WSS, the cross-correlation function will be independent of time, and therefore one can write:

$$\boxed{R_{\mathbf{yx}}(\tau) = \mathbb{E}[\mathbf{y}(t)\mathbf{x}(t-\tau)]}$$

Properties:

1. $R_{\mathbf{yx}}(\tau) = R_{\mathbf{xy}}(-\tau)$
2. $|R_{\mathbf{yx}}(\tau)| \leq \sqrt{R_{\mathbf{x}}(0)R_{\mathbf{y}}(0)}$.
   With the use of the Schwartz inequality (A.12) it follows that
   $$[\mathbb{E}(\mathbf{y}(t)\mathbf{x}(t-\tau))]^2 \leq \mathbb{E}\mathbf{y}^2(t) \cdot \mathbb{E}\mathbf{x}^2(t-\tau).$$
3. $|R_{\mathbf{yx}}(\tau)| \leq \frac{1}{2}[R_{\mathbf{x}}(0) + R_{\mathbf{y}}(0)]$
   The proof of this equation is obtained from considering
   $$\left\{ \begin{array}{lcll} \mathbb{E}[\mathbf{y}(t) - \mathbf{x}(t-\tau)]^2 & = & R_{\mathbf{x}}(0) + R_{\mathbf{y}}(0) - 2R_{\mathbf{yx}}(\tau) & \geq 0 \\ \mathbb{E}[\mathbf{y}(t) + \mathbf{x}(t-\tau)]^2 & = & R_{\mathbf{x}}(0) + R_{\mathbf{y}}(0) + 2R_{\mathbf{yx}}(\tau) & \geq 0 \end{array} \right. \Rightarrow$$
   $$\left\{ \begin{array}{lcl} R_{\mathbf{yx}}(\tau) & \leq & \frac{1}{2}[R_{\mathbf{x}}(0) + R_{\mathbf{y}}(0)] \\ R_{\mathbf{yx}}(\tau) & \geq & -\frac{1}{2}[R_{\mathbf{x}}(0) + R_{\mathbf{y}}(0)] \end{array} \right.$$
4. If $(\mathbf{x}, \mathbf{y})$ orthogonal, then $R_{\mathbf{yx}}(\tau) = 0$ for all τ.

For two stochastic processes that are jointly WSS, the cross-covariance function will also be independent of time, and therefore one can write:

$$C_{\mathbf{yx}}(\tau) = \mathbb{E}[(\mathbf{y}(t) - \mu_{\mathbf{y}})(\mathbf{x}(t-\tau) - \mu_{\mathbf{x}})]$$

If $(\mathbf{x}, \mathbf{y})$ independent or uncorrelated, then $C_{\mathbf{yx}}(\tau) = 0$ implying $R_{\mathbf{yx}}(\tau) = \mu_{\mathbf{x}}\mu_{\mathbf{y}}$ for all τ.

3.7.6 Cross-power spectral densities

We focus on the discrete time. The Fourier transform of the cross-correlation function is referred to as the cross-power spectral density:

$$\Phi_{\mathbf{yx}}(\omega) = \sum_{k=-\infty}^{\infty} R_{\mathbf{yx}}(k)e^{-i\omega k} \tag{3.67}$$

$$R_{\mathbf{yx}}(k) = \frac{1}{2\pi}\int_{-\pi}^{\pi} \Phi_{\mathbf{yx}}(\omega)e^{i\omega k}d\omega \tag{3.68}$$

Unlike the situation of the power spectral density, the cross-power spectral density is in general a complex-valued function, characterized by an amplitude and a phase function. The following properties can directly be derived from its definition:

- $\Phi_{\mathbf{yx}}(\omega)$ is complex-valued.
- $\Phi_{\mathbf{yx}}(\omega) = \Phi^*_{\mathbf{yx}}(-\omega)$ i.e., its real part is even and its imaginary part is odd.
- $\Phi_{\mathbf{yx}}(\omega) = \Phi^*_{\mathbf{xy}}(\omega)$; this is a result of the property that $R_{\mathbf{yx}}(\tau) = R_{\mathbf{xy}}(-\tau)$.
- $\Phi_{\mathbf{yx}}(\omega) = 0$ if $\mathbf{x}$ and $\mathbf{y}$ are orthogonal.

The cross-power spectral density function describes the statistical relationship between two processes $\mathbf{x}$ and $\mathbf{y}$ in the frequency domain.
The question can be raised if there is a notion of (mean) power related to

$$R_{\mathbf{yx}}(0) = \mathbb{E}[\mathbf{y}(t)\mathbf{x}(t)]$$

similar to the situation of power spectral densities?
In order to illustrate this, let $\mathbf{y}(t) = \alpha\mathbf{x}(t) + \mathbf{w}(t)$ with $\mathbb{E}[\mathbf{x}(t)\mathbf{w}(t)] = 0$ and α a scalar. Then

$$\mathbb{E}[\mathbf{y}(t)\mathbf{x}(t)] = \alpha \cdot \mathbb{E}[\mathbf{x}^2(t)] + \mathbb{E}[\mathbf{x}(t)\mathbf{w}(t)] \tag{3.69}$$

$$= \alpha \cdot \mathbb{E}[\mathbf{x}^2(t)] = \mathbb{E}[\sqrt{\alpha} \cdot \mathbf{x}(t)]^2 \tag{3.70}$$

is the power of $\sqrt{\alpha}\cdot\mathbf{x}(t)$, a process that is the square root of the product (i.e., the geometric "mean") of $\mathbf{x}(t)$ and $\alpha \cdot \mathbf{x}(t)$. This latter process actually is the "projection" of $\mathbf{y}(t)$ onto $\mathbf{x}(t)$.

Cross-correlation functions and cross-power spectral densities are used to characterize statistical relationships between two stochastic processes. Consider, e.g., the vibration behaviour of a car when it travels on a bumpy road. Since the source of the vibration (the particular road profile) will essentially generate all vibrations on several locations on the car frame, the acceleration signals that can be measured from different parts of the frame will show statistical dependencies.

3.8 Time-averaging and ergodicity

In this section we briefly consider the relation between averaging over time (of one time signal), and averaging over the ensemble.

Consider for example a constant c which is measured at several time instants, subject to a zero-mean disturbance signal $w(t)$, leading to a time signal

$$x(t) = c + w(t).$$

In order to estimate c, one can e.g. just pick a measurement at any time instant t_1, i.e., $\hat{c} = x(t_1)$. However since $\mathbf{w}(t)$ is zero-mean, one may expect that averaging over time, i.e.,

$$\hat{c} = \frac{1}{m}\sum_{i=1}^{m} x(t_i)$$

will lead to an estimate that is much more reliable.

The averaging over ensemble samples will in general be quite complicated. It requires an (infinite) repetition of experiments. In general one will need to estimate properties of a stochastic process on the basis of a single measured time series. The question will be considered under which conditions separate time series contain characteristic information on the underlying stochastic process.

The analysis of time-averaged estimators for first and second order moments of WSS stochastic processes will be analyzed in detail in one of the future chapters. Here only brief attention will be given to some general notions and concept definitions.

We consider the following definitions and notation for any function g of a stochastic process:

Time averages:

$$< g[\mathbf{x}(t)] >_T \;:=\; \frac{1}{T}\int_{-T/2}^{T/2} g[\mathbf{x}(t)]\; dt \quad \text{continuous-time} \tag{3.71}$$

$$:=\; \frac{1}{m}\sum_{i=1}^{m} g[\mathbf{x}(i)] \quad \text{discrete-time} \tag{3.72}$$

Ensemble-means:

$$\mathbb{E}\left[g[\mathbf{x}(t)]\right] \;:=\; \int_{-\infty}^{\infty} g(\alpha) f_{\mathbf{x}}(\alpha) d\alpha \quad \text{continuous r.v.} \tag{3.73}$$

$$:=\; \sum_i g(x_i) P(\mathbf{x} = x_i) \quad \text{discrete r.v.} \tag{3.74}$$

Time-averaged mean:

$$\langle \mathbf{x}(t) \rangle_T \;:=\; \langle \mu_{\mathbf{x}} \rangle_T \;:=\; \frac{1}{T}\int_{-T/2}^{T/2} \mathbf{x}(t) dt \quad \text{continuous-time} \tag{3.75}$$

$$\langle \mathbf{x}(t) \rangle_N \;:=\; \langle \mu_{\mathbf{x}} \rangle_N \;:=\; \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}(i) \quad \text{discrete-time} \tag{3.76}$$

Time-averaged autocorrelation: (sample autocorrelation)

$$\langle \mathbf{x}(t)\mathbf{x}(t-\tau)\rangle_T \;:=\; \langle R_{\mathbf{x}}(\tau)\rangle_T := \frac{1}{T}\int_{-T/2}^{T/2} \mathbf{x}(t)\mathbf{x}(t-\tau)dt \tag{3.77}$$

$$\langle \mathbf{x}(t)\mathbf{x}(t-k)\rangle_N \;:=\; \langle R_{\mathbf{x}}(k)\rangle_N := \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}(i)\mathbf{x}(i-k) \tag{3.78}$$

Note that a principal difference between time-averages and ensemble-means is that a time-average is a random variable, while an ensemble-mean is a fixed (deterministic) quantity.

Example 3.11 Consider the stochastic process containing 6 signals:

$$x_1(t) = 5;\; x_2(t) = 3;\; x_3(t) = 1;\; x_4(t) = -1;\; x_5(t) = -3;\; x_6(t) = -5,$$

while the probability of each realization is 1/6. For this process, $\mathbb{E}\mathbf{x}(t) = 0$ for all t, but time-averaging delivers a different value of $\langle \mathbf{x}(t)\rangle_T$ for every realization.
Consequently a single time series from this stochastic process does not contain full information on the stochastic process.

Stochastic processes of which each time signal has the characteristics of the full stochastic process are referred to as ergodic processes.

Ergodicity
A stationary stochastic process $\mathbf{x}(t)$ is *ergodic* if its ensemble means are equal (in mean square sense) to corresponding time-averages.
This implies:

- Expected valued of time-average is equal to the ensemble-mean, and
- The variance of the time-average tends to 0 for $N \to \infty$.

The following two forms of ergodicity are distinguished:

- Ergodicity in the mean, characterized by the properties that

$$\mathbb{E}\left[\langle \mathbf{x}(t)\rangle_N\right] = \mathbb{E}[\mathbf{x}], \quad \text{and } var\{\langle \mathbf{x}(t)\rangle_N\} \underset{N\to\infty}{\longrightarrow} 0$$

- Ergodicity in the autocorrelation function, characterized by the property that

$$\mathbb{E}\left[\langle \mathbf{x}(t)\mathbf{x}(t-k)\rangle_N\right] = R_{\mathbf{x}}(k), \quad \text{and } var\{\langle \mathbf{x}(t)\mathbf{x}(t-k)\rangle_N\} \underset{N\to\infty}{\longrightarrow} 0$$

Questions concerning ergodicity actually reflect whether estimates of characteristic properties of stochastic processes that are based on averaging over one realization of the process, have the desirable property that they provide accurate knowledge of the underlying process itself. More will be said about this when discussing estimation procedures for stochastic process models.

3.9 Summary

In this chapter a probabilistic modelling framework is presented for characterizing signals that or not exactly reproducible, i.e., signals that typically vary under experiment repetition. The framework of stochastic processes has been introduced, leading to the modelling of a measurement signal as one realization of a whole ensemble of realizations that are considered as possible outcomes of the measurement experiment. Mean values and autocorrelation functions are the important notions for characterizing typical properties of the underlying process. Two different signals/processes can be related through their interdependence that appears in the cross-correlation function. The power spectral density function characterizes which frequency regions contribute to the power of the signals, and this gives input to the design of filters that are designed towards influencing this power distribution.
Linear filtering of stochastic processes is the topic of the next chapter.

Chapter 4

Linear filtering of stochastic processes

Like deterministic signals, stochastic signals can be processed by linear dynamic filters. Filters allow to influence the characteristic properties of the signals with interpretations in both time and frequency domain. Additionally the concept of linear filters can be used to model the dynamic properties of stochastic processes as the output of a linear filter being excited by white noise.

4.1 Stochastic processes related through linear dynamical systems

One of the basic operations that can be applied to a stochastic process is that of filtering the stochastic process with a linear time-invariant finite-dimensional dynamical system. This means that a given stochastic process $\mathbf{u}(t)$ is applied as input to a dynamical system G, leading to an output stochastic process $\mathbf{y}$, as schematically depicted in the block diagram in figure 4.1.

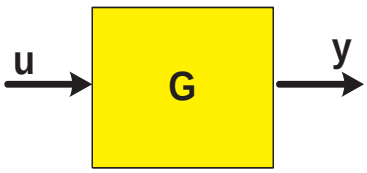

Figure 4.1: Filtering a stochastic process $\mathbf{u}$ by a linear dynamical system G.

As in the case of deterministic signals, one can shape the properties of the output signal/process by the appropriate design of the system G.
Following the notation and concepts for linear discrete-time dynamical systems, as presented in section 2.5, a linear discrete-time causal system can be represented by the convolution expression

$$y(t) = \sum_{k=0}^{\infty} g(k)u(t-k) \qquad t = 0, 1, 2, \cdots \tag{4.1}$$

where $\{g(k)\}_{k=0,1,\cdots}$ is the pulse response of the system, and y and u are (deterministic) output and input signals.
For deterministic signals for which a Fourier transform exists[1], the input-output relation (4.1) can be written in the frequency domain as:

$$Y(e^{i\omega}) = G(e^{i\omega})U(e^{i\omega}) \tag{4.2}$$

with $Y(e^{i\omega})$, $U(e^{i\omega})$ and $G(e^{i\omega})$ determined by the Fourier/z-transform:

$$\begin{aligned} Y(z) &= \sum_{k=-\infty}^{\infty} y(k)z^{-k} \\ U(z) &= \sum_{k=-\infty}^{\infty} u(k)z^{-k} \\ G(z) &= \sum_{k=0}^{\infty} g(k)z^{-k} \end{aligned} \tag{4.3}$$

For random input signals (stochastic processes) these relations can not be used straightforwardly, as the Fourier transform may not exist for these signals.
However the convolutional description (4.1) remains applicable. This implies that for a causal filter G with transfer function $G(z)$, an input-output relationship is determined given by

$$\boxed{\mathbf{y}(t) = \sum_{k=0}^{\infty} g(k)\mathbf{u}(t-k) \quad t = 0, 1, 2, \cdots} \tag{4.4}$$

specifying a stochastic process $\mathbf{y}(t)$, based upon the input stochastic process $\mathbf{u}(t)$, while $g(k)$ is defined by (4.3). In this way a transfer function $G(z)$ can be used to describe the relationship between two stochastic processes.
Next some of the most important properties of the stochastic process $\mathbf{y}(t)$ are discussed.

Wide sense stationarity of y
For the mean value $\mathbb{E}[\mathbf{y}(t)]$ it follows directly that

$$\mathbb{E}[\mathbf{y}(t)] = \sum_{k=0}^{\infty} g(k) \cdot \mathbb{E}[\mathbf{u}(t-k)]. \tag{4.5}$$

If $\mu_{\mathbf{u}} = \mathbb{E}[\mathbf{u}(t)]$ is independent of time t, this reduces to

$$\mathbb{E}[\mathbf{y}(t)] = \sum_{k=0}^{\infty} g(k) \cdot \mu_{\mathbf{u}} = G(z)|_{z=1} \cdot \mu_{\mathbf{u}}. \tag{4.6}$$

A direct consequence is that if $\mathbf{u}(t)$ is a zero mean stochastic process, then $\mathbf{y}(t)$ is also zero-mean. If $\mathbf{u}(t)$ has a non-zero mean value, this mean value gets amplified by the factor $G(z)|_{z=1}$, which is the *static gain* of the system G.

[1] The requirements for a discrete-time Fourier transform of a signal x to exist, are that either $\sum_{k=-\infty}^{\infty} |x(k)| < \infty$ or that $\sum_{k=-\infty}^{\infty} |x(k)|^2 < \infty$.

For the autocorrelation function we write:

$$R_{\mathbf{y}}(t_1, t_2) = \mathbb{E}[\mathbf{y}(t_1)\mathbf{y}(t_2)] = \tag{4.7}$$

$$= \mathbb{E}\{\sum_{k=0}^{\infty} g(k)\mathbf{u}(t_1 - k) \sum_{\ell=0}^{\infty} g(\ell)\mathbf{u}(t_2 - \ell)] \tag{4.8}$$

$$= \sum_{k=0}^{\infty}\sum_{\ell=0}^{\infty} g(k)g(\ell) \cdot R_{\mathbf{u}}(t_1 - k, t_2 - \ell). \tag{4.9}$$

If $R_{\mathbf{u}}(t_1, t_2)$ is independent of time, but dependent on $t_1 - t_2$ only, then

$$R_{\mathbf{y}}(t_1, t_2) = \sum_{k=0}^{\infty}\sum_{\ell=0}^{\infty} g(k)g(\ell) \cdot R_{\mathbf{u}}(t_1 - t_2 + \ell - k) \tag{4.10}$$

which shows that in that case $R_{\mathbf{y}}(t_1, t_2)$ is a also function of $t_1 - t_2$ only, and therefore is independent of absolute time.

The results for mean value and autocorrelation function directly imply that if $\mathbf{u}(t)$ is wide sense stationary (WSS), then also $\mathbf{y}(t)$ is WSS.

Distribution functions

In the general situation it is very hard to make statements about the distribution functions of $\mathbf{y}(t)$ on the basis of distribution functions of $\mathbf{u}(t)$. Even in the situation that the pulse response of the dynamical system has finite length, and $\mathbf{y}(t)$ simply becomes a finite linear combination of elements from $\mathbf{u}(t)$, the resulting distribution functions can not generally be expressed in closed–form / analytical expressions. There is an important exception for the situation of a process $\mathbf{u}(t)$ having Gaussian distribution functions. In that case $\mathbf{y}(t)$ actually is a linear combination of Gaussian distributions, which again is a Gaussian distribution, according to appendix A.3.2.

Correlation functions and power spectral density

If $\mathbf{u}(t)$ (and consequently also $\mathbf{y}(t)$) are wide sense stationary, then the following relationships for the correlation function are found:

$$R_{\mathbf{yu}}(\tau) = \mathbb{E}[\mathbf{y}(t)\mathbf{u}(t - \tau)] \tag{4.11}$$

$$= \mathbb{E}[\sum_{k=0}^{\infty} g(k)\mathbf{u}(t - k)\mathbf{u}(t - \tau)] \tag{4.12}$$

$$= \sum_{k=0}^{\infty} g(k)R_{\mathbf{u}}(\tau - k) = g(\tau) \star R_{\mathbf{u}}(\tau) \tag{4.13}$$

$$R_{\mathbf{uy}}(\tau) = \mathbb{E}[\mathbf{u}(t)\mathbf{y}(t - \tau)] \tag{4.14}$$

$$= \mathbb{E}[\sum_{k=0}^{\infty} g(k)\mathbf{u}(t)\mathbf{u}(t - \tau - k)] \tag{4.15}$$

$$= \sum_{k=0}^{\infty} g(k)R_{\mathbf{u}}(\tau + k) = g(-\tau) \star R_{\mathbf{u}}(\tau) \tag{4.16}$$

$$
\begin{aligned}
R_{\mathbf{y}}(\tau) &= \mathbb{E}[\mathbf{y}(t)\mathbf{y}(t-\tau)] && (4.17)\\
&= \mathbb{E}[\sum_{k=0}^{\infty} g(k)\mathbf{u}(t-k)\mathbf{y}(t-\tau)] && (4.18)\\
&= \sum_{k=0}^{\infty} g(k)R_{\mathbf{uy}}(\tau-k) = g(\tau) \star R_{\mathbf{uy}}(\tau) && (4.19)
\end{aligned}
$$

$$
\begin{aligned}
R_{\mathbf{y}}(\tau) &= g(\tau) \star R_{\mathbf{uy}}(\tau) && (4.20)\\
&= g(\tau) \star g(-\tau) \star R_{\mathbf{u}}(\tau). && (4.21)
\end{aligned}
$$

As a result we have the following expressions:

- $R_{\mathbf{yu}}(\tau) = g(\tau) \star R_{\mathbf{u}}(\tau)$
- $R_{\mathbf{uy}}(\tau) = g(-\tau) \star R_{\mathbf{u}}(\tau)$
- $R_{\mathbf{y}}(\tau) = g(\tau) \star g(-\tau) \star R_{\mathbf{u}}(\tau)$.

Note that the dynamical system G not only maps the input process $\mathbf{u}(t)$ to the output process $\mathbf{y}(t)$, it constitutes in a similar way the relation between $R_{\mathbf{u}}(t)$ and $R_{\mathbf{yu}}(t)$, as depicted in figure 4.2.

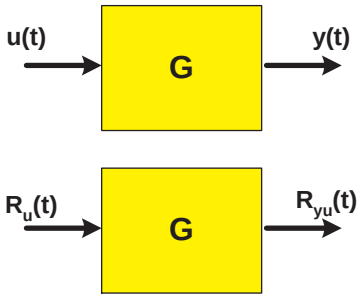

Figure 4.2: System G relates the processes $\mathbf{u}(t)$ and $\mathbf{y}(t)$, and similarly the correlation functions $R_{\mathbf{u}}(t)$ and $R_{\mathbf{yu}}(t)$.

By combining the operation of G and its complex conjugate form G^*, the linear relation is obtained between $R_{\mathbf{u}}(t)$ and $R_{\mathbf{y}}(t)$, as sketched in figure 4.3.

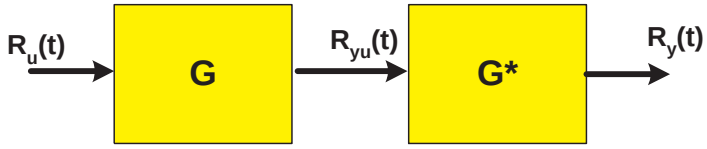

Figure 4.3: The dynamical system relation between $R_{\mathbf{u}}(t)$ and $R_{\mathbf{y}}(t)$ is represented by a concatenation of system G and its complex conjugate form G^*.

Consequences for the relationships between spectral density functions of $\mathbf{u}(t)$ and $\mathbf{y}(t)$ can now simply be obtained by applying Fourier transformation to the expressions for the

autocorrelation functions. Use has to be made of the Fourier transform properties that for two functions g and f, $g(\tau)\star f(\tau)$ transforms to $G(e^{i\omega})\cdot F(e^{i\omega})$ and $g(-\tau)\star f(\tau)$ transforms to $G(e^{-i\omega})\cdot F(e^{i\omega})$. Consequently

$$\begin{aligned}\Phi_{\mathbf{y}}(\omega) &= G(e^{-i\omega})G(e^{i\omega})\Phi_{\mathbf{u}}(\omega) = G^*(e^{i\omega})G(e^{i\omega})\Phi_{\mathbf{u}}(\omega) &(4.22)\\ \Phi_{\mathbf{y}}(\omega) &= |G(e^{i\omega})|^2\cdot\Phi_{\mathbf{u}}(\omega) &(4.23)\end{aligned}$$

and through the same reasoning:

$$\begin{aligned}\Phi_{\mathbf{yu}}(\omega) &= G(e^{i\omega})\cdot\Phi_{\mathbf{u}}(\omega) &(4.24)\\ \Phi_{\mathbf{uy}}(\omega) &= G(e^{-i\omega})\cdot\Phi_{\mathbf{u}}(\omega). &(4.25)\end{aligned}$$

Note that the first equation provides a means to construct the dynamical relation $G(e^{i\omega})$ between $\mathbf{u}$ and $\mathbf{y}$ on the basis of signal properties, by way of:

$$G(e^{i\omega}) = \frac{\Phi_{\mathbf{yu}}(\omega)}{\Phi_{\mathbf{u}}(\omega)}. \tag{4.26}$$

If the auto- and cross-spectral densities of the input and output process are known, or estimated, the dynamical relation between the two processes can be constructed. This powerful tool will be used later when discussing the estimation of spectral densities and transfer functions.

Example 4.1 Let $\mathbf{u}$ be a WSS stochastic process with

$$\begin{aligned}\mu_{\mathbf{u}} &= 0 &(4.27)\\ R_{\mathbf{u}}(\tau) &= 1 \quad \tau = 0 &(4.28)\\ &= 0 \quad \tau \neq 0. &(4.29)\end{aligned}$$

This is a white noise process.
Consider a linear filter G with the properties

$$\begin{aligned}g(\tau) &= 1 \quad \tau = 0, 1 &(4.30)\\ &= 0 \quad \text{elsewhere.} &(4.31)\end{aligned}$$

The pulse response of G is depicted in figure 4.4.
Let $\mathbf{y}(t) = \sum_{k=0}^{\infty} g(k)\mathbf{u}(t-k)$. Determine $\mu_{\mathbf{y}}$, $R_{\mathbf{y}}$ and $\Phi_{\mathbf{y}}$.

Taking the mean value of the convolutional expression for $\mathbf{y}(t)$ it follows that

$$\begin{aligned}\mu_{\mathbf{y}} &= \mathbb{E}[\sum_{k=0}^{\infty} g(k)\mathbf{u}(t-k)] = \sum_{k=0}^{\infty} g(k)\mu_{\mathbf{u}} = \\ &= \mu_{\mathbf{u}}\cdot G(1) = 0. &(4.32)\end{aligned}$$

$$\Phi_{\mathbf{u}}(\omega) = \sum_{k=-\infty}^{\infty} R_{\mathbf{u}}(k)e^{-i\omega k} = 1$$

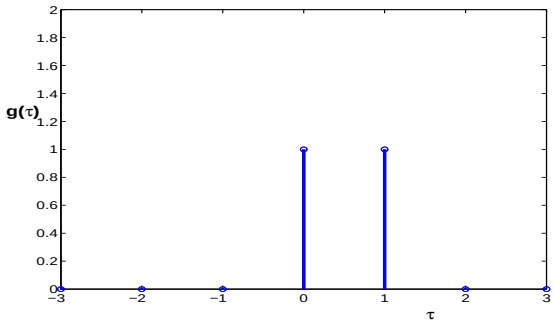

Figure 4.4: Pulse response $g(\tau)$ of system G.

Since $G(z) = 1 + z^{-1}$, it follows that $G(e^{i\omega}) = 1 + e^{-i\omega}$ and so

$$|G(e^{i\omega})|^2 = [1 + e^{-i\omega}][1 + e^{i\omega}] = 2 + 2cos(\omega).$$

As a result

$$\Phi_{\mathbf{y}}(\omega) = |G(e^{i\omega})|^2 \cdot \Phi_{\mathbf{u}}(\omega) = 2[1 + cos(\omega)]$$

$R_{\mathbf{y}}(k)$ can simply be obtained by inverse Fourier transform of $\Phi_{\mathbf{y}}(\omega)$. Since

$$\Phi_{\mathbf{y}}(\omega) = e^{-i\omega} + 2 + e^{i\omega}$$

and $\Phi_{\mathbf{y}}(\omega) = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k)e^{-i\omega k}$ it follows that

$$\begin{aligned} R_{\mathbf{y}}(k) &= 1 \quad k = \pm 1 && (4.33)\\ &= 2 \quad k = 0 && (4.34)\\ &= 0 \quad \text{elsewhere.} && (4.35) \end{aligned}$$

The autocorrelation function and the spectral density function of $\mathbf{y}(t)$ are sketched in figure 4.5.

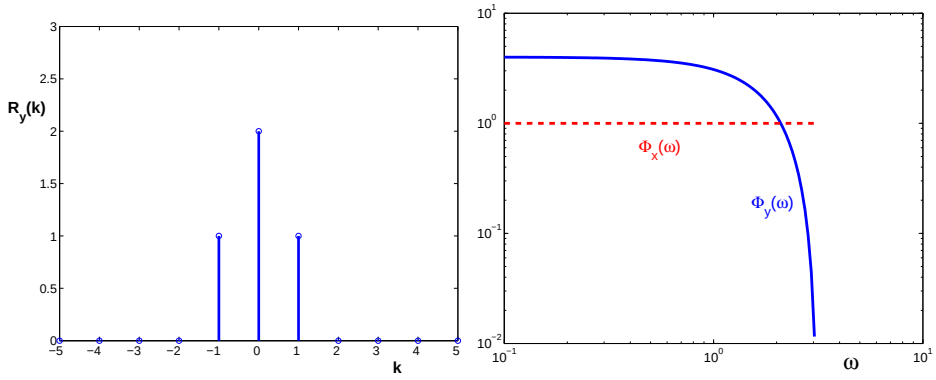

Figure 4.5: Left: Autocorrelation function $R_{\mathbf{y}}(k)$, and Right: spectral density function $\Phi_{\mathbf{y}}(\omega)$ (solid, blue) and $\Phi_{\mathbf{u}}(\omega)$ (dashed, red) of the processes in example 4.1.

Example 4.1 clearly illustrates how a dynamical system can influence the power spectral density properties of the output process. Whereas the input process is a white noise with

a spectral density that is constant over frequency, the output process shows an increased power contribution of the low-frequency range, while for higher frequencies the power density is reduced. This phenomenon is due to the low-pass characteristic of the dynamical system G, of which the squared amplitude $|G(e^{i\omega})|^2$ exactly describes the power spectral density of the output $\mathbf{y}(t)$.

The relations that apply for cross-correlation functions and cross-spectral density functions also extend to other filter operations. If two WSS stochastic processes $\mathbf{x}$ and $\mathbf{u}$ have cross-correlation function $R_{\mathbf{ux}}(\tau)$ and cross-spectral density function $\Phi_{\mathbf{ux}}(\omega)$ and if $\mathbf{y} = G\mathbf{u}$, then

$$\begin{aligned} R_{\mathbf{yx}}(\tau) &= g(\tau) \star R_{\mathbf{ux}}(\tau) && (4.36)\\ \Phi_{\mathbf{yx}}(\omega) &= G(e^{i\omega}) \cdot \Phi_{\mathbf{ux}}(\omega). && (4.37) \end{aligned}$$

Similarly, if $\mathbf{v} = G\mathbf{x}$ then

$$\begin{aligned} R_{\mathbf{uv}}(\tau) &= g(-\tau) \star R_{\mathbf{ux}}(\tau) && (4.38)\\ \Phi_{\mathbf{uv}}(\omega) &= G(e^{-i\omega}) \cdot \Phi_{\mathbf{ux}}(\omega). && (4.39) \end{aligned}$$

If there are two dynamical systems G_1, G_2 as depicted in Figure 4.6, that generate $\mathbf{y}_1$, $\mathbf{y}_2$, then application of the formulas above lead to the results:

$$\mathbf{y}_2 = G_2 G_1^{-1} \mathbf{y}_1$$

and therefore

$$\begin{aligned} \Phi_{\mathbf{y}_2\mathbf{y}_1} &= G_2 G_1^{-1} \Phi_{\mathbf{y}_1} && (4.40)\\ &= G_2 G_1^{-1} |G_1|^2 \Phi_{\mathbf{u}} = G_2 G_1^{-1} [G_1 G_1^*] \Phi_{\mathbf{u}} && (4.41)\\ &= G_2 G_1^* \Phi_{\mathbf{u}} && (4.42) \end{aligned}$$

so that

$$\Phi_{\mathbf{y}_2\mathbf{y}_1}(\omega) = G_2(e^{i\omega}) G_1(e^{-i\omega}) \Phi_{\mathbf{u}}(\omega).$$

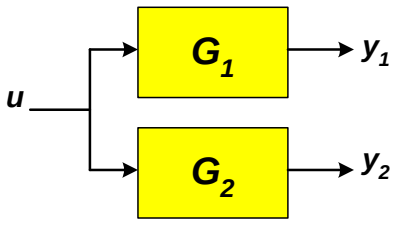

Figure 4.6: Two signals $\mathbf{y}_1$, $\mathbf{y}_2$ originating from the same source signal $\mathbf{u}$.

Covariance functions of $\mathbf{y}(t)$ and $\mathbf{u}(t)$

When considering a WSS stochastic process with non-zero mean $\mu_{\mathbf{u}}$, the systems relation (4.4) can be supplemented by adding a term $-\sum_{k=0}^{\infty} g(k)\mu_{\mathbf{u}}$ to both sides of the equation, leading to

$$\mathbf{y}(t) - \sum_{k=0}^{\infty} g(k)\mu_{\mathbf{u}} = \sum_{k=0}^{\infty} g(k)[\mathbf{u}(t-k) - \mu_{\mathbf{u}}]. \quad (4.43)$$

Since the process $\mathbf{u}(t) - \mu_{\mathbf{u}}$ is zero-mean, the resulting output $\mathbf{y}(t) - \sum_{k=0}^{\infty} g(k)\mu_{\mathbf{u}}$ is necessarily zero-mean also, showing that

$$\mu_{\mathbf{y}} = \sum_{k=0}^{\infty} g(k)\mu_{\mathbf{u}}.$$

Then

$$\mathbf{y}(t) - \mu_{\mathbf{y}} = \sum_{k=0}^{\infty} g(k)[\mathbf{u}(t-k) - \mu_{\mathbf{u}}].$$

Exactly the same line of reasoning as applied to the autocorrelation functions can now be followed to arrive at the same expressions for the autocovariance functions. In this case one considers squared forms and products of $\mathbf{y}(t) - \mu_{\mathbf{y}}$ and $\mathbf{u}(t-k) - \mu_{\mathbf{u}}$, leading to:

- $C_{\mathbf{yu}}(\tau) = g(\tau) \star C_{\mathbf{u}}(\tau)$
- $C_{\mathbf{uy}}(\tau) = g(-\tau) \star C_{\mathbf{u}}(\tau)$
- $C_{\mathbf{y}}(\tau) = g(\tau) \star g(-\tau) \star C_{\mathbf{u}}(\tau)$

4.2 Spectral factorization

In the previous section it has been indicated how a dynamical system G determines the spectral density $\Phi_{\mathbf{y}}(\omega)$ of the output signal, starting from a unit-variance white noise input process, leading the relation

$$\Phi_{\mathbf{y}}(\omega) = |G(e^{i\omega})|^2. \tag{4.44}$$

In this section the converse route will be followed by considering the question, which stable and causal dynamical system will generate a prespecified spectral density function as output process, when driven by unit variance white noise as input.
When considering this question one first has to specify which class of spectral density functions is being considered; e.g., any positive continuous function in ω, or any positive rational function in $e^{i\omega}$. In the sequel of this section the second case will be studied, i.e. assuming that $\Phi_{\mathbf{y}}(\omega)$ as any positive rational (and therefore continuous) function in $e^{i\omega}$. This excludes spectral density functions that contain Dirac functions.
As already appearing in (4.44) it seems rather natural that $\Phi_{\mathbf{y}}(\omega)$ can also be written in terms of the argument $e^{i\omega}$ rather than ω. In this section it will appear attractive (and actually necessary) to write spectral densities as $\Phi_{\mathbf{y}}(e^{i\omega})$ rather than $\Phi_{\mathbf{y}}(\omega)$. This allows also to generalize spectral densities to be written as functions over the complex indeterminate z. In line with (3.64) we will denote

$$\Phi_{\mathbf{y}}(z) = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k) z^{-k} \tag{4.45}$$

while the power spectral density function is obtained by evaluating $\Phi_{\mathbf{y}}(z)$ over $z = e^{i\omega}$. When appropriate, $\Phi_{\mathbf{y}}(z)$ will be referred to as the complex power spectral density function.

Similarly, in line with (3.67), complex cross-power spectral density functions can be defined as

$$\Phi_{\mathbf{yx}}(z) = \sum_{k=-\infty}^{\infty} R_{\mathbf{yx}}(k) z^{-k}, \tag{4.46}$$

where the cross-power spectral density function is obtained by evaluating $\Phi_{\mathbf{yx}}(z)$ over $z = e^{i\omega}$. Note that, unlike $\Phi_{\mathbf{yx}}(e^{i\omega})$, $\Phi_{\mathbf{y}}(e^{i\omega})$ will always be real-valued though, as is formulated as one of the properties of a spectral density function in section 3.7.4. Through the relation (4.44) it appears that the transfer function G only influences the spectral density $\Phi_{\mathbf{y}}$ through its amplitude function. The phase properties of G do not appear in $\Phi_{\mathbf{y}}$. In order to handle this indeterminateness of the phase the notion of minimum-phase systems is considered.

Definition 4.2 (Minimum-phase system) A linear dynamical system is called minimum-phase if the system and its inverse are causal and stable.

A rational transfer function $H(z)$ is representing a minimum-phase system if it has all poles and zeros strictly inside the unit circle. In that case

- $H(z)$ has a series expansion
$$H(z) = \sum_{k=0}^{\infty} h(k) z^{-k}$$
that is convergent for $|z| \geq 1$, i.e. $\sum_{k=0}^{\infty} |h(k)| < \infty$.
- $H(z)^{-1}$ also has all zeros and poles strictly inside the unit circle as can simply be verified by realizing that in the systems inverse zero's and pole's interchange roles. Therefore $H(z)^{-1}$ also has a stable and causal expansion.

Minimum-phase systems have a "minimum-phase" according to the following interpretation: when the amplitude of a system frequency response is fixed, the pole-zero locations related to the transfer function are not uniquely determined. E.g., pole- and zero-locations can be mirrored with respect to the unit circle without affecting the shape of the amplitude frequency response. See, e.g., the construction of the amplitude frequency response as sketched in Figure 2.2. It is more easily observed in the continuous-time case, where zeros and poles can be mirrored with respect to the imaginary axis $s = i\omega$ without affecting the amplitude function. Since we will generally deal with causal and stable filters, it is convenient to restrict the pole locations to be inside the unit circle. The zeros however are not determined yet. When a system zero is mirrored from a location outside the unit circle z_0 to a related location $1/z_0^*$ within the unit circle, the phase contribution of this zero is reduced. Therefore a system with all zeros within the unit circle is referred to as a "minimum-phase system".
Next attention has to be paid to the class of (complex) power spectral density functions that are rational (i.e. a fraction of polynomials) in z. The following property is shown.

Proposition 4.3 For a power spectral density function $\Phi_{\mathbf{y}}(z)$ that is rational, it holds that all poles and zeros of $\Phi_{\mathbf{y}}(z)$ come in complex conjugate reciprocal pairs, i.e.

- If z_0 is a pole of $\Phi_{\mathbf{y}}(z)$ then $1/z_0^*$ is also a pole of $\Phi_{\mathbf{y}}(z)$;
- If z_0 is a zero of $\Phi_{\mathbf{y}}(z)$ then $1/z_0^*$ is also a zero of $\Phi_{\mathbf{y}}(z)$.

Proof: The result is direct from (4.45) by the following reasoning:
Because $R_{\mathbf{y}}(k)$ is an even function, $R_{\mathbf{y}}(k) = R_{\mathbf{y}}(-k)$, and therefore

$$\Phi_{\mathbf{y}}(z^{-1}) = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k) z^{k} = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(-k) z^{k} = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k) z^{-k} = \Phi_{\mathbf{y}}(z).$$

Additionally because $R_{\mathbf{y}}(k)$ is real-valued, it follows that

$$\Phi_{\mathbf{y}}(z^{*}) = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k) (z^{*})^{-k} = \sum_{k=-\infty}^{\infty} R_{\mathbf{y}}(k) (z^{-k})^{*} = \Phi_{\mathbf{y}}^{*}(z).$$

Combining the two expressions leads to $\Phi_{\mathbf{y}}(z) = \Phi_{\mathbf{y}}^{*}(1/z^{*})$. This shows that any pole/zero of $\Phi_{\mathbf{y}}$ will be accompanied by a pole/zero in its complex conjugate reciprocal location. □

Since poles and zeros of $\Phi_{\mathbf{y}}(z)$ also come in (ordinary) complex conjugate pairs, i.e. z_0 and z_0^{*}, due to the fact that the series expansion of $\Phi_{\mathbf{y}}(z)$ (4.45) has real-valued coefficients, every pole and zero of $\Phi_{\mathbf{y}}(z)$ appears four times. This is illustrated in the plot of figure 4.7.

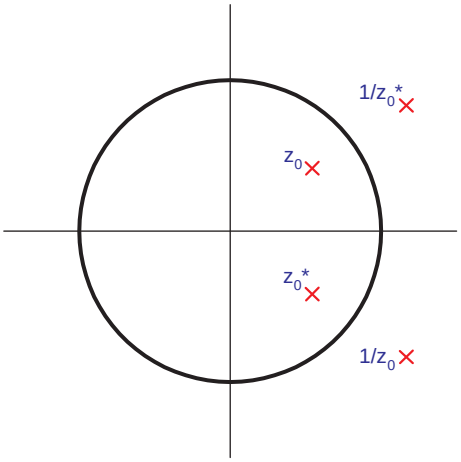

Figure 4.7: Repetition of pole locations of $\Phi_{\mathbf{y}}(z)$ in z_0, $1/z_0$, z_0^{*} and $1/z_0^{*}$.

Now sufficient ingredients have been collected to solve our originally posed problem in terms of the spectral factorization theorem.

Theorem 4.4 (Spectral Factorization) For any rational spectral density function[2] $\Phi_{\mathbf{y}}(z)$ that is analytic (i.e. has no poles) on the unit circle, there exists a minimum-phase system represented by $H_m(z)$, and a scalar σ^2, such that

$$\Phi_{\mathbf{y}}(z) = \sigma^2 \cdot H_m(z) H_m(1/z).$$

If additionally $H_m(z)$ is restricted to be monic, i.e.

$$H_m(z) = 1 + \sum_{k=1}^{\infty} h_m(k) z^{-k}$$

then this factorization is unique.

[2] Actually the theorem applies also to nonrational spectral density functions, provided that the Paley-Wiener condition $\int_{-\pi}^{\pi} |ln \Phi_{\mathbf{y}}(e^{i\omega})| d\omega < \infty$ is satisfied.

Whereas $H_m(z)$ has all its poles and zeros strictly inside the unit circle, $H_m(1/z)$ will necessarily have all its poles and zeros strictly outside the unit circle.
Using the pole/zero location symmetry as formulated in proposition 4.3, a procedure for the construction of H_m is simply obtained as follows:

- Evaluate the pole/zero locations of $\Phi_{\mathbf{y}}(z)$;
- Compose $H_m(z)$ be taking all the poles and zeros of $\Phi_{\mathbf{y}}(z)$ that are located strictly inside the unit circle;
- By proposition 4.3, all poles and zeros that remain are necessarily located outside the unit disc at the inverse locations of the selected poles/zeros;
- Adjust the positive scalar σ^2 to satisfy $\Phi_{\mathbf{y}}(z) = \sigma^2 H_m(z)H_m(1/z)$.

Example 4.5 Consider the spectral density function

$$\Phi_{\mathbf{y}}(e^{i\omega}) = \frac{1 - 2.5e^{-i\omega} + e^{-2i\omega}}{1 - 2.05e^{-i\omega} + e^{-2i\omega}}.$$

First we write the function in its z-domain representation:

$$\Phi_{\mathbf{y}}(z) = \frac{1 - 2.5z^{-1} + z^{-2}}{1 - 2.05z^{-1} + z^{-2}},$$

which can be factored as:

$$\Phi_{\mathbf{y}}(z) = \frac{(1 - 0.5z^{-1})(1 - 2z^{-1})}{(1 - 0.8z^{-1})(1 - 1.25z^{-1})} = \frac{(1 - 0.5z^{-1})}{(1 - 0.8z^{-1})} \cdot \frac{(1 - 2z^{-1})}{(1 - 1.25z^{-1})}.$$

The first factor has a pole in $z = 0.8$ and a zero in $z = 0.5$; the second factor has a pole in $z = 1.25 = 1/0.8$, and a zero in $z = 2 = 1/0.5$.
As a result the minimum-phase spectral factor of $\Phi_{\mathbf{y}}(z)$ is

$$H_m(z) = \frac{(1 - 0.5z^{-1})}{(1 - 0.8z^{-1})}.$$

The minimum-phase spectral factor of a spectral density function $\Phi_{\mathbf{y}}$ has a particular and important property. First it serves as the causal and stable linear dynamical system, that, when using a white noise as input process, generates an output process with the specified spectral density function $\Phi_{\mathbf{y}}(e^{i\omega})$.
However secondly, through the inverse $1/H_m$ which also is minimum-phase by construction, it generates a stable and causal linear dynamical system, that, when using $\mathbf{y}$ as an input process, generates a white noise output process. This representation is called the *innovations representation* of a WSS stochastic process $\mathbf{y}$, and is illustrated in figure 4.8.
As a result of the previous analysis it appears that for a stochastic process generated according to the structure in Figure 4.8 (left), the complex spectral density function is determined by

$$\Phi_{\mathbf{y}}(z) = H_m(z)H_m(1/z)\sigma_e^2.$$

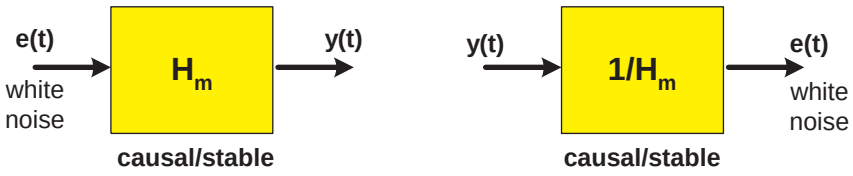

Figure 4.8: Innovations representation of WSS stochastic process $\mathbf{y}(t)$, with signal model (left) and inverse filter (right).

The spectral density functions considered in this section are not allowed to have poles and/or zeros exactly on the unit circle. When there are poles on the unit circle, $\Phi_{\mathbf{y}}(e^{i\omega})$ becomes unbounded and the Paley-Wiener condition mentioned in the footnote with theorem 4.4 can not be satisfied. Note that a pole on the unit circle implies the presence of a periodic component in $\mathbf{y}(t)$, and therefore a component that can be exactly (deterministically) described. Zeros that occur exactly on the unit circle require special care, as they can lead to inverse filters $1/H_m$ that are unstable.

4.3 Autoregressive moving average processes

4.3.1 Introduction

This section is based on Priestley (1981) and Shanmugan and Breipohl (1988). More detailed information can be found in these references.

A special class of random processes is formed by the so-called autoregressive (AR), moving-average (MA) and autoregressive moving average process (ARMA). These models have each been found to provide useful descriptions for a large number of "real life" processes. A discrete-time autoregressive process has been introduced already in section 3.4.7. It is represented by a difference equation of the form:

$$\mathbf{x}(t) + a_1\mathbf{x}(t-1) + a_2\mathbf{x}(t-2) + \cdots + a_n\mathbf{x}(t-n) = \mathbf{e}(t) \tag{4.47}$$

where $\mathbf{x}(t)$ is a real random sequence, $a_i, i = 1, \ldots, n$ are real-valued constant coefficients, and $\mathbf{e}$ is a (stationary) white noise process with variance $\sigma^2_{\mathbf{e}}$. If the difference equation is of the order n (i.e., $a_n \neq 0$), then the process is called an nth order autoregressive process (denoted by AR(n)).
An mth order moving average process (denoted by MA(m)) is described by:

$$\mathbf{x}(t) = \mathbf{e}(t) + b_1\mathbf{e}(t-1) + \ldots + b_m\mathbf{e}(t-m) \tag{4.48}$$

where, again, $\mathbf{e}$ is a white noise process, and $b_i, i = 1, \ldots, m, b_m \neq 0$ are real valued constants. In the autoregressive case the value of $\mathbf{e}(t)$ influences all future values $\mathbf{x}(t), \mathbf{x}(t+1), \ldots$. In the moving average case $\mathbf{x}(t)$ is expressed directly as a linear combination of the present and past values of the white noise process $\mathbf{e}$ but of finite extent, so that $\mathbf{e}(t)$ influences only q future values of $\mathbf{x}(t)$, namely $\mathbf{x}(t+1), \ldots, \mathbf{x}(t+m)$. This feature accounts for the fact that whereas the autocorrelation function of an autoregressive function "dies out gradually", the autocorrelation function of an MA(l) process, as we will show, cuts off after the point m, i.e., $R_{\mathbf{x}}(\tau) = 0, |\tau| > m$.

The mixed autoregressive/moving average model of order (n, m) is obtained by combining the above equations, leading to a model of the form

$$\mathbf{x}(t) + a_1\mathbf{x}(t-1) + \cdots + a_n\mathbf{x}(t-n) = \mathbf{e}(t) + b_1\mathbf{e}(t-1) + \ldots + b_m\mathbf{e}(t-m). \tag{4.49}$$

The mixed model includes of course, both the autoregressive and moving average models as special cases, these being obtained, respectively, either by setting $b_1 = b_2 = \ldots b_m = 0$, or by setting $a_1 = a_2 = \ldots a_n = 0$. Note that the ARMA model may be written more concisely in the operational form:

$$A(q^{-1})\mathbf{x}(t) = B(q^{-1})\mathbf{e}(t) \tag{4.50}$$

with q^{-1} the *backward shift operator*:

$$q^{-1}\mathbf{x}(t) = \mathbf{x}(t-1) \tag{4.51}$$

and

$$A(z^{-1}) = 1+a_1z^{-1}+a_2z^{-2}+\cdots+a_nz^{-n}, \quad B(z^{-1}) = 1+b_1z^{-1}+b_2z^{-2}+\cdots+b_mz^{-m} \tag{4.52}$$

It can be shown that an ARMA process corresponds to the output obtained by passing white noise through a linear time-invariant finite-dimensional causal dynamical system, as schematically depicted in figure 4.9, with a rational transfer function. Indeed, the ARMA

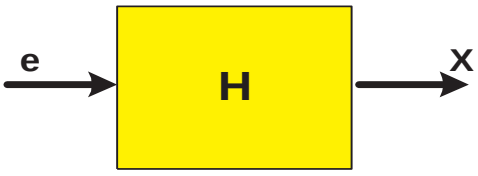

Figure 4.9: Filtering a stochastic process $\mathbf{e}$ by a linear dynamical system H.

process $\mathbf{x}$ can be represented by the convolution expression

$$\mathbf{x}(t) = \sum_{k=0}^{\infty} h(k)\mathbf{e}(t-k), \qquad t = 0, 1, 2, \cdots \tag{4.53}$$

where the transfer function, which is defined by

$$H(z) = \sum_{k=0}^{\infty} h(k)z^{-k}, \tag{4.54}$$

can be written as the rational function

$$H(z) = \frac{B(z^{-1})}{A(z^{-1})}. \tag{4.55}$$

This transfer function reduces to

$$H(z) = \frac{1}{A(z^{-1})} \tag{4.56}$$

and

$$H(z) = B(z^{-1}) \tag{4.57}$$

for AR processes and MA processes, respectively. The system will be stable if and only if all poles of the transfer function lie inside the unit circle in the complex plane. It follows from the analysis in the previous sections, that the following relationships for the complex (cross-)power spectral density functions of the ARMA process hold:

$$\begin{aligned} \Phi_{\mathbf{xe}}(z) &= H(z) \cdot \Phi_{\mathbf{e}}(z) = H(z) \cdot \sigma_{\mathbf{e}}^2 && (4.58)\\ \Phi_{\mathbf{ex}}(z) &= H(z^{-1}) \cdot \Phi_{\mathbf{e}}(z) = H(z^{-1}) \cdot \sigma_{\mathbf{e}}^2 && (4.59)\\ \Phi_{\mathbf{x}}(z) &= H(z)H(z^{-1}) \cdot \Phi_{\mathbf{e}}(z) = H(z)H(z^{-1}) \cdot \sigma_{\mathbf{e}}^2 && (4.60) \end{aligned}$$

and the power spectral density function of an ARMA process is given by:

$$\Phi_{\mathbf{x}}(\omega) = |H(e^{i\omega})|^2 \cdot \Phi_{\mathbf{e}}(\omega) = |H(e^{i\omega})|^2 \cdot \sigma_{\mathbf{e}}^2 \tag{4.61}$$

which is a rational function of $e^{i\omega}$.

4.3.2 First order autoregressive process

Let us now consider the AR(1) process. Recall that this process is described by the following difference equation:

$$\mathbf{x}(t) + a_1\mathbf{x}(t-1) = \mathbf{e}(t) \tag{4.62}$$

where $\mathbf{e}$ is a white noise process, and a_1 is a real-valued constant. It follows from (4.56) and (4.52) that the transfer function of the system described by model (4.62) is given by

$$H(z) = \frac{1}{1 + a_1 z^{-1}}. \tag{4.63}$$

The mean and the autocorrelation function of an AR(1) process have already been derived in section 3.4.7. There it was found that $\mathbb{E}[\mathbf{x}(t)] = 0$ and, for $|a_1| < 1$ and t sufficiently large, the autocorrelation function $R_{\mathbf{x}}(t+\tau, t)$ is a function of τ only, being given by

$$R_{\mathbf{x}}(\tau) = a_1^{|\tau|}\sigma_{\mathbf{x}}^2. \tag{4.64}$$

Hence, we may say that $\mathbf{x}$ is asymptotically WSS if $|a_1| < 1$, i.e. if the pole of the transfer function (4.63) lies inside the unit circle. The power spectral density function of an AR(1) process follows from (4.61) and (4.63), yielding:

$$\Phi_{\mathbf{x}}(\omega) = \left|\frac{1}{1 + a_1 e^{-i\omega}}\right|^2 \sigma_{\mathbf{e}}^2 = \frac{\sigma_{\mathbf{e}}^2}{1 + a_1^2 + 2a_1\cos(\omega)}. \tag{4.65}$$

Autocorrelation functions and corresponding power spectral density functions of three AR(1) processes are sketched in Figure 4.10.
Note that the autocorrelation functions of AR(1) processes show an exponential decay of the function for increasing values of τ.

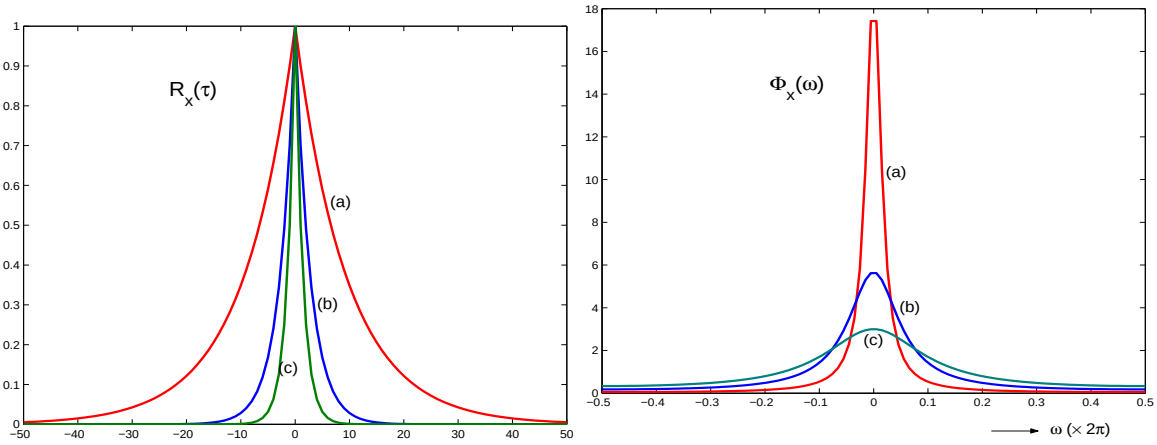

Figure 4.10: Autocorrelation functions (left) and Power spectral density functions (right) of the AR(1) process $\mathbf{x}$ described by (4.62) for $a = 0.9$ (a)/red, $a = 0.7$ (b)/blue, and $a = 0.5$ (c)/green.

4.3.3 Autoregressive process of general order

Recall that $\mathbf{x}$ is an autoregressive process of order n (denoted by AR(n)) if it satisfies the difference equation (4.47), while the transfer function of the system described by this difference equation is given by equation (4.56), that is,

$$H(z) = \frac{1}{A(z^{-1})} = \frac{1}{1 + a_1 z^{-1} + a_2 z^{-2} + \ldots + a_n z^{-n}} \tag{4.66}$$

Reasoning as in the first order case, it can be shown that an AR(n) process is asymptotically stationary (in wide sense) if all the poles of (4.66) lie inside the unit circle (Priesley, 1981). If we assume stationarity, multiplying both sides of equation (4.47) by $\mathbf{x}(t-\tau)$ and taking the expectations yields:

$$R_{\mathbf{x}}(\tau) + \sum_{i=1}^{n} a_i \cdot R_{\mathbf{x}}(\tau - i) = 0, \quad \tau \geq 1. \tag{4.67}$$

which, for $\tau = 1, 2, \ldots, n$, is an nth order difference equation, which can be expressed in matrix form as

$$\begin{bmatrix} R_{\mathbf{x}}(1) \\ R_{\mathbf{x}}(2) \\ \vdots \\ R_{\mathbf{x}}(n) \end{bmatrix} = \begin{bmatrix} R_{\mathbf{x}}(0) & R_{\mathbf{x}}(1) & \cdots & R_{\mathbf{x}}(n-1) \\ R_{\mathbf{x}}(1) & R_{\mathbf{x}}(0) & \ddots & R_{\mathbf{x}}(n-2) \\ \vdots & \ddots & \ddots & \vdots \\ R_{\mathbf{x}}(n-1) & \cdots & R_{\mathbf{x}}(1) & R_{\mathbf{x}}(0) \end{bmatrix} \cdot \begin{bmatrix} -a_1 \\ -a_2 \\ \vdots \\ -a_n \end{bmatrix} \tag{4.68}$$

This matrix equation is called the Yule-Walker equation. Since the $n \times n$ matrix in (4.68) is invertible, this equation can be used to express the coefficients of the AR(k) process in terms of autocorrelation function values, which is of considerable importance in data analysis.

It follows from (4.52), (4.56) and (4.61) that the power spectral density function of an AR(n) process is given by

$$\Phi_{\mathbf{x}}(\omega) = \frac{\sigma_{\mathbf{e}}^2}{|1 + \sum_{k=1}^{n} a_k e^{-i\omega k}|^2}. \tag{4.69}$$

4.3.4 First order moving average process

A moving average process of first order (an MA(1) proces) is described by the following difference equation:

$$\mathbf{x}(t) = \mathbf{e}(t) + b_1 \mathbf{e}(t-1), \tag{4.70}$$

where b_1 is a constant and $\mathbf{e}$ is a white noise process. It follows from (4.52) and (4.57) that the transfer function of the system described by (4.70) is given by:

$$H(z) = 1 + b_1 z^{-1} \tag{4.71}$$

The expected value, the variance and the autocorrelation function of the MA(1) process are given by

$$\mu_{\mathbf{x}} := \mathbb{E}[\mathbf{x}(t)] = \mathbb{E}\left[\mathbf{e}(t) + b_1 \mathbf{e}(t-1)\right] = 0, \tag{4.72}$$

$$\sigma_{\mathbf{x}}^2 = \mathbb{E}[\mathbf{x}^2(t)] = (b_1^2 + 1)\sigma_{\mathbf{e}}^2, \tag{4.73}$$

and (for $|\tau| \geq 1$)

$$\begin{aligned} R_{\mathbf{x}}(\tau) &= \mathbb{E}\left[\mathbf{x}(t)\mathbf{x}(t-\tau)\right] \\ &= \mathbb{E}\left[\left(b_1 \mathbf{e}(t-1) + \mathbf{e}(t)\right)\left(b_1 \mathbf{e}(t-\tau-1) + \mathbf{e}(t-\tau)\right)\right] \\ &= \begin{cases} b_1 \sigma_{\mathbf{e}}^2, & |\tau| = 1, \\ 0, & |\tau| > 1, \end{cases} \end{aligned} \tag{4.74}$$

respectively. Since the first two moments of the MA(1) process are independent of time, the process is WSS. The power spectral density function of the MA(1) process is

$$\Phi_{\mathbf{x}}(\omega) = \sum_{\tau=-1}^{1} R_{\mathbf{x}}(\tau) e^{-i\omega\tau} = \sigma_{\mathbf{e}}^2 \left(b_1^2 + 2b_1 \cos(\omega) + 1\right). \tag{4.75}$$

Note that this power spectral density function can also be derived from (4.71) and (4.61), yielding the same result:

$$\begin{aligned} \Phi_{\mathbf{x}}(\omega) &= \left|1 + b_1 e^{-i\omega}\right|^2 \cdot \sigma_{\mathbf{e}}^2 \\ &= \sigma_{\mathbf{e}}^2 \left(1 + b_1 e^{-i\omega} + b_1 e^{i\omega} + b_1^2\right)\right) \\ &= \sigma_{\mathbf{e}}^2 \left(b_1^2 + 2b_1 \cos(\omega) + 1\right). \end{aligned} \tag{4.76}$$

Autocorrelation functions and corresponding power spectral density functions of two MA(1) processes are sketched in Figure 4.11.
Note that the autocorrelation function has finite support, i.e. it becomes exactly zero for $|\tau| > 1$.

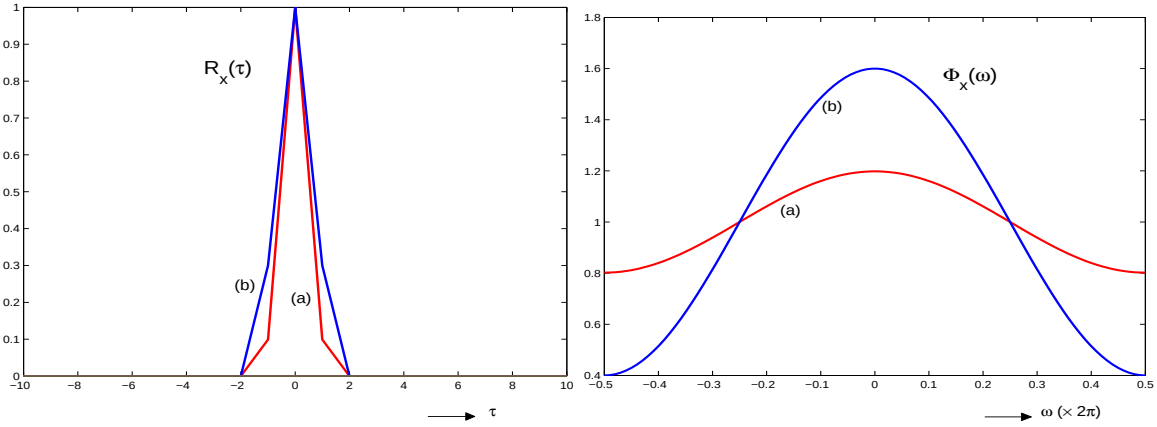

Figure 4.11: Autocorrelation functions (left) and Power spectral density functions (right) of the MA(1) process $\mathbf{x}$ described by (4.70) for $b_1 = 0.1$ (a)/red, and $b_1 = 1/3$ (b)/blue.

4.3.5 Moving average processes of general order

Recall that an MA(m) process may be expressed in the form:

$$\mathbf{x}(t) = \mathbf{e}(t) + b_1\mathbf{e}(t-1) + \ldots + b_m\mathbf{e}(t-m) \tag{4.77}$$

where, again, $\mathbf{e}$ is a white noise process, and $b_i, i = 1, \ldots, m, b_m \neq 0$ are real valued constants. It is easy to see that an MA process is always WSS, irrespective of the values of $b_1, b_2, \ldots, b_m$. Since $\mathbf{e}$ is a white noise process, we have $\mathbb{E}[\mathbf{x}(t)] = 0$, and

$$\sigma^2_{\mathbf{x}} = \sigma^2_{\mathbf{e}}(1 + b_1^2 + \cdots + b_m^2). \tag{4.78}$$

Furthermore, if we define $b_0 = 1$, it can be shown that the autocorrelation function is given by

$$\begin{aligned} R_{\mathbf{x}}(\tau) &= \mathbb{E}\left[\left(\sum_{i=0}^{m} b_i\mathbf{e}(t-i)\right)\left(\sum_{j=0}^{m} b_j\mathbf{e}(t-\tau-j)\right)\right] \\ &= \begin{cases} \sigma^2_{\mathbf{e}}\left(b_\tau b_0 + b_{\tau+1}b_1 + \ldots + b_m b_{m-\tau}\right), & 0 \leq \tau \leq m, \\ 0, & \tau > m, \end{cases} \\ R_{\mathbf{x}}(\tau) &= R_{\mathbf{x}}(-\tau), \quad \tau < 0 \end{aligned} \tag{4.79}$$

It follows from Equations (4.61), (4.57) and (4.52) that the power spectral density function of an MA(l) process is given by

$$\Phi_{\mathbf{x}}(\omega) = \left|B(e^{-i\omega})\right|^2 \sigma^2_{\mathbf{e}} = \left|1 + \sum_{k=1}^{m} b_k e^{-i\omega k}\right|^2 \cdot \sigma^2_{\mathbf{e}}. \tag{4.80}$$

4.3.6 Mixed autoregressive/moving average processes

Recall that an ARMA(n, m) process is describe by the following difference equation

$$\mathbf{x}(t) + a_1\mathbf{x}(t-1) + \cdots + a_n\mathbf{x}(t-n) = \mathbf{e}(t) + b_1\mathbf{e}(t-1) + \ldots + b_m\mathbf{e}(t-m). \tag{4.81}$$

The transfer function of the system represented by (4.81) is described by (4.55) where $A(z^{-1})$ and $B(z^{-1})$ are defined by (4.52). For stationarity, we require that all poles of the transfer function (4.55) lie inside the unit circle. The autocorrelation function of the ARMA(n, m) process can be derived as follows:

$$\begin{aligned} R_{\mathbf{x}}(\tau) &= \mathbb{E}\left[\mathbf{x}(t-\tau)\mathbf{x}(t)\right] \\ &= \mathbb{E}\left[\mathbf{x}(t-\tau)\cdot\left\{-\sum_{k=1}^{n} a_k\mathbf{x}(t-k)+\sum_{j=1}^{m} b_j\mathbf{e}(t-j)+\mathbf{e}(t)\right\}\right] \\ &= -\sum_{k=1}^{n} a_k R_{\mathbf{x}}(\tau-k)+\mathbb{E}[\mathbf{x}(t-\tau)\mathbf{e}(t)] \\ &\quad +\sum_{j=1}^{m} b_j\mathbb{E}[\mathbf{x}(t-\tau)\mathbf{e}(t-j)]. \end{aligned} \tag{4.82}$$

Because

$$\mathbb{E}[\mathbf{x}(t-\tau)\mathbf{e}(t)]=0, \quad \tau\geq 1 \tag{4.83}$$

the preceding equation reduces to

$$R_{\mathbf{x}}(\tau)=-\sum_{k=1}^{n} a_k R_{\mathbf{x}}(\tau-k), \tag{4.84}$$

for $\tau \geq m+1$. Thus for an ARMA(n, m) model $R_{\mathbf{x}}(0), R_{\mathbf{x}}(1), \ldots, R_{\mathbf{x}}(m)$ will depend upon both the autoregressive and the moving average parameters. The remainder of the autocorrelation function, that is $R_{\mathbf{x}}(\tau), \tau > m$ is determined by the nth order difference equation (4.84). Finally, it follows from Equation (4.61) that the power spectral density function of an ARMA(n, m) process is given by

$$\Phi_{\mathbf{x}}(\omega) = \left|\frac{1+\sum_{j=1}^{m} b_j e^{-i\omega j}}{1+\sum_{k=1}^{n} a_k e^{-i\omega k}}\right|^2 \cdot \sigma_{\mathbf{e}}^2. \tag{4.85}$$

4.4 Summary

Random signals resulting from stochastic processes can be filtered by linear filters. It has been shown that relatively simple theory exists for quantifying the effect of these linear filters. Auto- and cross-correlation functions are obtained by convolution with the filters' pulse response, and spectral densities are obtained by multiplication with the filters' frequency response. A general class of WSS stochastic processes (all rational spectra) can be modelled as the output of a linear filter that is driven by white noise at the input. This theory will be used in subsequent chapters when designing optimal filters for signal estimation and reconstruction.

Appendix 4A

Proof of the Paley-Wiener condition for spectral factorization

The Paley-Wiener condition:

$$\int_{-\pi}^{\pi} |ln\Phi_{\mathbf{y}}(e^{i\omega})| d\omega < \infty$$

implies that

$$ln\Phi_{\mathbf{y}}(z) = \sum_{k=-\infty}^{\infty} c_k z^{-k} \tag{4A.1}$$

is analytic and therefore converges in an annulus $\rho < |z| < 1/\rho$ that contains the unit circle. The sequence $\{c_k\}$ is called the cepstrum of the process $\mathbf{y}$. c_k may be viewed as Fourier coefficients of the periodic function $ln\Phi_{\mathbf{y}}(e^{i\omega})$, and therefore

$$c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} ln\Phi_{\mathbf{y}}(e^{i\omega}) e^{i\omega k} d\omega \tag{4A.2}$$

and

$$c_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} ln\Phi_{\mathbf{y}}(e^{i\omega}) d\omega. \tag{4A.3}$$

From (4A.1) the following decomposition can be constructed:

$$\Phi_{\mathbf{y}}(z) = exp\{ \sum_{k=-\infty}^{\infty} c_k z^{-k} \} \tag{4A.4}$$

$$= exp\{c_0\} \cdot exp\{\sum_{k=1}^{\infty} c_k z^{-k}\} \cdot exp\{ \sum_{k=-\infty}^{-1} c_k z^{-k}\}. \tag{4A.5}$$

The function

$$H(z) = exp\{\sum_{k=1}^{\infty} c_k z^{-k}\}$$

is the z-transform of a stable and causal system, that can be written in a power series of the form:

$$H(z) = 1 + h(1)z^{-1} + h(2)z^{-2} + \cdots$$

which is analytic for $|z| > \rho$ and therefore is causal and stable.
Performing a similar analysis for $1/\Phi_{\mathbf{y}}(z)$, it can be shown that the factor $H(z)$ will also have a causal and stable inverse, and therefore is minimum-phase.
Since $\ln \Phi_{\mathbf{y}}(e^{i\omega})$ is positive and real-valued, it follows that c_k is real-valued and $c_{-k} = c_k$, showing that

$$exp\{ \sum_{k=-\infty}^{-1} c_k z^{-k}\} = exp\{\sum_{k=1}^{\infty} c_k z^{k}\} = H(1/z).$$

As a result

$$\Phi_{\mathbf{y}}(z) = e^{c_0} \cdot H(z)H(1/z)$$

with $H(z)$ minimum-phase and $e^{c_0} > 0$.

Chapter 5

Estimation

In this chapter the basic concepts are introduced for estimating physical quantities on the basis of (random) measurement data. It is discussed how the quality of estimators can be assessed in terms of bias and variance. Additionally basic estimation principles as least-squares, linear regression, weighted and total least squares, and maximum likelihood are presented.

5.1 Introduction

In estimation theory it is the objective to determine/estimate an unknown physical quantity or variable, on the basis of data that is available from measurements, that are generally subject to stochastic uncertainties. Because of the stochastic uncertainties, the measurement data is actually described by a set (or sequence) of random variables.
An estimator (or sometimes called *statistic*)[1] is any (deterministic) function of a (set of) random variable(s); the estimator itself does not contain any unknown parameters.
Let, e.g., $\mathbf{x}_1, \cdots, \mathbf{x}_N$ be a random sample from the probability density function $f_{\mathbf{x}}(x)$, then

$$\mathbf{y} = g(\mathbf{x}_1, \cdots, \mathbf{x}_N)$$

with any function g is an estimator. In other words, an estimator is simply a function of random variables, and attempts to reproduce/estimate an unknown (physical) quantity on the basis of measured data.
In this way

$$\bar{\mathbf{x}} := \frac{1}{N} \sum_{i=1}^{N} \mathbf{x}_i$$

is an estimator of the mean value of $\mathbf{x}$, but also

$$\tilde{\mathbf{x}} := \frac{1}{2}[\max_i \mathbf{x}_i + \min_i \mathbf{x}_i]$$

can be an estimator of this quantity.

[1] A statistic is any deterministic function of observable random variables, which is itself an observable random variable, which does not contain any unknown parameters (Mood et al., 1974). When a statistic is used to estimate a function $\tau(\theta)$ of an unknown parameter θ, it is called an estimator of $\tau(\theta)$.

Estimators are not bound to have any direct relation with the quantity/variable that is estimated. Even $\min_i \mathbf{x}_i$ can be an estimator of the mean value of $\mathbf{x}$, although one may guess that the properties of this estimator will be rather poor, unless $\mathbf{x}_i$ is fixed for all values of i.
In this chapter several approaches will be discussed on how to design estimators for particular objects and with particular properties. The discussion is directed towards general estimators and their properties. Estimation of second order moments (correlation functions and spectral density function) of stochastic processes will be the subject of a later chapter.

5.2 Characteristics of estimators

5.2.1 Basic definitions

We will first discuss several notions that characterize the "quality" of estimators. As a notation we will consider an estimator $\hat{\boldsymbol{\theta}}_N$ of an underlying quantity θ_0, based on N measurements. Since we will generally assume that both $\hat{\boldsymbol{\theta}}_N$ and θ_0 are real-valued vectors in $\mathbb{R}^d$, the definitions will be given for real-valued parameters. However, they can easily be extended to include complex-valued parameters.

Bias
The estimator $\hat{\boldsymbol{\theta}}_N$ is called unbiased if

$$\boxed{\mathbb{E}[\hat{\boldsymbol{\theta}}_N] = \theta_0}$$

i.e., the estimator is delivering the "right" quantity "on average" and for a particular finite value of N. Additionally the estimator is called asymptotically unbiased if

$$\lim_{N\to\infty} \mathbb{E}[\hat{\boldsymbol{\theta}}_N] = \theta_0.$$

Estimators that are not unbiased are called biased, and the bias is given by

$$bias(\hat{\boldsymbol{\theta}}_N) = \mathbb{E}[\hat{\boldsymbol{\theta}}_N] - \theta_0.$$

Variance
The variance of $\hat{\boldsymbol{\theta}}_N$ is given by the mean deviation from its expected value, measured in a quadratic sense, i.e.,

$$\boxed{var(\hat{\boldsymbol{\theta}}_N) = \mathbb{E}\left[(\hat{\boldsymbol{\theta}}_N - \mathbb{E}\hat{\boldsymbol{\theta}}_N)(\hat{\boldsymbol{\theta}}_N - \mathbb{E}\hat{\boldsymbol{\theta}}_N)^T.\right]}$$

It determines the variation of outcomes of $\hat{\boldsymbol{\theta}}_N$ about its mean value. For vector-valued parameters the resulting matrix is referred to as the covariance matrix, also denoted as

$$cov(\hat{\boldsymbol{\theta}}_N)$$

having dimensions $d \times d$ when $\hat{\boldsymbol{\theta}}_N$ has dimension n. For scalar-valued parameters this covariance matrix reduces to the variance of the random variable $\hat{\boldsymbol{\theta}}_N$.

Mean squared error (MSE)
The mean squared error (MSE) is given by[2]

$$\boxed{\text{MSE}(\hat{\boldsymbol{\theta}}_N) := \mathbb{E}[(\hat{\boldsymbol{\theta}}_N - \theta_0)^2].}$$

Denote $\mathbb{E}[\hat{\boldsymbol{\theta}}_N] = m$, then

$$\begin{aligned} \text{MSE}(\hat{\boldsymbol{\theta}}_N) &= \mathbb{E}[(\hat{\boldsymbol{\theta}}_N - m + m - \theta_0)^2] \\ &= \mathbb{E}[(\hat{\boldsymbol{\theta}}_N - m)^2] + (m - \theta_0)^2 + 2\mathbb{E}[(\hat{\boldsymbol{\theta}}_N - m)(m - \theta_0)]. \end{aligned}$$

Because $\mathbb{E}[\hat{\boldsymbol{\theta}}_N] = m$ the last term on the right hand side is 0, and therefore

$$\text{MSE}(\hat{\boldsymbol{\theta}}_N) = var(\hat{\boldsymbol{\theta}}_N) + [bias(\hat{\boldsymbol{\theta}}_N)]^2. \tag{5.1}$$

This renders the measure $\text{MSE}(\hat{\boldsymbol{\theta}}_N)$ into a notion that includes aspects of both bias and variance. It describes the mean deviation of the estimator from its exact value, measured in a quadratic sense.

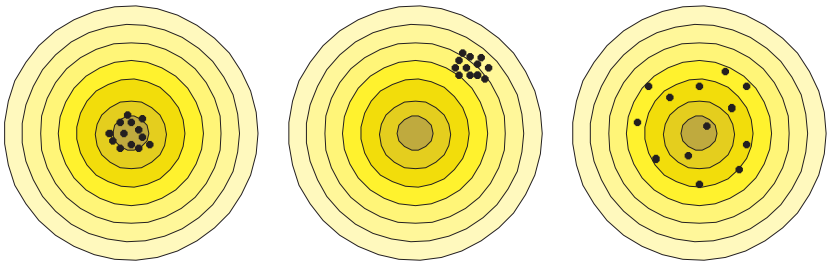

Figure 5.1: Hitting a target, where the bull's eye reflects θ_0; left: an unbiased estimator with small variance; middle: a biased estimator with small variance; right: an unbiased estimator with large variance.

For vector-valued parameters a natural extension exists, writing

$$\text{MSE}(\hat{\boldsymbol{\theta}}_N) := trace\mathbb{E}[(\hat{\boldsymbol{\theta}}_N - \theta_0)(\hat{\boldsymbol{\theta}}_N - \theta_0)^T]$$

then through a similar reasoning as above,

$$\begin{aligned} \text{MSE}(\hat{\boldsymbol{\theta}}_N) &= trace[\mathbb{E}(\hat{\boldsymbol{\theta}}_N - m)(\hat{\boldsymbol{\theta}}_N - m)^T + \mathbb{E}(m - \theta_0)(m - \theta_0)^T] \\ &= trace[cov(\hat{\boldsymbol{\theta}}_N) + bias(\hat{\boldsymbol{\theta}}_N) \cdot bias(\hat{\boldsymbol{\theta}}_N)^T] \qquad (5.2) \\ &= \sum_i var(\hat{\boldsymbol{\theta}}_N^{(i)}) + \sum_i bias(\hat{\boldsymbol{\theta}}_N^{(i)})^2 \qquad (5.3) \end{aligned}$$

where $(\cdot)^{(i)}$ refers to the i-th component of a vector. As a result the MSE for a vector-valued parameter is equal to the sum of the MSE's of each of its components.

[2] We first consider the situation of scalar real-valued parameters.

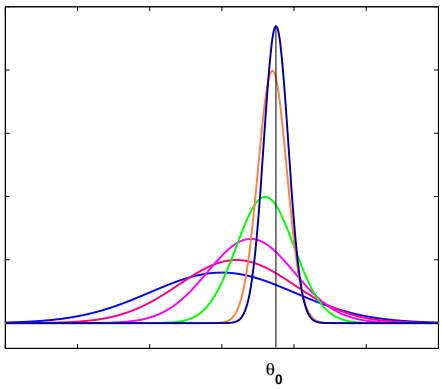

Figure 5.2: Probability density function $f_{\hat{\boldsymbol{\theta}}_N}(\theta)$ of a scalar-valued $\boldsymbol{\theta}$ for increasing N.

Distribution of estimator - probability density function
The most complete description of the properties of an estimator (which is a random variable), is of course its distribution function

$$F_{\hat{\boldsymbol{\theta}}_N}(\theta) = P[\hat{\boldsymbol{\theta}}_N \leq \theta].$$

or probability density function $f_{\hat{\boldsymbol{\theta}}_N}(\theta)$. For the considered situation, this pdf is a function of the number of measurements N that is taken into account.

If $f_{\hat{\boldsymbol{\theta}}_N}(\theta)$ tends to a Gaussian or normal distribution for $N \to \infty$, then $\hat{\boldsymbol{\theta}}_N$ is being referred to as being asymptotically normal (Gaussian), also denoted as

$$\hat{\boldsymbol{\theta}}_N \in As\mathcal{N}(\theta^*, \Sigma_{\boldsymbol{\theta}})$$

with θ^* the mean of the asymptotic distribution, and $\Sigma_{\boldsymbol{\theta}}$ the asymptotic covariance matrix.

Consistency
An estimator $\hat{\boldsymbol{\theta}}_N$ is called (weakly) consistent if for every $\delta > 0$,

$$\lim_{N\to\infty} Pr[\|\hat{\boldsymbol{\theta}}_N - \theta_0\| > \delta] = 0$$

also denoted as $\text{plim}_{N\to\infty} \hat{\boldsymbol{\theta}}_N = \theta_0$.
Weak consistency is also denoted as convergence in probability. It means that for every θ with $\|\theta_0 - \theta\| > \delta$ with δ arbitrarily small, the probability density function $f_{\hat{\boldsymbol{\theta}}_N}(\theta)$ disappears for $N \to \infty$. particular when the probability density function converges to a dirac-pulse for $N \to \infty$. This situation is illustrated in figure 5.2. Note that a (weakly) consistent estimator is not necessarily unbiased for finite values of N, as is also illustrated in figure 5.2. However an estimator for which holds that

$$\left.\begin{array}{c} \mathbb{E}[\hat{\boldsymbol{\theta}}_N] \to \theta_0 \\ cov(\hat{\boldsymbol{\theta}}_N) \to 0 \end{array}\right\} \text{ for } N \to \infty \tag{5.4}$$

is weakly consistent. Here the expression $cov(\hat{\boldsymbol{\theta}}_N) \to 0$ should be interpreted as an element-wise convergence to 0.

Efficiency
$\hat{\boldsymbol{\theta}}_N$ is called an efficient estimator of θ_0 if

$$cov(\hat{\boldsymbol{\theta}}_N) \leq cov(\bar{\boldsymbol{\theta}}_N) \tag{5.5}$$

for all unbiased estimators $\bar{\boldsymbol{\theta}}_N$. This means that it concerns an estimator that has the smallest possible variability (variance), measured in terms of its covariance matrix, of all unbiased estimators. Note that the inequality in (5.5) is a matrix inequality, requiring that the matrix $cov(\bar{\boldsymbol{\theta}}_N) - cov(\hat{\boldsymbol{\theta}}_N)$ is positive semi-definite.
For scalar-valued estimators, the (relative) efficiency of an estimator $\hat{\boldsymbol{\theta}}_1$ with respect to another estimator $\hat{\boldsymbol{\theta}}_2$ is sometimes denoted by

$$\frac{var(\hat{\boldsymbol{\theta}}_2)}{var(\hat{\boldsymbol{\theta}}_1)}.$$

Example 5.1 (Sample-average estimator of the mean value of a stochastic process)
Consider a discrete-time WSS stochastic process $\mathbf{x}(t)$ with

$$\mathbb{E}[\mathbf{x}] = \mu_{\mathbf{x}} \tag{5.6}$$
$$\mathbb{E}[(\mathbf{x} - \mu_{\mathbf{x}})^2] = \sigma^2_{\mathbf{x}}. \tag{5.7}$$

Determine the statistical properties (bias, variance, consistency) of the sample-average [3]

$$\bar{\mathbf{x}}_N := \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}(i)$$

as an estimator for the mean value $\mu_{\mathbf{x}}$.

- $\mathbb{E}[\bar{\mathbf{x}}_N] = \frac{1}{N}\sum_{i=1}^{N} \mu_{\mathbf{x}} = \mu_{\mathbf{x}}$. As a result the estimator is unbiased for any value of N.
- $var(\bar{\mathbf{x}}_N)$ is given by

$$\begin{aligned}\mathbb{E}[(\frac{1}{N}\sum_{i=1}^{N} \mathbf{x}(i) - \mu_{\mathbf{x}})^2] &= \frac{1}{N^2}\mathbb{E}[(\sum_{i=1}^{N} \mathbf{x}(i) - N\mu_{\mathbf{x}})^2] \\ &= \frac{1}{N^2}\mathbb{E}\left[\{(\mathbf{x}(1) - \mu_{\mathbf{x}}) + (\mathbf{x}(2) - \mu_{\mathbf{x}}) + \cdots + (\mathbf{x}(N) - \mu_{\mathbf{x}})\}^2\right]\end{aligned} \tag{5.8}$$

  If $\mathbf{x}(i), \mathbf{x}(j)$ are uncorrelated random variables for $i \neq j$, as , e.g., is the case if $\mathbf{x}$ is a white noise stationary stochastic process, then the previous expression directly leads to

$$var(\bar{\mathbf{x}}_N) = \frac{1}{N}\sigma^2_{\mathbf{x}}. \tag{5.9}$$

  As a result, the variance decays with $1/N$ and therefore tends to 0 for $N \to \infty$. Since the bias of the estimator is 0 for all values of N, this shows that the sample-average estimator is a consistent estimator of $\mu_{\mathbf{x}}$.

[3] In the literature this is often referred to as the sample-mean. Here we choose to use "mean" for expectation over the ensemble, and "average" as averaging over sequence/time.

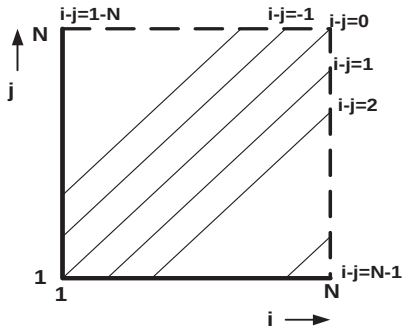

Figure 5.3: Variable substitution $i - j = k$ in (5.11).

- If $\mathbf{x}(i), \mathbf{x}(j)$ are not uncorrelated for $i \neq j$, then

$$var(\bar{\mathbf{x}}_N) = \mathbb{E}[\frac{1}{N^2}\sum_i\sum_j(\mathbf{x}(i)-\mu_{\mathbf{x}})(\mathbf{x}(j)-\mu_{\mathbf{x}})] \tag{5.10}$$

$$= \frac{1}{N^2}\sum_i\sum_j C_{\mathbf{x}}(|i-j|). \tag{5.11}$$

By drawing the i, j-grid over which the summation is taken, and by making the variable substitution $i - j = k$, as illustrated in Figure 5.3, it can be noted that for each value of k there are $N - |k|$ grid points to be taken into account. Therefore the previous expression becomes

$$var(\bar{\mathbf{x}}_N) = \frac{1}{N^2}\sum_{k=1-N}^{N-1}(N-|k|)C_{\mathbf{x}}(|k|) \tag{5.12}$$

$$= \frac{1}{N}\sum_{k=1-N}^{N-1}(1-\frac{|k|}{N})C_{\mathbf{x}}(|k|). \tag{5.13}$$

Note that the situation of $\mathbf{x}$ being a white noise process again results as a special case, i.e., if $C_{\mathbf{x}}(k) = 0$ for $k \neq 0$, then

$$var(\bar{\mathbf{x}}_N) = \frac{1}{N}C_{\mathbf{x}}(0) = \frac{\sigma_{\mathbf{x}}^2}{N}.$$

In order for the sample-average estimator to be a consistent estimator of the mean value it is sufficient that

$$\lim_{N\to\infty}\frac{1}{N}\sum_{k=1-N}^{N-1}(1-\frac{|k|}{N})C_{\mathbf{x}}(k) = 0. \tag{5.14}$$

This condition implies that $C_{\mathbf{x}}(k)$ should converge to 0 sufficiently fast for growing values of k. As a counter example consider the situation that $C_{\mathbf{x}}(k) = c$, a constant, then

$$var(\bar{\mathbf{x}}_N) = \frac{c}{N}\left[\sum_{k=1}^{N-1}(1-\frac{k}{N})+1+\sum_{k=1}^{N-1}(1+\frac{k}{N})\right] = \frac{c}{N}[2(N-1)+1] = \frac{c(2N-1)}{N}$$

which tends to $2c$ for $N \to \infty$, and therefore does not satisfy the condition (5.14). Apparently, properties of the covariance function of the process fully determine whether the mean value of the process can be estimated consistently on the basis of a single measurement sequence. Processes for which the sample-average estimator is a consistent estimator of the ensemble mean are referred to as *ergodic in the mean*, see also section 3.8.

Example 5.2 (Estimating the variance of a random variable) Let $\{\mathbf{x}_i\}_{i=1,\cdots N}$ be N independent observations of a random variable $\mathbf{x}$. We investigate the bias properties of the estimator

$$\widehat{\boldsymbol{\sigma}^2} := \frac{1}{N}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})^2 \tag{5.15}$$

with $\bar{\mathbf{x}} = (1/N)\sum_{i=1}^{N}\mathbf{x}_i$, as an estimator of the mean of $\mathbf{x}$:

$$\sigma_{\mathbf{x}}^2 = \mathbb{E}[(\mathbf{x} - \mathbb{E}\mathbf{x})^2].$$

The mean value of the estimator can be analyzed as follows:

$$\mathbb{E}[\widehat{\boldsymbol{\sigma}^2}] \;=\; \mathbb{E}[\frac{1}{N}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})^2] = \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}[(\mathbf{x}_i - \bar{\mathbf{x}})^2]. \tag{5.16}$$

Analyzing one term:

$$\begin{aligned}
\mathbb{E}[(\mathbf{x}_i - \bar{\mathbf{x}})^2] &= \mathbb{E}[(\mathbf{x}_i - \frac{1}{N}\sum_{j=1}^{N}\mathbf{x}_j)^2] = \mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}} + \mu_{\mathbf{x}} - \frac{1}{N}\sum_{j=1}^{N}\mathbf{x}_j)^2] \\
&= \mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}} - \frac{1}{N}\sum_{j=1}^{N}(\mathbf{x}_j - \mu_{\mathbf{x}}))^2] \\
&= \mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})^2] + \frac{1}{N^2}\sum_{j=1}^{N}\sigma_{\mathbf{x}}^2 - 2\mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})\frac{1}{N}\sum_{j=1}^{N}(\mathbf{x}_j - \mu_{\mathbf{x}})]
\end{aligned} \tag{5.17}$$

where the expression for the second term follows from equation (5.9). With respect to the last term on the right hand side of the last equation, we can write:

$$\mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})\sum_{j=1}^{N}(\mathbf{x}_j - \mu_{\mathbf{x}})] \;=\; \mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})^2] + \underbrace{\sum_{j=1,j\neq i}^{N}\mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})(\mathbf{x}_j - \mu_{\mathbf{x}})]}_{=0},$$

where the second term is zero because of the fact that $\mathbf{x}_i$ and $\mathbf{x}_j$ are independent random variable for $i \neq j$. As a result

$$\mathbb{E}[(\mathbf{x}_i - \mu_{\mathbf{x}})\sum_{j=1}^{N}(\mathbf{x}_j - \mu_{\mathbf{x}})] = \sigma_{\mathbf{x}}^2. \tag{5.18}$$

Substitution of this result in (5.17) delivers:

$$\begin{aligned}\mathbb{E}[(\mathbf{x}_i - \bar{\mathbf{x}})^2] &= \sigma_{\mathbf{x}}^2 + \frac{1}{N^2} \cdot N\sigma_{\mathbf{x}}^2 - \frac{2}{N}\sigma_{\mathbf{x}}^2 \\ &= \sigma_{\mathbf{x}}^2(1 - \frac{1}{N}) = \frac{\sigma_{\mathbf{x}}^2(N-1)}{N}. \end{aligned} \tag{5.19}$$

As a result:

$$\mathbb{E}[\widehat{\boldsymbol{\sigma}^2}] = \frac{N-1}{N}\sigma_{\mathbf{x}}^2.$$

Consequently the estimator is biased for finite values of N. The bias disappears when $N \to \infty$, as in that case $\frac{N}{N-1} \to 1$. However the analysis also shows that one can construct an unbiased estimator

$$\widetilde{\boldsymbol{\sigma}^2} := \frac{N}{N-1}\widehat{\boldsymbol{\sigma}^2} = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})^2.$$

The division by $N-1$ can be understood by realizing that one degree of freedom in the set of data is used to calculate $\bar{\mathbf{x}}$. As a result there remain $N-1$ degrees of freedom to be used in the variance estimation.

5.2.2 Diagonalizing the covariance matrix

The covariance matrix $\Sigma_{\boldsymbol{\theta}} = cov(\hat{\boldsymbol{\theta}})$ of a parameter estimator determines the statistical relation between the several components of the parameter vector $\hat{\boldsymbol{\theta}}$. If $\Sigma_{\boldsymbol{\theta}}$ is diagonal this implies that the several components of $\hat{\boldsymbol{\theta}}$ are uncorrelated. In the case that $\Sigma_{\boldsymbol{\theta}}$ has off-diagonal terms, this indicates a statistical relationship between the several variables.
In many situations it is convenient, and helpful for purposes of interpretation, to consider uncorrelated estimators, corresponding to a diagonal covariance matrix. This situation can simply be obtained by performing a linear operation of rotation and scaling to the parameter vector, as will be explained here.
Since $\Sigma_{\boldsymbol{\theta}}$ is positive semi-definite and symmetric all its eigenvalues will be real-valued and ≥ 0. Additionally the eigenvalue decomposition of $\Sigma_{\boldsymbol{\theta}}$ can be written as

$$\Sigma_{\boldsymbol{\theta}} = W\Lambda W^T$$

where W is a unitary matrix of eigenvectors w_i:

$$W = \begin{bmatrix} | & | & & | \\ w_1 & w_2 & \cdots & w_n \\ | & | & & | \end{bmatrix} \tag{5.20}$$

and Λ is a diagonal matrix with positive (≥ 0) eigenvalues λ_i:

$$\Lambda = \begin{bmatrix} \lambda_1 & & & 0 \\ & \lambda_2 & & \\ & & \ddots & \\ 0 & & & \lambda_n \end{bmatrix}. \tag{5.21}$$

Since W is unitary, its columns are orthonormal vectors, i.e., $W^T W = I_n$.
When defining the transformed parameter:

$$\hat{\boldsymbol{\eta}} := W^T \hat{\boldsymbol{\theta}}$$

it follows from simple substitution in the definitions that

$$\begin{aligned} \mathbb{E}[\boldsymbol{\eta}] &:= \mu_{\boldsymbol{\eta}} = W^T \cdot \mu_{\boldsymbol{\theta}} && (5.22) \\ cov(\hat{\boldsymbol{\eta}}) &:= \Sigma_{\boldsymbol{\eta}} = W^T \Sigma_{\boldsymbol{\theta}} W = \Lambda. && (5.23) \end{aligned}$$

The second equation can be verified by noting that

$$\Sigma_{\boldsymbol{\eta}} := \mathbb{E}[\boldsymbol{\eta} - \mu_{\boldsymbol{\eta}}][\boldsymbol{\eta} - \mu_{\boldsymbol{\eta}}]^T = \mathbb{E}[W^T(\hat{\boldsymbol{\theta}} - \mu_{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}} - \mu_{\boldsymbol{\theta}})^T W] = W^T \Sigma_{\boldsymbol{\theta}} W.$$

Consequently, the estimator $\hat{\boldsymbol{\eta}}$ has uncorrelated components, regardless of the probability density function of the "original" estimator $\hat{\boldsymbol{\theta}}$. If $\hat{\boldsymbol{\theta}}$ has a Gaussian probability density function, then it follows from section A.3.2 that the transformed random variable $\hat{\boldsymbol{\eta}}$ also is Gaussian distributed. In this case the components of $\hat{\boldsymbol{\eta}}$ are even statistically independent. The results for the covariance matrix can be used to plot the contour lines of the density function, defined by the relation:

$$f_{\boldsymbol{\theta}}(\theta) = constant.$$

In the case of a Gaussian distribution $\mathcal{N}(\theta_0, \Sigma_{\boldsymbol{\theta}})$, these contour lines are defined by the relation

$$(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma_{\boldsymbol{\theta}}^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) = c$$

with $c \in \mathbb{R}$, and $c \geq 0$.
Applying now the transformation from $\hat{\boldsymbol{\theta}}$ to $\hat{\boldsymbol{\eta}}$ then shows that

$$(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma_{\boldsymbol{\theta}}^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) = (\hat{\boldsymbol{\eta}} - \mu_{\boldsymbol{\eta}})^T \Lambda^{-1} (\hat{\boldsymbol{\eta}} - \mu_{\boldsymbol{\eta}}) = c \tag{5.24}$$

or equivalently

$$\sum_{i=1}^{n} \frac{|\hat{\boldsymbol{\eta}}^{(i)} - \mu_{\boldsymbol{\eta}}^{(i)}|^2}{\lambda_i} = c \tag{5.25}$$

with $(\cdot)^{(i)}$ denoting the i-th component of the vector considered. Equation (5.25) is the characterization of an ellipsoid. An ellipsoid in an (x, y)-plane is characterized by an equation of the type

$$\frac{x^2}{a} + \frac{y^2}{b} = c$$

or equivalently

$$\left(\frac{x}{\sqrt{a}}\right)^2 + \left(\frac{y}{\sqrt{b}}\right)^2 = c$$

determining an ellipsoid with the x and y axis as principal axes, and axes crossings $x = \sqrt{ac}$ and $y = \sqrt{bc}$.

In the considered situation (5.25) it concerns an ellipsoid in the orthogonal basis spanned by the components of $\hat{\boldsymbol{\eta}}$, with the center point $\mu_{\boldsymbol{\eta}}$. The orthogonal basis is determined by the relation

$$\hat{\boldsymbol{\eta}} = W^T \hat{\boldsymbol{\theta}} = \begin{pmatrix} w_1^T \\ w_2^T \\ \vdots \\ w_n^T \end{pmatrix} \hat{\boldsymbol{\theta}}_N$$

The first basis vector is determined by

$$\begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} = \begin{pmatrix} w_1^T \\ w_2^T \\ \vdots \\ w_n^T \end{pmatrix} \hat{\boldsymbol{\theta}}_N$$

leading to $\hat{\boldsymbol{\theta}} = w_1$, etcetera. Therefore the principal axes of the ellipsoid are aligned with the orthogonal eigenvectors $w_1, w_2, \cdots$ of $\Sigma_{\boldsymbol{\theta}}$. The principal axes of the ellipsoid are determined in size by $2\sqrt{c\lambda_i}$, $i = 1, \cdots n$. This is illustrated in figure 5.4 for a 2-dimensional example.

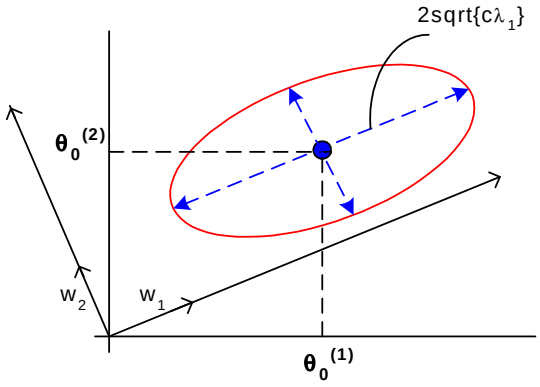

Figure 5.4: Ellipsoid indicating levels of equal probability density function for a normally distributed estimator $\hat{\boldsymbol{\theta}}$ with covariance matrix $\Sigma_{\boldsymbol{\theta}}$ having eigenvalues λ_1, λ_2 and eigenvectors w_1, w_2.

5.3 Linear regression

5.3.1 Introduction

One of the most simple examples of estimation problems is the problem of estimating a linear relationship between two different (random) variables, on the basis of multiple observations of the two variables. Consider for instance the situation sketched in figure 5.5(left) where the (blue) dots reflect measurement pairs (u_i, y_i). The question of finding a linear relationship between y and u can then be rephrased by finding the "best" straight line through the cloud of measurement points.
In order to solve this problem one generally has to make some assumptions on the source of the randomness in the data. We will first treat the standard case, where it is assumed that

one variable (u) is measured noise free, and the other variable (y) is noise disturbed. Later we will comment on this and discuss generalizations of this paradigm. In the considered situation the relation between the measurements, $\{u_i, y_i\}_{i=1,\cdots n}$, is hypothesized by the model:

$$y_i = b_0 + b_1 u_i + e_i \tag{5.26}$$

where b_0, b_1 are unknown coefficients (parameters) that are to be estimated, and e_i is an error term that accounts for the fact that the measured points do not lie exactly on a straight line. The term e_i can be considered a realization of a random variable $\mathbf{e}$ with a particular probability density function. More attention will be paid to this when analyzing the properties of the linear least squares estimator and the weighted linear least squares estimator.

5.3.2 Linear least-squares estimation

The linear least-squares estimate, constructs a solution to the sketched problem, by looking for an estimate $\hat{b}_0$ and $\hat{b}_1$ such that

$$\sum_{i=1}^{n} e_i^2$$

is minimal. This implies that the "errors" between the straight line and measurement points are minimal (in squared sense) if one considers the "errors" to lie in the y-direction of the plot, as sketched in figure 5.5(right).

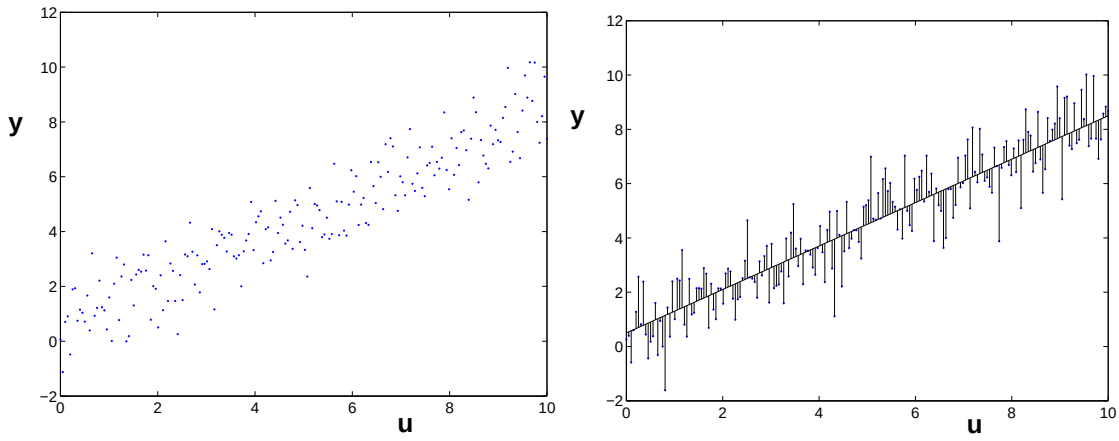

Figure 5.5: Observed data points $\{u_i, y_i\}$ (left) and measure of fit to a straight line, by considering the errors in the y variables (right).

Starting from the model equation

$$y_i = b_0 + b_1 u_i + e_i \tag{5.27}$$

we write

$$y_i = \phi_i^T \theta + e_i \tag{5.28}$$

with

$$\phi_i = \begin{bmatrix} 1 \\ u_i \end{bmatrix} \quad \text{and } \theta = \begin{bmatrix} b_0 \\ b_1 \end{bmatrix}.$$

In this formulation the vector ϕ_i is called the *regressor* or regression variable. The components of ϕ_i are referred to as the *independent variables* in the regression problem, while the y_i are denoted the *dependent variables*.
The cost function V that has to be minimized is written as

$$V := \sum_{i=1}^{n} e_i^2 = \sum_{i=1}^{n} (y_i - \phi_i^T \theta)^2$$

When the data is given, the function V becomes a function of θ, $V(\theta)$. Besides the fact that $V(\theta)$ is quadratic in e_i, it is also quadratic in θ. This implies that it is a convex function, as illustrated in figure 5.6 for a scalar-valued θ, having a unique global minimum that can be obtained by setting the derivative of $V(\theta)$ to zero.

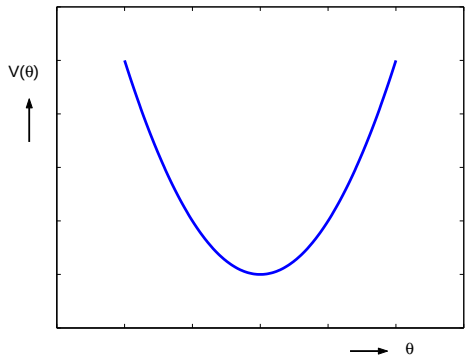

Figure 5.6: V is a quadratic function θ, which can be visualized for the situation that θ is scalar-valued.

Since[4]

$$\frac{\partial V(\theta)}{\partial \theta} = \begin{pmatrix} \frac{\partial V}{\partial b_0} \\ \frac{\partial V}{\partial b_1} \end{pmatrix} = \begin{pmatrix} -2\sum_i (y_i - \phi_i^T \theta) \\ -2\sum_i u_i (y_i - \phi_i^T \theta) \end{pmatrix}$$

it follows that

$$\left.\frac{\partial V}{\partial \theta}\right|_{\theta=\hat{\theta}} = 0 \rightarrow \sum_i \phi_i (y_i - \phi_i^T \hat{\theta}) = 0.$$

The resulting equations

$$\sum_i \phi_i (y_i - \phi_i^T \hat{\theta}) = 0 \tag{5.29}$$

are called the normal equations. Note that ϕ_i and θ are both 2-dimensional vectors, leading to a set of 2 (normal) equations with 2 unknowns. They deliver the following analytical solution for $\hat{\theta}$:

$$\left[\sum_i \phi_i \phi_i^T\right] \hat{\theta} = \sum_i \phi_i y_i \tag{5.30}$$

or equivalently:

$$\boxed{\hat{\theta} = \left[\sum_i \phi_i \phi_i^T\right]^{-1} \sum_i \phi_i y_i} \tag{5.31}$$

[4] As a notational convention, the derivative of a scalar with respect to a column vector is again a column vector.

provided that the inverse of the corresponding 2×2-matrix exists.
Note that the right hand side of the expression is only dependent on measurement data; once measurement data is available, the solution to the least-squares problem is simply obtained.
Existence of the matrix inverse in (5.31) is directly coupled to the question whether a sufficiently informative experiment has been done in order to uniquely determine the LS-solution. Consider , e.g., the situation that in the experiment all measurements u_i are the same, so

$$u_i = c \quad \text{for all } i.$$

Then $\phi_i = (1\ c)^T$ and

$$\sum_i \phi_i \phi_i^T = \sum_i \begin{bmatrix} 1 \\ c \end{bmatrix} [1 \quad c] = \begin{pmatrix} n & cn \\ cn & nc^2 \end{pmatrix}$$

As the second column of this matrix is obtained by scalar multiplication of the first column by c, the matrix is singular and its inverse will not exist. Therefore the LS-solution will not be unique; there will exists many solutions of the equation (5.30). The non-uniqueness of the solution can also simply be understood by considering the problem in the scope of figure 5.5; if the cloud of points is concentrated around one value of u_i, then there does not exists a unique "best" straight line[5] that relates u to y.

Least-squares solution in matrix form

For convenience, often use will be made of a more extensive matrix notation, replacing the summations of i.
The normal equations

$$\sum_i \phi_i (y_i - \phi_i^T \hat{\theta}) = 0 \tag{5.32}$$

can actually be rewritten as:

$$[\phi_1 \cdots \phi_n] \begin{pmatrix} y_1 - \phi_i^T \hat{\theta} \\ \vdots \\ y_n - \phi_n^T \hat{\theta} \end{pmatrix} = 0$$

which can be rewritten in short as

$$X^T (Y - X\hat{\theta}) = 0 \tag{5.33}$$

with

$$X = \begin{bmatrix} \phi_1^T \\ \vdots \\ \phi_n^T \end{bmatrix}; \quad Y = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

Note that X is a $n \times 2$ matrix and Y an $n \times 1$ vector. The solution is given by

$$\boxed{\hat{\theta} = (X^T X)^{-1} X^T Y}$$

where again the assumption has to be made that the 2×2 matrix inverse indeed exists.

[5] Note that the straight line $u = c$ does not relate u to y and therefore is not a solution to the problem.

5.3.3 Linear regression as a statistical estimation problem

Through imposing relation (5.27), the least-squares problem considered here has become a fully deterministic problem of drawing a straight line through a cloud of points, by minimizing the deviation to the line in the y-direction.

Considering the problem as a statistical estimation problem, we have to take care of the random character of the measurements y_i. If in (5.26) the variables e_i are random, then we have to write:

$$\mathbf{y}_i = b_0 + b_1 u_i + \mathbf{e}_i \tag{5.34}$$

showing that $\mathbf{y}_i$ actually are random variables as well.
The corresponding estimator of θ_0 then becomes

$$\hat{\boldsymbol{\theta}} = (X^T X)^{-1} X^T \mathbf{Y},$$

with $\mathbf{Y}$ now a vector of random variables:

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_n \end{bmatrix}$$

and therefore $\hat{\boldsymbol{\theta}}$ also becomes a random variable. As a result we can study the statistical properties of the estimator $\hat{\boldsymbol{\theta}}$.

Bias of a linear least squares estimator

Assume that the measured data u_i, y_i is generated by an equation of the form

$$\mathbf{Y} = X\theta_0 + \mathbf{E} \qquad \mathbf{E} = \begin{bmatrix} \mathbf{e}_1 \\ \vdots \\ \mathbf{e}_n \end{bmatrix}.$$

This implies that there exists a noise free output vector $Y = X\theta_0$ that is observed through the random variable $\mathbf{Y} = Y + \mathbf{E}$ where $\mathbf{E}$ is a vector of random variables.
Substituting $\mathbf{Y}$ into the expression for $\hat{\boldsymbol{\theta}}$, it follows that

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= (X^T X)^{-1} X^T (X\theta_0 + \mathbf{E}) \\ &= \theta_0 + (X^T X)^{-1} X^T \mathbf{E}. \end{aligned} \tag{5.35}$$

The estimator is unbiased if $\mathbb{E}[\hat{\boldsymbol{\theta}}] = \theta_0$. If X is deterministic, as considered here, this condition is simply satisfied if $\mathbb{E}[\mathbf{E}] = 0$, or equivalently if $\mathbb{E}[\mathbf{e}_i] = 0$ for all i.

Conclusion:

> The linear least squares estimator is unbiased if the noise terms on the output variables are zero-mean random variables.

Variance of a linear least squares estimator

For an unbiased estimator:

$$cov(\hat{\boldsymbol{\theta}}) = \mathbb{E}\left[(\hat{\boldsymbol{\theta}} - \theta_0)(\hat{\boldsymbol{\theta}} - \theta_0)^T\right] = \mathbb{E}\left[(X^T X)^{-1} X^T \mathbf{E}\mathbf{E}^T X (X^T X)^{-1}\right]$$

If $\mathbf{e}_i$ is a sequence of zero-mean uncorrelated random variables with equal variance σ^2, then

$$\mathbb{E}[\mathbf{E}\mathbf{E}^T] = \sigma^2 \cdot I$$

and when the components of X are deterministic (no random variables), this will imply that

$$\begin{aligned} cov(\hat{\boldsymbol{\theta}}) &= (X^T X)^{-1} X^T \sigma^2 I X (X^T X)^{-1} && (5.36) \\ &= \sigma^2 (X^T X)^{-1}. && (5.37) \end{aligned}$$

One of the important and appealing observations here is that the variance of $\hat{\boldsymbol{\theta}}$ increases with increasing values of σ^2: the higher the noise level, the larger the parameter variance.

A graphical interpretation of the covariance matrix is obtained when considering the situation when $\hat{\boldsymbol{\theta}}$ has a Gaussian distribution with covariance matrix

$$\Sigma = \sigma^2 (X^T X)^{-1}.$$

In this case

$$(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) \in \chi^2_d$$

which means that the left hand side expression follows a χ^2 distribution with d degrees of freedom (with d the dimension of θ_0). [6] As a result the probability that

$$(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) \geq \alpha$$

is specified by the χ^2_d distribution. The sets determined by this latter expression determine ellipsoids in $\mathbb{R}^2$, as visualized in figure 5.7 (for the case $d = 2$) and elaborated in section 5.2.2. Note that when giving confidence intervals for the two parameters on the basis of a particularly chosen level of probability α, the covariance result (corresponding to the ellipsoidal area in figure 5.7) is less conservative than when bounding the two parameters separately (corresponding to the rectangular area in the figure). We will come back to the subject of confidence intervals in subsection 5.3.3.

Variance for incorrect models

The expression for the variance of $\hat{\boldsymbol{\theta}}$ remains the same if the chosen model does not match the underlying data generating equations. Consider , e.g., the situation where data originates from the relation

$$\mathbf{Y} = X\theta_0 + \mathbf{E}_0 = X_r \theta_0^{(1)} + X_e \theta_0^{(2)} + \mathbf{E}_0$$

while a model is used with only a restricted number of regression variables:

$$\mathbf{Y} = X_r \theta + \mathbf{E}.$$

[6] A χ^2_d distribution is typically obtained when taking the sum of d squared terms of independent normally distributed random variables with mean 0 and variance 1.

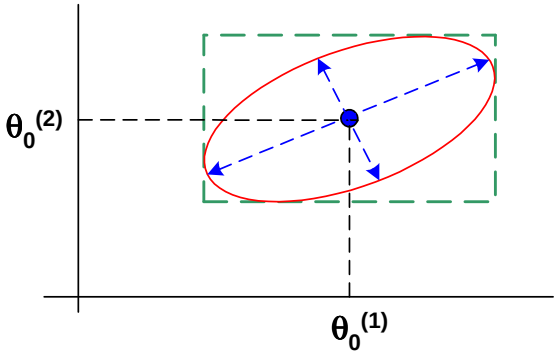

Figure 5.7: Ellipsoid indicating levels of equal probability density function for a normally distributed estimator $\hat{\boldsymbol{\theta}}$ with covariance matrix Σ. The principal axes of the ellipsoid are determined by the eigenvectors and eigenvalues of Σ.

In this situation the regressor X_e is not incorporated in the model, possibly due to the fact that its relevance for the considered data was unknown to the user. Now the model estimator becomes:

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= (X_r^T X_r)^{-1} X_r^T \mathbf{Y} = (X_r^T X_r)^{-1} X_r^T [X_r \theta_0^{(1)} + X_e \theta_0^{(2)} + \mathbf{E}_0] && (5.38) \\ &= \theta_0^{(1)} + (X_r^T X_r)^{-1} X_r^T X_e \theta_0^{(2)} + (X_r^T X_r)^{-1} X_r^T \mathbf{E}_0. && (5.39) \end{aligned}$$

If $\mathbf{E}_0$ is zero-mean, and X_r and X_e are deterministic then $\mathbb{E}[\hat{\boldsymbol{\theta}}] = \theta_0^{(1)} + (X_r^T X_r)^{-1} X_r^T X_e \theta_0^{(2)}$ and as a result

$$\begin{aligned} cov(\hat{\boldsymbol{\theta}}) &= \mathbb{E}\left[(\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}})^T\right] && (5.40) \\ &= (X_r^T X_r)^{-1} X_r^T \sigma^2 I X_r (X_r^T X_r)^{-1} && (5.41) \\ &= \sigma^2 (X_r^T X_r)^{-1}. && (5.42) \end{aligned}$$

The resulting estimator will be biased now, but its variance expression still matches the general formula (5.37), where X then has to be interpreted as the regressor matrix that is used in the chosen model.

Example 5.3 (Estimation of a physical variable from 5 different measurements) Consider a physical variable θ_0 that is measured by 5 different instruments, each having a different level of measurement noise. One can think , e.g., of the measurement of a temperature with different instruments. The available instruments are given by

$$\begin{aligned} \mathbf{y}_1 &= \theta_0 + \mathbf{e}_1 \\ \mathbf{y}_2 &= \theta_0 + \mathbf{e}_2 \\ \mathbf{y}_3 &= \theta_0 + \mathbf{e}_3 \\ \mathbf{y}_4 &= \theta_0 + \mathbf{e}_4 \\ \mathbf{y}_5 &= \theta_0 + \mathbf{e}_5 \end{aligned}$$

where $\mathbf{e}_i$ is the random (additive) error that is induced by instrument number i. These "errors" are supposed to be zero-mean random variables with variance σ_i^2, such that

$$\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_5.$$

There are several candidate estimators for estimating the true parameter value θ_0:

- Choose the measurement of the "best" instrument, i.e., the instrument with the smallest variance error. This implies:
$$\hat{\boldsymbol{\theta}} = \mathbf{y}_1. \tag{5.43}$$
This estimator is unbiased, and has a variance of σ_1^2.

- Combine all measurement into a simple linear least squares estimator. The corresponding model is: $\mathbf{y}_i = \theta + \mathbf{e}_i$, and the least squares solution is obtained by minimizing $\sum_{i=1}^{5} e_i^2$ is given by
$$\hat{\boldsymbol{\theta}} = (X^T X)^{-1} X^T \mathbf{Y}$$
with $X = [1 \cdots 1]^T$ and $\mathbf{Y} = [\mathbf{y}_1 \cdots \mathbf{y}_5]^T$, leading to
$$\hat{\boldsymbol{\theta}} = \frac{1}{5} \sum_{i=1}^{5} \mathbf{y}_i.$$
The estimator is simply obtained as the average of the 5 different instrument measurements. Note that again $\mathbb{E}\hat{\boldsymbol{\theta}} = \theta_0$, and so the estimator is again unbiased.
The variance of the estimator is determined by
$$var(\hat{\boldsymbol{\theta}}) = \mathbb{E}\left[(X^T X)^{-1} X^T \mathbf{E}\mathbf{E}^T X (X^T X)^{-1}\right]$$
which with X as given above, reduces to
$$\frac{1}{25}[1 \cdots 1] \begin{bmatrix} \sigma_1^2 & & \\ & \ddots & \\ & & \sigma_5^2 \end{bmatrix} \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}$$
so that
$$var(\hat{\boldsymbol{\theta}}) = \frac{1}{25} \sum_{i=1}^{5} \sigma_i^2.$$

Note that it is not automatic that the variance of the second estimator is smaller than that of the first estimate. If all σ_i are equal, then averaging the 5 results improves the final variance with a factor 5; however when instruments 2 until 5 are much worse in quality than instrument number 1, averaging may even deteriorate the final estimator variance to become worse than the first one.

Confidence regions and intervals

A $100(1-\alpha)\%$ confidence region is a region in the parameter space that covers the true but unknown parameter vector θ_0 with probability $1-\alpha$. Similarly, a $100(1-\alpha)\%$ confidence interval for an element $\theta_0^{(r)}$ of θ_0 is an interval that covers the true element $\theta_0^{(r)}$ of θ_0 with probability $1-\alpha$. If we consider the situation that the noise terms on the output variables

are zero-mean, uncorrelated, Gaussian distributed random variables with equal variance σ^2, then the linear regression estimator

$$\hat{\boldsymbol{\theta}} = (X^T X)^{-1} \cdot X^T \mathbf{Y}$$

has a Gaussian distribution $\mathcal{N}(\theta_0, \Sigma)$ with covariance matrix

$$\Sigma = \sigma^2 \cdot (X^T X)^{-1}.$$

This also implies that

$$\hat{\boldsymbol{\theta}}^{(r)} \in \mathcal{N}({\theta_0}^{(r)}, [\Sigma]_{rr})$$

with $\hat{\boldsymbol{\theta}}_0^{(r)}$ the rth element of $\hat{\boldsymbol{\theta}}$ and $[\Sigma]_{rr}$ the rth diagonal element of Σ. In this case

$$\frac{\hat{\boldsymbol{\theta}}^{(r)} - \theta_0^{(r)}}{\sqrt{[\Sigma]_{rr}}} \in \mathcal{N}(0,1)$$

and

$$(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) \in \chi_d^2$$

which means that the left hand side expression follows a χ^2 distribution with d degrees of freedom, with d the dimension of the parameter vector θ_0. Hence, the following probability statements hold true:

$$P\left[-\lambda_{1-\alpha/2} < \frac{\hat{\boldsymbol{\theta}}^{(r)} - \theta_0^{(r)}}{\sqrt{[\Sigma]_{rr}}} < \lambda_{1-\alpha/2}\right] = 1 - \alpha,$$

and

$$P\left[(\hat{\boldsymbol{\theta}} - \theta_0)^T \Sigma_{\hat{\boldsymbol{\theta}}}^{-1} (\hat{\boldsymbol{\theta}} - \theta_0) < \chi_{d,1-\alpha}^2\right] = 1 - \alpha$$

with $\chi_{d,1-\alpha}^2$ the $(1-\alpha)$ quantile of a chi-square distribution with d degrees of freedom, and $\lambda_{1-\alpha/2}$ the $(1-\alpha/2)$ quantile of the standard normal distribution. The α quantile of the distribution of a continuous random variable $\mathbf{x}$ is defined as the smallest number ξ satisfying $F_{\mathbf{x}}(\xi) = \alpha$, with $F_{\mathbf{x}}(x)$ the distribution function of $\mathbf{x}$. Consequently, an (ellipsoidal) $100(1-\alpha)\%$ confidence region for θ_0 is given by:

$$\left\{\theta | (\hat{\boldsymbol{\theta}} - \theta)^T \Sigma^{-1} (\hat{\boldsymbol{\theta}} - \theta) < \chi_{d,1-\alpha}^2\right\} \tag{5.44}$$

and a $100(1-\alpha)\%$ confidence interval for the parameter $\theta_0^{(r)}$ is given by

$$\hat{\boldsymbol{\theta}}^{(r)} \pm \lambda_{1-\alpha/2} \sqrt{[\Sigma]_{rr}}. \tag{5.45}$$

5.3.4 Extensions

The linear regression estimator or the linear least squares (LS) estimator, has been derived for the situation of a two dimensional regressor

$$\phi_i^T = [1 \quad u_i].$$

However all expressions considered in this analysis can straightforwardly be extended to the situation of any number of d regressors and d unknown parameters, by extending ϕ_i as , e.g.,

$$\phi_i^T = [1 \; u_i \; v_i \; w_i \; \cdots \; x_i] \in \mathbb{R}^d.$$

The several variables in the regressor are variables that are considered to be related to the observed variable y_i that one wants to explain from the independent variables u_i, v_i etc. The corresponding parameter vector θ will then become a $d \times 1$ vector, and the expressions for the parameter estimator stay as they were:

$$\hat{\boldsymbol{\theta}} = [\sum_i \phi_i \phi_i^T]^{-1} \sum_i \phi_i \mathbf{y}_i \tag{5.46}$$

$$= (X^T X)^{-1} X^T \mathbf{Y} \tag{5.47}$$

now with X a $n \times d$ matrix and Y a $n \times 1$ vector.

So far we have considered the situation that the regressor variables u_i are exactly known, while the measured variables y_i are assumed to be noise disturbed.
However the converse assumption is of course also possible: y_i exactly known, and u_i disturbed by noise. In that case a similar analysis as given in this section can be applied leading to a least-squares criterion that minimizes the distance from the measured points to a straight line in the direction of the u variable. This is illustrated in figure 5.8(left). Note that when starting with a cloud of measurement points it will clearly make a difference, and lead to different parameter estimators, if one chooses for one or the other option.
A symmetric treatment of the y and u variables leads to the suggestion to determine the straight line that minimizes the distance to the measurement points when taking orthogonal projections onto the straight line. This is illustrated in figure 5.8(right) and will be further analyzed in section 5.4.

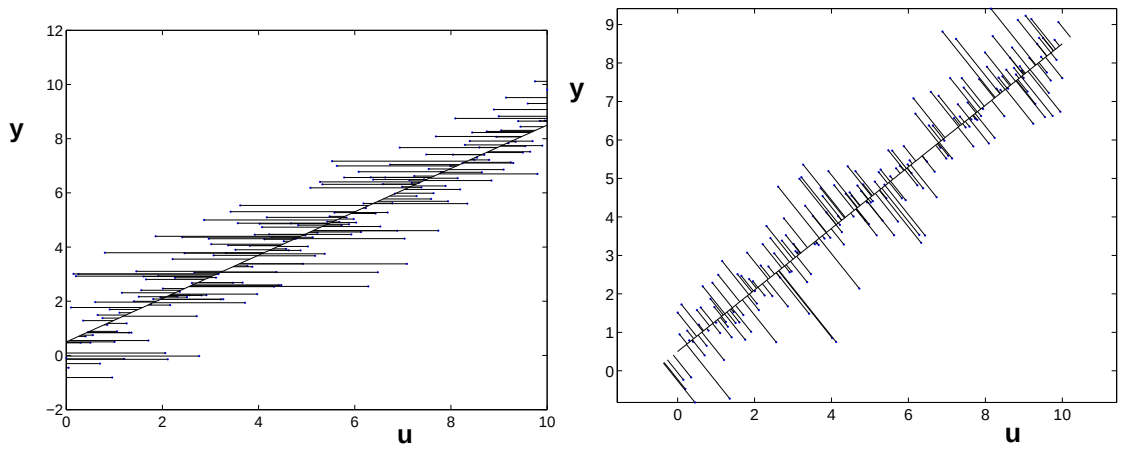

Figure 5.8: Linear regression problem with y as independent (exactly known) variable, and u as dependent variable (left); regression problem where an orthogonal projection to the straight line is used to determine least distance (right).

Example 5.4 (Absorption coefficient of an optical fiber (Frieden, 2001)) A physical measurement problem does not always look as a regression problem, but can often be written in this format.
Consider the problem of experimentally determining the absorptance coefficient of an optical fiber. This can be done by measuring the light intensity $i(x)$ at several measurement locations x along the fiber. The -theoretical- relation between light intensity and location is given by

$$I(x) = I_0 e^{-\alpha x}$$

with $I(x)$ the light intensity at distance x from the light source, α the absorption coefficient, and I_0 the light intensity of the source.
If the initial intensity I_0 is unknown, there are basically two unknown parameters: I_0 and α. Measurement of light intensity at location x is done by cutting the fiber, and measuring the light intensity. This is done for N different values of x.
In this situation it does not seem to be possible to write the system equation in the form:

$$I(x) = \phi^T \theta$$

as is required for a linear least squares estimator to be constructed.
However if we take the logarithm of the relation then:

$$\begin{aligned} \log I(x) &= \log I_0 - \alpha x && (5.48) \\ &= \phi^T \theta && (5.49) \end{aligned}$$

with

$$\phi = \begin{pmatrix} 1 \\ -x \end{pmatrix} \quad \text{and} \quad \theta = \begin{pmatrix} \log I_0 \\ \alpha \end{pmatrix}.$$

By taking N measurements $\{x_i, I(x_i)\}_{i=1,\cdots N}$ one can now construct a linear least squares estimator:

$$\hat{\theta}_N = [\sum_i \phi_i \phi_i^T]^{-1} \sum_i \phi_i \log I(x_i).$$

In order for this estimator to be unbiased the noise on the measurement data has to be such that

$$\log I(x_i) = \log I_0 - \alpha x_i + e_i$$

with e_i realizations of a zero-mean stochastic process. In terms of the original relation for $I(x)$, this implies that there is a multiplicative noise contribution

$$I(x) = I_0 e^{-\alpha x} e^{\mathbf{e}}$$

with $\mathbf{e}$ a zero-mean stochastic process.

5.3.5 Weighted linear least squares estimation

In the linear regression problem the residual "error" at every measurement point is weighted with a constant weighting over all measurements. In other words: all data is equally weighted. In the weighted least squares approach, an additional weighting factor is introduced that allows for different weighting of the several measurements. This will be shown to provide means to arrive at estimators with improved statistical properties.

The *weighted least squares criterion* is formulated as:

$$V(\theta) = (Y - X\theta)^T W (Y - X\theta)$$

with W a symmetric positive definite $n \times n$ matrix that is called the *weighting matrix*. The weighted least squares criterion may alternatively be written as

$$V(\theta) = \sum_i \sum_j w_{ij}(y_i - \phi_i^T\theta)(y_j - \phi_j^T\theta).$$

with $\{w_{i,j}\}_{i,j=1,\cdots,n}$ the elements of W. Note that for the special case of a diagonal weighting matrix $W = diag(w_1, \cdots, w_n)$, the weighted least squares criterion can also be written as

$$V(\theta) = \sum_i w_i(y_i - \phi_i^T\theta)^2.$$

The weighted least squares solution is found by minimizing the weighted least squares criterion with respect to θ. Setting the derivative of $V(\theta)$ with respect to θ to zero, the normal equations result:

$$X^T W(Y - X\hat{\theta}) = 0$$

If Y is again considered as a random variable (i.e., $Y = \mathbf{Y}$), the solution to the weighted least squares problem becomes:

$$\boxed{\hat{\boldsymbol{\theta}} = (X^T W X)^{-1} X^T W \mathbf{Y}}$$

which reduces to the simple least squares estimator when $W = I$.

Bias and variance of the weighted linear least squares estimator

For a data-generating system

$$\mathbf{Y} = X\theta_0 + \mathbf{E}_0$$

the weighted least squares estimator is given by:

$$\hat{\boldsymbol{\theta}} = (X^T W X)^{-1} X^T W \mathbf{Y} = \theta_0 + (X^T W X)^{-1} X^T W \mathbf{E}_0$$

so the estimator is unbiased if $\mathbb{E}\mathbf{E}_0 = 0$, i.e., if the noise process is zero-mean.

For the variance of the unbiased estimator one can write:

$$\begin{aligned} cov(\hat{\boldsymbol{\theta}}) &= \mathbb{E}[(\hat{\boldsymbol{\theta}} - \theta_0)(\hat{\boldsymbol{\theta}} - \theta_0)^T] \\ &= \mathbb{E}[(X^T W X)^{-1} X^T W \mathbf{E}_0 \mathbf{E}_0^T W X (X^T W X)^{-1}] \end{aligned}$$

Note that in the particular situation that $W = (\mathbb{E}[\mathbf{E}_0 \mathbf{E}_0^T])^{-1}$, the expression for the variance simplifies to

$$cov(\hat{\boldsymbol{\theta}}) = (X^T W X)^{-1}. \tag{5.50}$$

This particular choice of weighting matrix is not just chosen for simplicity; it appears to have an important property. It is the - optimal - weighting that minimizes the variance over all possible choices of weighting matrices. More precisely, the weighted least squares estimator with as weighting matrix the inverse of the covariance matrix of the noise (i.e., with $W = \Sigma^{-1}$) has minimum variance within the class of weighted least squares estimators, that is, the difference of the covariance matrix of any estimator of this class and its covariance matrix is positive semidefinite.

Lemma 5.5 Let Σ be a positive definite matrix, and define

$$P(W) = (X^T W X)^{-1} X^T W \Sigma W X (X^T W X)^{-1}.$$

Then for all symmetric, positive semi-definite W, it holds that

$$P(\Sigma^{-1}) \leq P(W).$$

This inequality expresses that the difference between the left-hand and right-hand member is negative semi-definite. A property of a negative semi-definite matrix is that its diagonal elements cannot be positive. This means that the diagonal elements of $P(\Sigma^{-1})$ are smaller than or equal to the corresponding diagonal elements of $P(W)$.

Proof: The matrix

$$\begin{bmatrix} X^T \\ X^T W \Sigma \end{bmatrix} \Sigma^{-1} \begin{bmatrix} X^T \\ X^T W \Sigma \end{bmatrix}^T = \begin{bmatrix} X^T \Sigma^{-1} X & X^T W X \\ X^T W X & X^T W \Sigma W X \end{bmatrix}$$

is positive semi-definite by construction.
Consider any positive semi-definite matrix

$$H = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix}$$

with C invertible. Then considering xHx^T with $x = [x_1 \quad -x_1 B C^{-1}]$ and x_1 arbitrary shows that

$$xHx^T = x_1 [A - BC^{-1}B^T] x_1.$$

If H is positive semi-definite, then $(A - BC^{-1}B^T)$ is also positive semi-definite. Applying this result to the above expressions, shows that

$$X^T \Sigma^{-1} X - X^T W X [X^T W \Sigma W X]^{-1} X^T W X \geq 0$$

or equivalently

$$X^T \Sigma^{-1} X \geq X^T W X [X^T W \Sigma W X]^{-1} X^T W X.$$

Taking inverses of both sides of the inequality leads to

$$(X^T \Sigma^{-1} X)^{-1} \leq (X^T W X)^{-1} [X^T W \Sigma W X] (X^T W X)^{-1}.$$

□

A weighted least-squares estimator, with the optimal weighting as indicated above, is also referred to as the Markov estimator or the *Best Linear Unbiased Estimator (BLUE)*. The reason for this is that it can be shown that this estimator has the minimum variance not only within the class of weighted least squares estimators, but also within the (broader) class of estimators that are linear in the observations ($\mathbf{Y}$) and are unbiased.

If we apply this BLUE to the problem as sketched in Example 5.3, then the "optimal" weighting should be chosen as

$$W = \Sigma^{-1} = \begin{bmatrix} \sigma_1^{-2} & & \\ & \ddots & \\ & & \sigma_5^{-2} \end{bmatrix}$$

The weighted least squares criterion now reads:

$$\min_{\theta} \sum_{i=1}^{5} \frac{(\mathbf{y}_i - \theta)^2}{\sigma_i^2},$$

leading to the weighted least squares solution

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= (X^T \Sigma^{-1} X)^{-1} X^T \Sigma^{-1} \mathbf{Y} \\ &= \frac{[\mathbf{y}_1/\sigma_1^2 + \cdots + \mathbf{y}_5/\sigma_5^2]}{1/\sigma_1^2 + \cdots + 1/\sigma_5^2}. \end{aligned} \tag{5.51}$$

The variance of this estimator is given by (5.50):

$$\begin{aligned} var(\hat{\boldsymbol{\theta}}) &= (X^T \Sigma^{-1} X)^{-1} \\ &= \frac{1}{1/\sigma_1^2 + \cdots + 1/\sigma_5^2}. \end{aligned} \tag{5.52}$$

Note for the first suggested estimator (5.43) the variance was given by

$$\sigma_1^2 = \frac{1}{1/\sigma_1^2 + 0}.$$

Therefore the estimator (5.51) will always have a smaller variance than the estimator (5.43).

Consequently: the use of all five measurement devices reduces the variance of the final temperature estimator, and is to be preferred over simply choosing the measurement from the device with the highest precision. However in order to find the optimal weightings that minimize the overall variance, knowledge of the variance of the several instruments is required.

5.4 Total least squares

5.4.1 Introduction

So far we have considered the problem of finding a linear relation between measured variables where we have chosen one of the variables to be given as deterministic, and one variable to contain a random error. In this way a linear regression estimator comes down to minimizing the sum of squared distances between the measurement points and the optimal straight line, where "distance" is measured in the direction of the variable that is supposed to contain the (random) measurement error.

However it may happen (and maybe it is even more realistic in many cases) that both measured variables will be contaminated with measurement noise. In that case the usual regression techniques, as discussed so far, do not suffice. First it will briefly be analyzed how a linear least squares estimator will behave when all data is contaminated with noise. Next an algorithm will be presented leading to an estimator that can appropriately handle the considered situation.

In order to keep the exposition of the mathematical details at a comprehensible level, a restriction will be made to linear regression models for which the regressor matrix contains measured variables only, and does not contain constants, as , e.g., the constant term 1 in the regressor ϕ_i in section 5.3.4. In the scope of Figure 5.8 this implies that in this section we will restrict attention to linear relationships that pass through the origin $(y, u) = (0, 0)$.

5.4.2 Properties of the linear least squares estimator in the presence of disturbances on all data

If both measurements of X and Y are contaminated with random noise the following model equations can be formalized:

$$\begin{aligned} Y_0 &= X_0\theta_0 && (5.53)\\ \mathbf{Y} &= Y_0 + \mathbf{E}_y && (5.54)\\ \mathbf{X} &= X_0 + \mathbf{E}_x && (5.55) \end{aligned}$$

where X_0, Y_0 are the exact (noise-free) variables that are assumed to be related through the exact parameter value θ_0, and $\mathbf{X}, \mathbf{Y}$ are the measurement variables that are available, being contaminated with the random variables $\mathbf{E}_x$ and $\mathbf{E}_y$ respectively.
When applying the linear least squares estimator, the following parameter estimator is obtained on the basis of n data points:

$$\hat{\boldsymbol{\theta}}_n = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T[X_0\theta_0 + \mathbf{E}_y] \tag{5.56}$$

For an analysis of this expression, assume that each row of the random matrix $[\mathbf{E}_x \;\; \mathbf{E}_y]$ is a vector-valued random variable with zero mean and covariance matrix

$$\Sigma = \begin{pmatrix} \Sigma_x & \Sigma_{xy} \\ \Sigma_{yx}^T & \Sigma_y \end{pmatrix},$$

and that the several rows are independent and identically distributed. Let $\mathbf{z}_i^T$ be the i-th row of $[\mathbf{E}_x \;\; \mathbf{E}_y]$. Then $\mathbb{E}\mathbf{z}_i\mathbf{z}_i^T = \Sigma$ for all i, and consequently

$$\mathbb{E}[\mathbf{E}_x \;\; \mathbf{E}_y]^T[\mathbf{E}_x \;\; \mathbf{E}_y] = \mathbb{E}[\mathbf{z}_1 \; \mathbf{z}_2 \; \cdots \mathbf{z}_n] \begin{pmatrix} \mathbf{z}_1^T \\ \mathbf{z}_2^T \\ \vdots \\ \mathbf{z}_n^T \end{pmatrix} = n \cdot \Sigma.$$

For analyzing the statistical properties of the estimator (5.56) we need to consider its expected value, which requires taking the expectation of an expression of products and quotients of correlated random variables. In the general case this is quite hard to do. However when using the notion of probability limit (plim) (see the definition of consistent estimators, page 82) an asymptotic analysis can be made.
For the estimator

$$\hat{\boldsymbol{\theta}}_n = [(X_0 + \mathbf{E}_x)^T(X_0 + \mathbf{E}_x)]^{-1}(X_0 + \mathbf{E}_x)^T(X_0\theta_0 + \mathbf{E}_y) \tag{5.57}$$

it follows that under the considered assumptions

$$\begin{aligned} \operatorname*{plim}_{n\to\infty} \frac{1}{n}[(X_0 + \mathbf{E}_x)^T(X_0 + \mathbf{E}_x)] &= \lim_{n\to\infty} \frac{1}{n} X_0^T X_0 + \Sigma_x && (5.58)\\ \operatorname*{plim}_{n\to\infty} \frac{1}{n}(X_0 + \mathbf{E}_x)^T(X_0\theta_0 + \mathbf{E}_y) &= \lim_{n\to\infty} (\frac{1}{n} X_0^T X_0)\theta_0 + \Sigma_{xy}. && (5.59) \end{aligned}$$

Denoting $A_0 = \lim_{n\to\infty}(\frac{1}{n}X_0^T X_0)$ and substituting the above expressions into (5.57) then delivers:[7]

$$\begin{aligned} \operatorname*{plim}_{n\to\infty} \hat{\boldsymbol{\theta}}_n &= [A_0 + \Sigma_x]^{-1}[A_0\theta_0 + \Sigma_{xy}] && (5.60)\\ &= [I + A_0^{-1}\Sigma_x]^{-1}[\theta_0 + A_0^{-1}\Sigma_{xy}], && (5.61) \end{aligned}$$

provided that the limit A_0 exists and that it is invertible.
This result shows that the parameter estimator will generally be biased, in contrast with the situation when the variables in X are noise free. The latter situation can be obtained by substituting $\Sigma_x = \Sigma_{xy} = 0$.
The bias of the estimator can be constructed by rewriting (5.60) as

$$\begin{aligned} \operatorname*{plim}_{n\to\infty} \hat{\boldsymbol{\theta}}_n &= (A_0 + \Sigma_x)^{-1}[(A_0 + \Sigma_x - \Sigma_x)\theta_0 + \Sigma_{xy}] && (5.62)\\ &= \theta_0 + (A_0 + \Sigma_x)^{-1}[-\Sigma_x\theta_0 + \Sigma_{xy}]. && (5.63) \end{aligned}$$

If the noise terms in $\mathbf{X}$ and $\mathbf{Y}$ are uncorrelated, i.e., $\Sigma_{xy} = 0$, the expression for the parameter estimator simplifies to

$$\operatorname*{plim}_{n\to\infty} \hat{\boldsymbol{\theta}}_n = [I + A_0^{-1}\Sigma_x]^{-1}\theta_0 \tag{5.64}$$

and the bias expression (5.63) becomes

$$\operatorname*{plim}_{n\to\infty} \hat{\boldsymbol{\theta}}_n = \theta_0 - (A_0 + \Sigma_x)^{-1}\Sigma_x\theta_0. \tag{5.65}$$

Apparently the least squares estimator is asymptotically biased now, and the bias is dependent on the noise-to-signal ratio in the measured variables $\mathbf{X}$; here Σ_x is a reflection of the noise signal, and $(A_0 + \Sigma_x)$ a reflection of the measured signal properties of $\mathbf{X}$.

5.4.3 Total least squares

The linear least squares problem can be generalized to a situation that both X and Y are contaminated by errors. This is reflected in the following (deterministic) model for the measurements:

$$Y + E_y = (X + E_x)\theta \tag{5.66}$$

where E_y and E_x are perturbations of the measured variables Y and X respectively.
As in the case of the linear least squares method, we will first consider all data to be deterministic, leading to a model-fit problem on the basis of a particular cost function. In the next subsection we will then analyse the related statistical estimation method with the corresponding total least squares estimator.
The total least squares problem, is formulated as to finding the solution for θ that satisfies (5.66) and that minimizes the Frobenius norm:

$$\|[E_y \mid E_x]\|_F = \sqrt{tr([E_y \mid E_x][E_y \mid E_x]^T)}.$$

[7] Use is made of the so-called Theorem of Slutsky, that states that for two estimators $\hat{\boldsymbol{\theta}}_n^{(1)}$, $\hat{\boldsymbol{\theta}}_n^{(2)}$ with probability limits $\theta^{(1)}$ and $\theta^{(2)}$ it holds that $\operatorname{plim}_{n\to\infty} \hat{\boldsymbol{\theta}}_n^{(1)} \cdot \hat{\boldsymbol{\theta}}_n^{(2)} = \theta^{(1)}\theta^{(2)}$, and for any continuous function h of $\hat{\boldsymbol{\theta}}_n^{(1)}$, $\operatorname{plim}_{n\to\infty} h(\hat{\boldsymbol{\theta}}_n^{(1)}) = h(\theta^{(1)})$.

I.e., on the basis of n measured data Y, X the total least squares estimate $\hat{\theta}_n$ is determined by

$$\hat{\theta}_n = \arg\min_{\theta, E_x, E_y} \|[E_y \mid E_x]\|_F \tag{5.67}$$

$$\text{under the constraint that } Y + E_y = (X + E_x)\theta \tag{5.68}$$

The squared Frobenius norm of a matrix is given by the sum of squared values of all matrix elements. If X contains only one regressor, as is the situation in figure 5.8(right) when the linear curve is restricted to pass through the origin, the Frobenius norm can simply be written as $\sqrt{\sum_{i=1}^{n} E_{x,i}^2 + E_{y,i}^2}$, where index i refers to the noise term on data sample i. In this situation it follows directly that minimizing this Frobenius norm is (with Pythagoras) equal to minimizing the sum of distances from the measurement points to the (perpendicular) projected points on the estimated line.

The equation (5.66) can be rewritten into the form

$$[\bar{X} + \bar{\Delta}] \begin{bmatrix} -1 \\ \theta \end{bmatrix} = 0 \tag{5.69}$$

with

$$\bar{X} = [Y \mid X] \tag{5.70}$$

$$\bar{\Delta} = [E_y \mid E_x]. \tag{5.71}$$

If the matrix $[\bar{X} + \bar{\Delta}]$ is regular (i.e., it has full column rank), the only -trivial- solution to the equation $[\bar{X} + \bar{\Delta}]x = 0$ will be $x = 0$. Therefore in order for a non-trivial θ to exist, the columns of the matrix $[\bar{X} + \bar{\Delta}]$ have to be linearly dependent, i.e., the matrix has to be singular. The problem is now to find a matrix $\bar{\Delta}$ with minimum Frobenius norm such that $[\bar{X} + \bar{\Delta}]$ becomes singular.
This latter problem can be solved through a singular value decomposition of $\bar{X}$.
Note that the dimensions of $\bar{X}$ are $n \times m$, with n the number of data points available, and $m = d+1$, d being the number of regressors in X. It is assumed that $n > m$. Now consider the singular value decomposition (see appendix B.2):

$$\bar{X} = \sum_{k=1}^{m} \sigma_k U_k V_k^T$$

then the matrix $\bar{\Delta}$ with minimum Frobenius norm that renders $\bar{X}$ singular, is given by

$$\bar{\Delta} = -\sigma_m U_m V_m^T.$$

In other words: compensating a matrix by its component that is induced by its smallest singular value, delivers a matrix which is reduced in rank by one, and at the same time it is the smallest compensating term (in Frobenius norm) that achieves this goal.
Substituting this into the equation (5.69) shows that

$$\sum_{k=1}^{m-1} \sigma_k U_k V_k^T \begin{bmatrix} -1 \\ \hat{\theta}_n \end{bmatrix} = 0.$$

The vector $(-1 \;\; \hat{\theta}_n^T)^T$ should be orthogonal to all singular vectors V_k, $k = 1, \cdots m-1$. The only vector for which this holds is given by V_m.
Write V_m as

$$V_m = \begin{bmatrix} v_1 \\ V_{m,r} \end{bmatrix}$$

with v_1 scalar, and $V_{m,r}$ a d dimensional vector, then it follows that

$$\hat{\theta}_n = -\frac{1}{v_1} V_{m,r}. \tag{5.72}$$

Summarizing, the total least squares estimate is obtained by applying an SVD to the data matrix $\bar{X}$, and by extracting the parameter estimate from the (right) singular vector V_m that is related to the smallest singular value.
The presented algorithm provides the best straight line in the cloud of measurement points y_i, u_i by minimizing the orthogonal distance of measurement points to the straight line. For this reason this method is also referred to as *orthogonal least squares* or *errors-in-variables (EIV)* method.

5.4.4 Statistical properties of TLS-estimators

If the measurement data X, Y are considered to be random variables, denoted by $\mathbf{X}$ and $\mathbf{Y}$, the total least squares solution (5.72) becomes a random variable, and therefore an estimator, that is determined by

$$\hat{\boldsymbol{\theta}}_n = \arg\min_{\theta} \min_{\theta, E_y, E_x} \|[E_y \mid E_x]\|_F \tag{5.73}$$

$$\text{under the constraint that } \mathbf{Y} + E_y = (\mathbf{X} + E_x)\theta. \tag{5.74}$$

The properties of this estimator can now be analyzed. The principal result is formulated in the following proposition.

Proposition 5.6 Assume that the measured data in $\mathbf{X}$ and $\mathbf{Y}$ satisfies the equations

$$Y_0 = X_0\theta_0 \tag{5.75}$$

$$\mathbf{Y} = Y_0 + \mathbf{E}_y \tag{5.76}$$

$$\mathbf{X} = X_0 + \mathbf{E}_x \tag{5.77}$$

where X_0, Y_0 are the exact (noise-free) variables that are assumed to be related through the exact parameter value θ_0, and that the rows of the matrix $[\mathbf{E}_y \mid \mathbf{E}_x]$ are independent and identically distributed vector random variables with zero mean and covariance $\Sigma = \sigma^2 I_m$. If $\lim_{n\to\infty} \frac{1}{n} X_0^T X_0$ exists and is positive definite, then for $n \to \infty$, $\hat{\boldsymbol{\theta}}_n$ (5.73) converges to θ_0 with probability one, and therefore the estimator is consistent.

The proof of this proposition is beyond the scope of the current text. It can be found in Gleser (1981). More extensive properties of the TLS estimator, including asymptotic distribution and variance properties can be found in Van Huffel and Vandewalle (1991).
This result shows that the TLS-estimator is consistent under the assumption that the noise disturbance on all regression and output variables are independent, with fixed and equal variance over all variables. In particular this latter situation is rather restrictive. If this

condition is not satisfied, i.e., if each row of $[\mathbf{E}_y \mid \mathbf{E}_x]$ has a covariance matrix that is not $\sigma^2 I$, the data $\mathbf{Y}$ and $\mathbf{X}$ can be scaled and transformed by a linear operation such that transformed data $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{X}}$ occurs that does satisfy the covariance matrix assumption.
Suppose that the rows of $[\mathbf{E}_y \mid \mathbf{E}_x]$ have a covariance matrix Σ, then Σ can be decomposed in an eigenvalue decomposition, as

$$\Sigma = W \Lambda W^T$$

equivalent to the situation in section 5.2.2.
By denoting

$$\Sigma = B B^T$$

with $B = W \Lambda^{\frac{1}{2}}$, the original model equations can be transformed from

$$[\bar{\mathbf{X}} + \bar{\mathbf{\Delta}}] \begin{bmatrix} -1 \\ \theta \end{bmatrix} = 0$$

with $\bar{\mathbf{X}} = [\mathbf{Y} \mid \mathbf{X}]$ and $\bar{\mathbf{\Delta}} = [\mathbf{E}_y \mid \mathbf{E}_x]$, into

$$[\bar{\mathbf{X}} + \bar{\mathbf{\Delta}}] B^{-T} B^T \begin{bmatrix} -1 \\ \theta \end{bmatrix} = 0$$

and

$$[\bar{\mathbf{X}} B^{-T} + \bar{\mathbf{\Delta}} B^{-T}] B^T \begin{bmatrix} -1 \\ \theta \end{bmatrix} = 0.$$

In this latter expression the covariance of the rows of $\bar{\mathbf{\Delta}}$ now will be the identity matrix, and thus satisfying the consistency conditions of the TLS estimator.
The algorithm now follows by taking the SVD of the transformed data matrix $\bar{X} B^{-T}$, leading to the m-th right singular vector $\mathbf{V}_m$, and by solving for the equation

$$B^T \begin{bmatrix} -1 \\ \hat{\boldsymbol{\theta}}_n \end{bmatrix} = \mathbf{V}_m$$

or equivalently

$$\begin{bmatrix} -1 \\ \hat{\boldsymbol{\theta}}_n \end{bmatrix} = B^{-T} \mathbf{V}_m.$$

With this algorithm the situation of noise disturbance on $\mathbf{Y}$ and $\mathbf{X}$ can be handled, leading to a consistent estimator of the model parameter. However knowledge of the covariance matrix of the several noise terms is required to construct this estimator.

5.5 The Cramér-Rao lower bound

The estimators presented so far show different properties for the variance (i.e., the covariance matrix) of the estimated parameter. This raises the fundamental question whether, in a particular estimation problem, there exists a lower bound on the reachable estimator variance. The answer to this question is one of the most powerful results in estimation

theory, and is known as the Cramér-Rao lower bound (CRLB), named after the work of Cramér (1946) and Rao (1945)[8].

Cramér-Rao lower bound (CRLB)

Consider observations from a random variable $\mathbf{y}$ with probability density function $f_{\mathbf{y}}(y, \theta)$, where θ is the unknown parameter. Then for *any* unbiased estimator $\hat{\boldsymbol{\theta}}$ of the parameter θ, its covariance matrix satisfies the inequality [9]

$$cov(\hat{\boldsymbol{\theta}}) \geq J^{-1} \tag{5.78}$$

with the *Fisher Information Matrix*:

$$J = \mathbb{E}\left[-\frac{\partial^2}{\partial \theta^2} \log f_{\mathbf{y}}(\mathbf{y}; \theta) \Big|_{\theta=\theta_0} \right] \tag{5.79}$$

The logarithm of the probability density function of the measurement data determines the lower bound on the variance of any unbiased parameter estimator.
Note that the measurable random variable $\mathbf{y}$ will generally be an n-dimensional vector, and the corresponding probability density function $f_{\mathbf{y}}$ a multivariate p.d.f. With θ being a d-dimensional vector of parameters, the second partial derivative of a scalar function g with respect to θ is a $d \times d$ matrix defined by

$$\frac{\partial^2}{\partial \theta^2} g = \begin{pmatrix} \frac{\partial^2 g}{\partial \theta_1 \partial \theta_1} & \frac{\partial^2 g}{\partial \theta_1 \partial \theta_2} & \cdots & \frac{\partial^2 g}{\partial \theta_1 \partial \theta_d} \\ \frac{\partial^2 g}{\partial \theta_2 \partial \theta_1} & \frac{\partial^2 g}{\partial \theta_2 \partial \theta_2} & \cdots & \frac{\partial^2 g}{\partial \theta_2 \partial \theta_d} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 g}{\partial \theta_d \partial \theta_1} & \frac{\partial^2 g}{\partial \theta_d \partial \theta_2} & \cdots & \frac{\partial^2 g}{\partial \theta_d \partial \theta_d} \end{pmatrix}.$$

Proof of the Cramér-Rao inequality

Since $f_{\mathbf{y}}(y; \theta)$ is a probability density function, it holds that

$$\int_{-\infty}^{\infty} f_{\mathbf{y}}(y; \theta) dy = 1 \quad \text{for all } \theta \tag{5.80}$$

where the integral over dy should be considered to be taken over the n-dimensional measurement space, when the measurement data is n-dimensional. In particular it will follow that

$$\int_{-\infty}^{\infty} \theta f_{\mathbf{y}}(y; \theta) dy = \theta \quad \text{for all } \theta. \tag{5.81}$$

[8] Actually Fisher (1922) had shown a similar result two decades earlier.
Ronald Aylmer Fisher (1890-1962) was a British mathematician who played a key role in the development of modern probability theory.

[9] The expression for the Fisher Information Matrix requires the second partial derivative of $\log f$ (the natural logarithm) with respect to θ to exist and to be absolutely integrable.

The expected value of an estimator $\hat{\theta}(\mathbf{y})$ is given by

$$\mathbb{E}[\hat{\theta}(\mathbf{y})] = \int_{-\infty}^{\infty} \hat{\theta}(y) f_{\mathbf{y}}(y;\theta_0) dy.$$

Consequently one can write that for every unbiased estimator $\hat{\boldsymbol{\theta}} = \hat{\theta}(\mathbf{y})$ it follows that

$$\mathbb{E}[\hat{\boldsymbol{\theta}} - \theta_0] = \int_{-\infty}^{\infty} [\hat{\theta}(y) - \theta_0] f_{\mathbf{y}}(y;\theta_0) dy = 0.$$

Basically this expression holds true for all θ_0, as the unbiasedness of the estimator $\hat{\boldsymbol{\theta}}$ is not dependent on the particular choice of θ_0. As a result the above equation should also hold when differentiated with respect to θ_0.
If the partial derivative $\frac{\partial f_{\mathbf{y}}(y;\theta)}{\partial \theta}$ exists and is absolutely integrable, then

$$\frac{\partial}{\partial \theta_0^T} \int_{-\infty}^{\infty} [\hat{\theta}(y) - \theta_0] f_{\mathbf{y}}(y;\theta_0) dy = \tag{5.82}$$

$$= \int_{-\infty}^{\infty} [\hat{\theta}(y) - \theta_0] \left[\frac{\partial f_{\mathbf{y}}(y;\theta_0)}{\partial \theta_0} \right]^T dy - \int_{-\infty}^{\infty} I \cdot f_{\mathbf{y}}(y;\theta_0) dy = 0 \tag{5.83}$$

showing that

$$\int_{-\infty}^{\infty} [\hat{\theta}(y) - \theta_0] \left[\frac{\partial f_{\mathbf{y}}(y;\theta_0)}{\partial \theta_0} \right]^T dy = I.$$

With the expression for the derivative of the logarithm, this leads to

$$\int_{-\infty}^{\infty} [\hat{\theta}(y) - \theta_0] \left[\frac{\partial \log f_{\mathbf{y}}(y;\theta_0)}{\partial \theta_0} \right]^T f_{\mathbf{y}}(y;\theta_0) dy = I,$$

or

$$\mathbb{E}\left[[\hat{\boldsymbol{\theta}} - \theta_0] \left. \frac{\partial \log f_{\mathbf{y}}(\mathbf{y};\theta)}{\partial \theta^T} \right|_{\theta=\theta_0} \right] = I. \tag{5.84}$$

In the scalar case ($d = 1$) one can now take the square of this relation, and apply Schwartz inequality $(\mathbb{E}[\mathbf{xy}])^2 \leq \mathbb{E}[\mathbf{x}^2] \cdot \mathbb{E}[\mathbf{y}^2]$, to show that

$$1 \leq \mathbb{E}[(\hat{\boldsymbol{\theta}} - \theta_0)^2] \cdot \mathbb{E}\left[\left(\left. \frac{\partial}{\partial \theta} \log f_{\mathbf{y}}(\mathbf{y};\theta) \right|_{\theta=\theta_0} \right)^2 \right]$$

leading to

$$\mathbb{E}[\hat{\boldsymbol{\theta}} - \theta_0]^2 \geq \left\{ \mathbb{E}\left[\left(\left. \frac{\partial}{\partial \theta} \log f_{\mathbf{y}}(\mathbf{y};\theta) \right|_{\theta=\theta_0} \right)^2 \right] \right\}^{-1}. \tag{5.85}$$

In the multivariate situation ($d > 1$) a different route has to be followed. In that situation we rewrite (5.84) into the form

$$\mathbb{E}[\mathbf{a}\mathbf{b}^T] = I \tag{5.86}$$

and we consider

$$\mathbb{E}\left[\begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix} \begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}^T \right] = \begin{pmatrix} \mathbb{E}[\mathbf{a}\mathbf{a}^T] & I \\ I & \mathbb{E}[\mathbf{b}\mathbf{b}^T] \end{pmatrix}$$

which is a positive semi-definite matrix by construction. Then with the property of positive semi-definite matrices, as explained in the proof of Lemma 5.5 (page 100) it follows that

$$\mathbb{E}[\mathbf{a}\mathbf{a}^T] \geq \left(\mathbb{E}[\mathbf{b}\mathbf{b}^T]\right)^{-1},$$

leading to

$$cov(\hat{\boldsymbol{\theta}}) \geq \left\{ \mathbb{E}\left[\left(\frac{\partial}{\partial \theta} \log f_{\mathbf{y}}(\mathbf{y};\theta)\Big|_{\theta=\theta_0} \right) \left(\frac{\partial}{\partial \theta} \log f_{\mathbf{y}}(\mathbf{y};\theta)\Big|_{\theta=\theta_0} \right)^T \right] \right\}^{-1}. \tag{5.87}$$

Differentiating (5.80) shows that

$$\int_{-\infty}^{\infty} \frac{\partial f_{\mathbf{y}}(y;\theta)}{\partial \theta} dy = 0$$

or equivalently

$$\int_{-\infty}^{\infty} \frac{\partial \log f_{\mathbf{y}}(y;\theta)}{\partial \theta} f_{\mathbf{y}}(y;\theta) dy = 0.$$

Differentiating this equation with respect to θ^T gives

$$\int_{-\infty}^{\infty} \left[\frac{\partial^2 \log f_{\mathbf{y}}(y;\theta)}{\partial \theta^2} + \left(\frac{\partial \log f_{\mathbf{y}}(y;\theta)}{\partial \theta} \right) \left(\frac{\partial \log f_{\mathbf{y}}(y;\theta)}{\partial \theta} \right)^T \right] f_{\mathbf{y}}(y;\theta) dy$$

and as a result

$$\mathbb{E}\left[\left(\frac{\partial \log f_{\mathbf{y}}(\mathbf{y};\theta)}{\partial \theta} \right) \left(\frac{\partial \log f_{\mathbf{y}}(\mathbf{y};\theta)}{\partial \theta} \right)^T \right] = -\mathbb{E}\left[\frac{\partial^2 \log f_{\mathbf{y}}(\mathbf{y};\theta)}{\partial \theta^2} \right].$$

Together with (5.87) this proves the result for the Cramér-Rao bound. □

Discussion

Some comments on the nature of the lower bounds:

- The calculation of the CRLB requires exact knowledge of the probability density function of the measured random variables. This is of course a quite strict requirement, that asks for detailed knowledge of the disturbances that act on the measured data.
- The calculation of the CRLB will generally require knowledge of the exact system parameter θ_0. There are exceptions though, when the second derivative of the $\log f_{\mathbf{y}}(\mathbf{y};\theta)$ is independent of θ. This typically happens in the situation of Gaussian probability density functions, in combination of a parameter θ which occurs linearly in the measured output. See also Example 5.7.
- The CRLB is a lower bound on the variance of unbiased estimators. It does not say anything about the variance of biased estimators, and so it is possible that there exists a biased estimator that has a smaller variance, and possibly even a smaller mean-squared error.

As a result of these remarks, calculation of the CRLB is often not feasible in practical situations. However the bound is very useful in analysis questions, e.g., when evaluating properties of several estimators and when comparing the covariance matrices of several estimators.
Notice that the CRLB is not related to a particular estimation method. It depends on the statistical properties of the observed variables, the measured data, and in most cases the hypothetical true values of the parameters. As first sight, this dependence on the true values looks as a serious impediment to the practical use of the bound. However, the expressions for the bound provide the experimenter with the means to compute numerical values for it, using nominal values of the parameters. This provides the experimenter with quantitative insight in what precision (i.e., variance) might be achieved from the available observations. In addition, it provides insight in the sensitivity of the precision to the parameter values.
Another important purpose for which the expressions for the CRLB can be used, is the optimization of the experiment design, i.e., the selection of which data to measure and to use as a basis for the estimation. By calculating the CRLB, the experimenter gets an impression if for a given experiment setup the precision attainable is sufficient for the purpose concerned. If not, the experiment design has to be changed. If this is not possible, it has to be concluded that the observations are not suitable for the purpose of the measurement procedure. In this way, the experiment design can be optimized so as to attain the highest precision, i.e., the smallest variance (see , e.g., Van den Bos, 1999).
As a final remark it should be noted that existence of a lower bound on the parameter variance, does not imply that an estimator can be found that reaches this lower bound. Especially in the situation of observing random variables with a finite number of observations, it appears very hard to find the minimum variance estimator.

Example 5.7 (Continuation of Example 5.3) Consider again the 5 measurements

$$\mathbf{y}_i = \theta_0 + \mathbf{e}_i \quad i = 1, \cdots 5.$$

For calculation of the CRLB of this estimation problem, we additionally assume that the random errors $\mathbf{e}_i$ are jointly Gaussian distributed, with mean value 0 and covariance matrix $\Sigma = diag(\sigma_1^2, \cdots, \sigma_5^2)$.
Using the multivariate Gaussian distribution (A.10) it follows that

$$f_{\mathbf{y}}(\mathbf{y}, \theta) = \frac{1}{(2\pi)^{5/2}\sqrt{det\Sigma}} exp[-\frac{1}{2}(\mathbf{y} - \underline{\theta})^T \Sigma^{-1}(\mathbf{y} - \underline{\theta})]$$

with $\mathbf{y} := [\mathbf{y}_1 \ \mathbf{y}_2 \cdots \mathbf{y}_5]^T$, and $\underline{\theta} = [\theta \cdots \theta]^T$. Consequently

$$\log f_{\mathbf{y}}(\mathbf{y}, \theta) = c - \sum_{i=1}^{5} \frac{1}{2} \frac{(\mathbf{y}_i - \theta)^2}{\sigma_i^2}.$$

Taking the partial derivative with respect to θ delivers

$$\frac{\partial \log f_{\mathbf{y}}(\mathbf{y}; \theta)}{\partial \theta} = \sum_{i=1}^{5} \frac{(\mathbf{y}_i - \theta)}{\sigma_i^2}$$

and the second derivative:

$$\frac{\partial^2 \log f_{\mathbf{y}}(\mathbf{y};\theta)}{\partial \theta^2} = \sum_{i=1}^{5} \frac{-1}{\sigma_i^2}.$$

Substituting this result in the expression for the CRLB, gives:

$$var(\hat{\boldsymbol{\theta}}) \geq \frac{1}{1/\sigma_1^2 + \cdots + 1/\sigma_5^2}.$$

This provides us, in the considered situation, with the best possible variance for any unbiased estimator.
Note that the weighted least squares estimator that was analyzed in section 5.3.5 reaches exactly this lower bound of the parameter variance in expression (5.52).
The conclusions that one can draw from this, are the following

- In the problem setting of the considered example, the weighted least squares estimator with the particular weights as chosen in section 5.3.5 leads to the smallest possible variance among all unbiased weighted least squares estimators, irrespective of the probability density function of the disturbances.
- If in the problem setting of the considered example, the disturbance terms are Gaussian distributed, then the weighted least squares estimator with the above mentioned weights, has the smallest possible variance among *all* unbiased estimators.

If the probability density function of the disturbances is unknown, the weighted least squares estimator can lead to satisfactory results; however there is no guarantee that it is the best possible estimator in terms of minimum variance.

5.6 Maximum likelihood estimator

In the previous section it was shown that the probability density function of the measured random variables plays a crucial role in the smallest possible variance that can be reached for any unbiased parameter estimator.
The parameter estimation methods considered so far are based on linear (regression) techniques, and therefore they are restricted to a well-defined class, characterized by the fact that the resulting estimators lead to simple analytical functions of the measured random variables. As a result, the computational tools required for the estimators are very straightforward.
In this section a general philosophy to parameter estimation will be presented that takes account of the probability density function of the measured random variables, and therefore has the potentials to come close to the covariance expressions as induced by the CRLB.

The so-called maximum likelihood principle is based on the following reasoning.
Suppose that a (vector) random variable $\mathbf{y}$ is observed, and that the underlying probability density function of $\mathbf{y}$ is given by

$$f_{\mathbf{y}}(y, \theta)$$

where θ reflects the unknown (vector) parameter that is to be estimated.

For a given value of θ the function $f_{\mathbf{y}}(y,\theta)$ is a probability density function. However for a fixed value of y and unknown θ, the function f is a (deterministic) function of θ and referred to as the *likelihood function*, indicated by

$$L(\theta; y)$$

Maximum likelihood principle

> For a given observation y of the measured random variable $\mathbf{y}$, determine the maximum likelihood estimate as that value of θ that maximizes the likelihood function, i.e.,
>
> $$\max_{\theta} L(\theta; y).$$
>
> The resulting maximum likelihood (ML) estimator is denoted as:
>
> $$\hat{\boldsymbol{\theta}}_{ml} = \arg\max_{\theta} L(\theta; \mathbf{y}).$$

For given observed values $\mathbf{y} = y$, the maximum likelihood estimate is determined by that value of θ for which the probability density function $f_{\mathbf{y}}(y,\theta)$ reaches its maximum value. In other words: that parameter θ is sought for that generates a probability density function for which the observed measurement data was the most probable data. The idea can best be visualized in a very simple example.

Example 5.8 (Visualization of the maximum likelihood principle) We consider the linear model between the observed variables $\mathbf{y}$ and u, where u is exactly measured and known, and $\mathbf{y}$ is a random variable that satisfies :

$$\mathbf{y} = \theta_0 u + \mathbf{e}$$

with $\mathbf{e}$ a random variable with a particular pdf $f_{\mathbf{e}}$, and θ_0 an unknown scalar constant.
We are going to estimate a model of the form

$$\mathbf{y} = \theta u + \mathbf{e}$$

and the pdf of the observed random variable $\mathbf{y}$ is then given by $f_{\mathbf{e}}(y - \theta u)$.
For one observed measurement of $\mathbf{y}$, i.e., $\mathbf{y}$ is a scalar, the pdf of $\mathbf{y}$ as a function of θ is depicted in Figure 5.9 for $u = 2$. It is a continuum of pdf's, since θ varies over a continuous region.
Now, if one observation of $\mathbf{y}$ is made, e.g., $y = 0$, then the likelihood function results: $L(\theta) = f_{\mathbf{e}}(y - \theta u)|_{y=0}$, being a function of θ only, and depicted in Figure 5.9 as the solid (red) curve. The maximum likelihood estimate of θ is that value of θ for which L reaches its maximum.
In Figure 5.9 the example is sketched of a normal pdf $f_{\mathbf{e}}$ with zero mean and unit variance. The likelihood function then becomes

$$L(\theta; y) = \frac{1}{\sqrt{2\pi}} e^{\frac{-(y-2\theta)^2}{2}}.$$

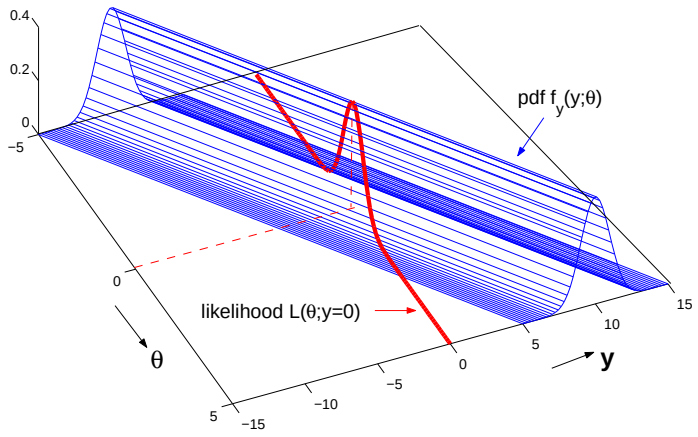

Figure 5.9: A prior probability density function of $\mathbf{y}$, as a function of θ, (blue continuum of curves), and the likelihood function L for observation $y = 0$ (red solid curve).

For the observation $y = 0$, maximization of $L(\theta; y = 0)$ comes down to maximizing

$$\frac{1}{\sqrt{2\pi}} e^{\frac{-4\theta^2}{2}}$$

which is obtained for $\theta = 0$, this being the maximum likelihood estimate.

Maximum likelihood estimation of the parameters of linear regression models from Gaussian distributed observations

The considered -very simple- example can straightforwardly be extended to the situation of multiple observations and vector parameters. Consider the previously used linear regression model:

$$\mathbf{y}_i = \phi_i^T \theta + \mathbf{e}_i \qquad i = 1, \cdots n$$

with

$$\phi_i = \begin{bmatrix} 1 \\ u_i \end{bmatrix} \quad \text{and } \theta = \begin{bmatrix} b_0 \\ b_1 \end{bmatrix}.$$

and assume that $\mathbf{e}_i$ is a set of independent Gaussian random variables with pdf $f_{\mathbf{e}}$ having mean value 0 and variance σ^2.
Then the joint probability density function of the measured variables $\mathbf{y}_1 \cdots \mathbf{y}_n$ is given by

$$f_{\mathbf{y}}(y) = \prod_{i=1}^{n} f_{\mathbf{e}}(y_i - \phi_i^T \theta)$$

and the likelihood function then becomes

$$L(\theta; Y) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y_i - \phi_i^T \theta)^2}{2\sigma^2}}.$$

Maximizing the likelihood function over θ leads to the same argument as maximizing $\log L(\theta, Y)$ over θ. This is due to the fact that the log is a monotone increasing function. Taking the logarithm is often advantageous, in particular for probability density functions that contain exponentials. Instead of maximizing $\log L$ one can equivalently minimize $-\log L$, which is given by

$$-\log L(\theta; Y) = \frac{n}{2}\log 2\pi + n\log\sigma + \frac{1}{2\sigma^2}\sum_{i=1}^{n}(\mathbf{y}_i - \phi_i^T\theta)^2. \tag{5.88}$$

Since the first two terms on the right hand side of this expression are constants and not functions of θ, consequently

$$\hat{\boldsymbol{\theta}}_n = \arg\min_\theta \frac{1}{2\sigma^2}\sum_{i=1}^{n}(\mathbf{y}_i - \phi_i^T\theta)^2 \tag{5.89}$$

$$= \arg\min_\theta \sum_{i=1}^{n}(\mathbf{y}_i - \phi_i^T\theta)^2 \tag{5.90}$$

and this expression is exactly the same as the simple least squares (linear regression) estimator that was developed in section 5.3. This leads to the following conclusion.

> For n independent observations from a Gaussian distribution with equal variance for all observations, the ML estimator is given by the simple least squares (LS) estimator.

If in the considered problem the noise terms $\mathbf{e}_i$ are independent Gaussian random variables with zero mean and with fixed and known variance σ_i^2, being different for the different measurements i, then the likelihood function simply generalizes to

$$L(\theta; Y) = \prod_{i=1}^{n}\frac{1}{\sqrt{2\pi}\sigma_i}e^{-\frac{(y_i-\phi_i^T\theta)^2}{2\sigma_i^2}}$$

and the corresponding maximum likelihood estimator is given by

$$\hat{\boldsymbol{\theta}}_n = \arg\min_\theta \sum_{i=1}^{n}\frac{(\mathbf{y}_i - \phi_i^T\theta)^2}{\sigma_i^2} \tag{5.91}$$

and this expression is the same as the weighted least squares estimator with a diagonal weighting matrix $W = diag(1/\sigma_i^2, \cdots, \sigma_n^2)$, which was developed in section 5.3.5. If in the considered problem the noise terms $\mathbf{e}_i$ are *correlated* Gaussian random variables with zero mean and $n \times n$ covariance matrix Σ, then the likelihood function generalizes to

$$L(\theta; Y) = \frac{1}{(2\pi)^{n/2}\sqrt{det(\Sigma)}}e^{-\frac{1}{2}(Y-X\theta)^T\Sigma^{-1}(Y-X\theta)}$$

with

$$X = \begin{bmatrix} \phi_1^T \\ \vdots \\ \phi_n^T \end{bmatrix}; \quad Y = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

and the corresponding maximum likelihood estimator is given by

$$\hat{\boldsymbol{\theta}}_n = \arg\min_{\theta}(\mathbf{Y} - X\theta)^T\Sigma^{-1}(\mathbf{Y} - X\theta), \tag{5.92}$$

which leads to

$$\hat{\boldsymbol{\theta}}_n = (X^T\Sigma^{-1}X)^{-1}X^T\Sigma^{-1}\mathbf{Y} \tag{5.93}$$

and this expression is the same as the weighted least squares estimator that was developed in section 5.3.5. As a result one can state the following:

> For n observations from a joint Gaussian distribution, the ML estimator is given by the weighted least squares (WLS) estimator, where the weighting matrix is given by the inverse of the covariance matrix of the observations.

Note that the considered weighted least squares estimator is identical to the BLUE (best linear unbiased estimator), discussed in section 5.3.5. So, for Gaussian disturbances (and regression models linear in the unknown parameters θ) , the BLUE and the ML estimator coincide, and from the results of the previous section it follows that their variance reaches the Cramér-Rao lower bound, i.e., the minimum possible variance over all unbiased estimators. Note that the equivalence between ML and LS/WLS estimators typically holds true for Gaussian distribution functions. However the maximum likelihood principle goes beyond this situation and applies also to other distributions.

Discussion

In the general case, the ML estimator will require solving the optimization problem

$$\hat{\boldsymbol{\theta}}_n = \arg\min_{\theta}(-\log L(\theta, \mathbf{Y}))$$

which might not be straightforward at all. In particular for Gaussian distributions, and a linear regression model, this optimization problem is simple, since it reduces to minimizing a quadratic function in θ. This convex optimization problem is simply solvable by efficient algorithms, relying on analytical expressions for its solution. In the general situation however the optimization problem will require nonlinear (gradient-type) optimization tools that suffer from possible poor convergence due to the existence of local minima, and the resulting lack of guarantee that a solution to the problem has been obtained. For an overview of optimization methods see , e.g., Fletcher (1980) and Miller (2000).

Properties of the ML estimator
The maximum likelihood estimator has several important properties:

> The ML-estimator has the property that for number of the observations n tending to infinity,
>
> $$\hat{\boldsymbol{\theta}}_n \rightarrow \mathcal{N}(\theta_0, J^{-1})$$
>
> meaning that the random variable $\hat{\boldsymbol{\theta}}_n$ converges in distribution to a Gaussian distribution with mean value θ_0 and covariance matrix equal to the Cramér-Rao bound.
> As a result of this, the ML-estimator is
>
> - asymptotically unbiased;
> - consistent
> - asymptotically efficient, i.e., it asymptotically reaches the minimum possible variance (CRLB) among all unbiased estimators.

These properties of the ML estimator are very powerful, and are a strong support for this estimation principle. With respect to the practical implications one should realize though that in order for an ML estimator to be applicable, detailed knowledge of the probability density function of the observations has to be available (as prior information, e.g., derived for physical considerations/analysis of the disturbances that are acting on the measurement data). Note also that the attractive properties of the ML estimator all hold asymptotically in the number of data n. There is no guarantee for unbiased and minimum variance estimators in the case of finite data.
This section on the ML estimator will be ended by providing some examples.

Example 5.9 (Estimation of mean value of Gaussian random variable) Consider n independent observations of a Gaussian random variable $\mathbf{y}$ with unknown mean μ and known variance σ^2. The likelihood function, as a function of μ is given by

$$L(\mu) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(\mathbf{y}_i - \mu)^2}{2\sigma^2}} \tag{5.94}$$

and

$$\log L(\mu) = n \log \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \sum_{i=1}^{n} (\mathbf{y}_i - \mu)^2. \tag{5.95}$$

Maximizing L leads to the same argument as minimizing $-\log L$, and so the maximum likelihood estimator for μ is determined by

$$\hat{\boldsymbol{\mu}}_n = \arg\min_{\mu} \left\{ -n \log \frac{1}{\sqrt{2\pi}\sigma} + \frac{1}{2\sigma^2} \sum_{i=1}^{n} (\mathbf{y}_i - \mu)^2 \right\} \tag{5.96}$$

$$= \arg\min_{\mu} \frac{1}{2\sigma^2} \sum_{i=1}^{n} (\mathbf{y}_i - \mu)^2. \tag{5.97}$$

Setting the derivative of the function in the right hand side to zero:

$$\left[-2\sum_{i=1}^{n}\mathbf{y}_i - \mu\right]_{\mu=\hat{\boldsymbol{\mu}}_n} = 0 \tag{5.98}$$

delivers $n\hat{\boldsymbol{\mu}}_n = \sum_{i=1}^{n}\mathbf{y}_i$ or equivalently

$$\hat{\boldsymbol{\mu}}_n = \frac{1}{n}\sum_{i=1}^{n}\mathbf{y}_i. \tag{5.99}$$

In other words, in the considered situation is the sample average the maximum likelihood estimator for the mean value of the random variable.

Example 5.10 (Estimation of variance of Gaussian random variable) Consider n independent observations of a Gaussian random variable $\mathbf{y}$ with known mean μ and unknown variance σ^2. The likelihood function, as a function of σ is given by

$$L(\sigma) = \prod_{i=1}^{n}\frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(\mathbf{y}_i-\mu)^2}{2\sigma^2}}$$

Since maximizing $L(\sigma)$ has the same optimum argument as minimizing $-\log L(\sigma)$, we consider

$$-\log L(\sigma) = \frac{1}{2\sigma^2}\sum_{i=1}^{n}(\mathbf{y}_i-\mu)^2 - n\log\frac{1}{\sqrt{2\pi}\sigma}$$

Minimizing $-\log L(\sigma)$ can be done by setting $\frac{\partial}{\partial\sigma} = 0$, i.e.,

$$\left[\frac{n}{\sigma} - 2\sum_{i=1}^{n}\frac{(\mathbf{y}_i-\mu)^2}{2\sigma^3}\right]_{\sigma=\hat{\boldsymbol{\sigma}}_n} = 0$$

or equivalently

$$n - \sum_{i=1}^{n}\frac{(\mathbf{y}_i-\mu)^2}{\hat{\boldsymbol{\sigma}}_n^2} = 0$$

from which follows that

$$\hat{\boldsymbol{\sigma}}_n^2 = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{y}_i-\mu)^2. \tag{5.100}$$

The difference of this estimator with analyzed in Example 5.2 is that here the exact expression μ is used, whereas in the earlier example an estimate of μ was substituted in the expression. The current estimator is unbiased, whereas the estimator in (5.15) is biased.

Example 5.11 (Estimating arrival rate of photons on a detection plate) The time T between the detection of two photons on a detection plate of an electron microscope is a random variable with exponential density function

$$f_{\mathbf{T}}(T) = \begin{cases} \alpha e^{-\alpha T} & T \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{5.101}$$

The coefficient α is known as the arrival rate. Additionally interarrival intervals are independent (see also the Poisson process discussed in section 3.4.3). For the estimation of the arrival rate α, we construct the likelihood function:

$$L(\alpha; T_1, \cdots, T_n) = \prod_{i=1}^{n} \alpha e^{-\alpha T_i} = \alpha^n e^{-\alpha \sum_{i=1}^{n} T_i}. \tag{5.102}$$

The maximum likelihood estimate is found by taking the logarithm of this function, and setting the derivative equal to 0:

$$\frac{\partial}{\partial \alpha}\left[n \log \alpha - \alpha \sum_{i=1}^{n} T_i\right] = \frac{n}{\alpha} - \sum_{i=1}^{n} T_i = 0. \tag{5.103}$$

Solving this equation for α leads to the maximum likelihood estimate

$$\hat{\alpha}_{ml} = \frac{1}{\frac{1}{n}\sum_{i=1}^{n} T_i}. \tag{5.104}$$

The estimate is the inverse of the average interarrival time.

In the examples given so far, the maximum likelihood estimate can always be obtained by simple analytic expressions. This is however not a structural property of this estimator. Consider, e.g., the emission of radioactive particles from two radioactive sources, generally modelled as

$$y_t = \lambda_1 e^{\mu_1 t} + \lambda_2 e^{\mu_2 t} + \eta + e_t$$

with λ_1, λ_2 the concentration of the two radioactive sources, μ_1, μ_2 the decay rate of the separate sources, η the effect of background radiation, and e_t a random error, and that the radiation y_t is observed over n time instants.
Suppose that e_t are independent Gaussian random variables with known and fixed mean and variance.
If μ_1 and μ_2 are known, and we intend to estimate λ_1, λ_2 and η, we can write

$$y_t = \phi_t^T \theta + e_t \tag{5.105}$$

with

$$\phi_t = [e^{\mu_1 t} \quad e^{\mu_2 t} \quad 1]^T \text{and } \theta = [\lambda_1 \ \lambda_2 \ \eta]^T$$

and with e_t having a Gaussian distribution, the ML-estimator of θ is then simply obtained as the linear least squares estimate. This is a direct consequence of the fact that the observations can be written in the linear regression form (5.105) which is linear in the unknown parameter θ.
If all variables are unknown the likelihood function to be optimized becomes:

$$L = \prod_{t=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} exp(-\frac{(y_t - \lambda_1 e^{\mu_1 t} - \lambda_2 e^{\mu_2 t} - \eta)^2}{2\sigma^2}). \tag{5.106}$$

Maximizing this function is equivalent to minimizing the function

$$-\log L = c + \frac{1}{2\sigma^2}\sum_{t=1}^{n}(y_t - \lambda_1 e^{\mu_1 t} - \lambda_2 e^{\mu_2 t} - \eta)^2. \tag{5.107}$$

Minimizing this function as a function of the unknown parameters $\lambda_1, \lambda_2, \eta, \mu_1$ and μ_2 is an optimization problem that cannot simply be solved analytically by setting the partial derivatives equal to 0. Finding the ML estimate is now equivalent to solving a nonlinear least squares problem. More complex nonlinear optimization algorithms, as , e.g., gradient type methods, are necessary to computationally find the optimal parameter value.

5.7 Parameter estimation and random variable estimation

So far we have considered estimation problems where the unknown parameters to be estimated are unknown constants, such as , e.g., the mean or variance of a random variable, but also the autocorrelation function and the spectral density function of a stationary stochastic process belong to this category. Estimators for the latter notions will be discussed in chapter 7. Also unknown (but fixed) physical parameters in physical models can be estimated with the methods presented in the current chapter.
A different situation occurs if the object that one intends to estimate itself is a random variable. This typically occurs in problems of prediction and filtering:

- Consider observations of a stationary stochastic process $\mathbf{y}(t)$ over a time sequence $t = 1, \cdots N-1$; determine an estimate of the future values of the process, $\mathbf{y}(t), t \geq N$. E.g., predict future values of the stock exchange index on the basis of past observations (prediction problem).

- Estimate the realization of a random variable that is not directly measurable, on the basis of related and measurable random variables.
  E.g., predict the weather forecast conditions on a fine grid over Europe, on the basis of measurements at a limited number of locations, together with an atmospheric model (filtering problem).

Random variable estimation will be considered in the next chapter.

5.8 Summary

In this chapter the basic theory for characterizing estimators has been presented. Principle estimation methods, such as least squares/linear regression have been analyzed, and have been shown to be easily extendable to more general estimators. This includes weighted least squares, total least squares, and finally maximum likelihood methods. While going to more advanced estimation methods an increasing amount of prior information on the statistics of the measurement data is required. Linear least squares methods are directed towards measurement data that are contaminated by zero-mean random variables. If the variance of the measurement variables is known, better estimators (with smaller variance) can be obtained by using a weighted least squares method. If the full *pdf* of the measurement data is known, the maximum likelihood method provides an estimator which asymptotically reaches the smallest variance possible, determined by the Cramér-Rao bound.

Chapter 6

Optimal filtering

The problem is addressed how to estimate an information carrying source signal, that is represented as a single realization of a stochastic process, on the basis of observations that are corrupted by stochastic noise. Optimal filter theory provides filters for extracting the source signal on the basis of the essential properties of the stochastic processes involved. The theory includes methods for estimating physical variables that are not directly measurable, but for which indirect measurements are available.

6.1 Introduction

The estimation of an unobservable information carrying signal from a measurable -possibly noise corrupted- signal is one of the central problems is signal processing. It refers to many practical examples as, e.g., in speech, radar, EEG, imaging, where the direct information carrying part of the signal can not be measured free of noise. This noise effect may be due to inaccuracies or limited resolution in the applied sensors, but the contamination of the signal can also be due to the underlying physical phenomena as , e.g., the propagation of the signal from the source to the receiver (as the channel characteristic in a digital communication system), or the propagation of an optical wave through a lens system.
Whereas in the previous chapter we have been considering the problem of estimating an unknown physical (but constant) parameter θ_0 on the basis of measurable random variables, the problems considered in this chapter concern the estimation of quantities that are random variables themselves. In other words: the quantities to be estimated are realizations of random variables, and estimators are sought that provide good estimates for all possible realizations of the random variable to be estimated.

> Whereas parameter estimation refers to the hitting of a fixed target on the basis of noisy observations, random variable estimation (or optimal filtering) refers to the hitting of a (randomly) moving target on the basis of noisy observations.

E.g., suppose that there exists a random variable $\mathbf{x}$ that can not be observed directly, and a related set of random variables $\mathbf{y}_1 \cdots \mathbf{y}_n$ that can be measured. For example, $\mathbf{x}$ may be the value of a random signal at a particular time, but the signal is embedded in a hostile environment and cannot be received without significant distortion. The observations may

represent various receptions of the signal corrupted by noise, dispersion, multipath, or other interfering signals. The goal is now to find an estimator

$$\hat{\mathbf{x}} = \hat{x}(\mathbf{y}_1, \cdots, \mathbf{y}_n)$$

that is optimal in some sense.
In time-sequenced signals, or random processes, the problems can be phrased as follows: Suppose an information carrying random process $\{\mathbf{x}(t)\}$ is not directly measurable, but its information is hidden in a measurable random process $\{\mathbf{y}(t)\}$ that is related / correlated to the process $\{\mathbf{x}(t)\}$. In terms of the global pictures presented in Chapter 1, this refers to the situation as sketched in Figure 6.1.

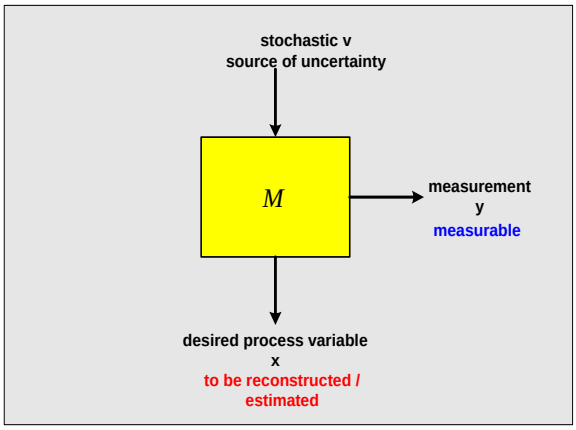

Figure 6.1: Filter problem

In Figure 6.2 an example of signals $\{x(t)\}$ and $\{y(t)\}$ is given. For given signal values $y(t)$, $t \leq t_0$, the optimal estimation of $x(\tau)$ is generally referred to as a
smoothing problem if $\tau < t_0$
filtering problem if $\tau = t_0$, and
prediction problem if $\tau > t_0$.
The presumed relation between $\mathbf{y}$ and $\mathbf{x}$ can be structured in different forms. The most simple situation is sketched in Figure 6.3(left). Here the observed process $\mathbf{y}$ is a noise disturbed version of the process $\mathbf{x}$, with $\mathbf{v}$ an additive noise disturbance, e.g. reflecting the finite precision of the measurement device.
If the measurement device does not only add a noise term, but also statically or dynamically influences the measurement, the situation as sketched in Figure 6.3(right) applies. Here the transfer function G refers to, e.g., the dynamical properties of the sensing device (finite bandwidth), or the physical relationship that exists between the process of interest $\mathbf{x}$ and the measurable process $\mathbf{y}$.
Consider, e.g., a piezo-element that is used to measure a displacement by way of the voltage across the element. In this case the transfer G refers to the relationship between voltage and displacement of the considered piezo-element. In a communication channel, $\mathbf{x}$ can be the transmitted signal, G the properties of the channel, and $\mathbf{y}$ the received signal. In an image analysis/processing problem, G can be the effect of blurring.
In order to solve the filtering problem, several parts of prior information are required. This will in particular include

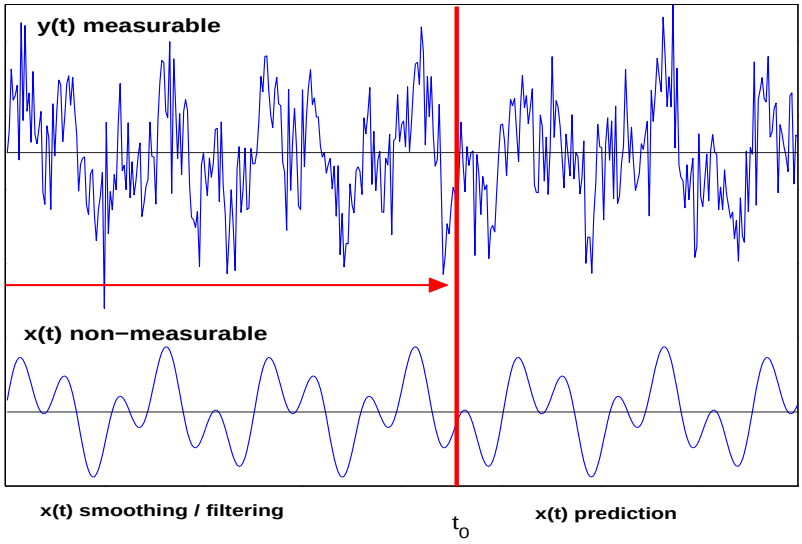

Figure 6.2: Smoothing, filtering and prediction

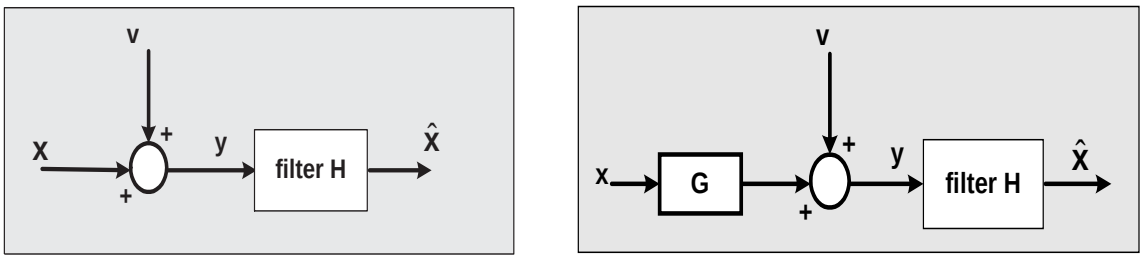

Figure 6.3: Simple relationships between measurable process $\mathbf{y}$ and to be reconstructed process $\mathbf{x}$; $\mathbf{v}$ is a stochastic disturbance term.

- Knowledge of the relationship between $\mathbf{x}$ and $\mathbf{y}$, and particularly
- Knowledge of G and knowledge of the second order properties of $\mathbf{x}$ and $\mathbf{v}$, i.e., their means and covariance functions (or spectral densities).

Of course in general this required information will not be known a priori, and will have to be estimated from measurements also. This will be the subject of later chapters. Here it is assumed that the information mentioned above is available.
In the situation of (time-sequenced) stochastic processes, the optimal filtering problem may seem to have a trivial solution. If the spectral density functions (frequency bands) of the information carrying process $\mathbf{x}$ and the disturbance process $\mathbf{v}$ are disjunct, as sketched in Figure 6.4, one might apply a simple low-pass filter H to reconstruct $\mathbf{x}$ from the measurement $\mathbf{y} = \mathbf{x} + \mathbf{v}$, and thus separate the two different signals. The resulting (optimal) filter operation is sketched in Figure 6.5. In this case $\hat{\mathbf{x}}(t)$ will be a perfect estimate of $\mathbf{x}(t)$ for any realization of $\mathbf{x}$ that is considered. However in most applications such a clear separation of frequency range will not appear, and more advanced techniques are required to optimally estimate the information carrying signal from the available measurements.
In this chapter the general random variable estimation problem will be considered:

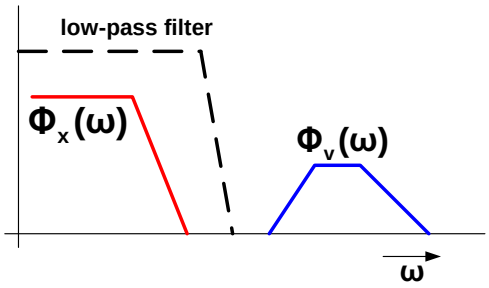

Figure 6.4: Information carrying process $\mathbf{x}$ and noise $\mathbf{v}$ with disjunct power spectral densities

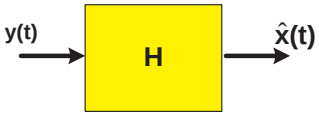

Figure 6.5: Filtering for separation of signals

- Construct an estimator $\hat{\mathbf{x}}$ of the random variable $\mathbf{x}$, on the basis of n measured related random variables $\mathbf{y}_1, \cdots, \mathbf{y}_n$.

This problem is tackled by considering so-called mean squared error (MSE) estimation. Subsequently, this problem is transferred to the situation of time-sequenced random processes, and related optimal filtering problems are considered. In this latter situation the measured random variables are considered to be time-sequenced random variables in a stationary stochastic process.

6.2 Random variable estimation and the mean squared error (MSE) criterion

In the random variable estimation problem we first consider n random variables to be available from measurements:

$$\mathbf{y}_1, \cdots \mathbf{y}_n$$

and the problem is to construct an estimator

$$\hat{\mathbf{x}} = g(\mathbf{y}_1, \cdots, \mathbf{y}_n)$$

to obtain a "best" estimator of the random variable $\mathbf{x}$ that somehow is related to the random variables $\mathbf{y}_i$. So random variable $\mathbf{x}$ is the target to be reached, and $\mathbf{y}_i$ are the random measurements that are used to estimate the target.
As a criterion of fit, many different cost functions are possible. The most common criterion is the mean squared error criterion

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] \tag{6.1}$$

and an estimator $\hat{\mathbf{x}} = g(\mathbf{y}_1, \cdots, \mathbf{y}_n)$ is sought for that minimizes this criterion.
In contrast with the use of the MSE of a parameter estimator, the MSE criterion (6.1) concerns *two* random variables, i.e., $\mathbf{x}$ (the target is random) and $\mathbf{y} = [\mathbf{y}_1, \cdots, \mathbf{y}_n]^T$ (the

measurements are random). An appropriate notation for this MSE criterion is therefore:

$$\mathbb{E}_{\mathbf{x},\mathbf{y}}[(\mathbf{x} - \hat{x}(\mathbf{y}))^2].$$

Consequently the expected value in (6.1) is taken over the joint probability density function $f_{\mathbf{x},\mathbf{y}}(x, y)$, i.e.,

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \mathbb{E}[(\mathbf{x} - \hat{x}(\mathbf{y}))^2] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(x - \hat{x}(y))^2 f_{\mathbf{x},\mathbf{y}}(x, y)dxdy \tag{6.2}$$

where of course the integration over y should be considered to be taken over an n-dimensional measurement space.
By using Bayes' relation (A.5) for joint probability density functions:

$$f_{\mathbf{x},\mathbf{y}}(x, y) = f_{\mathbf{x}|\mathbf{y}}(x|y) \cdot f_{\mathbf{y}}(y)$$

the MSE criterion becomes

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \int_{-\infty}^{\infty} J(\hat{x}, y) f_{\mathbf{y}}(y)dy \tag{6.3}$$

with

$$J(\hat{x}, y) = \int_{-\infty}^{\infty}(x - \hat{x})^2 f_{\mathbf{x}|\mathbf{y}}(x|y)dx. \tag{6.4}$$

A minimum for (6.3) can be obtained if we can find an estimate $\hat{x}$ that is minimal for all y. A necessary condition for $J(\hat{x}, y)$ to be minimal for all y, is that

$$\frac{\partial}{\partial \hat{x}} J(\hat{x}, y) = -2\int_{-\infty}^{\infty}(x - \hat{x}) f_{\mathbf{x}|\mathbf{y}}(x|y)dx = 0, \tag{6.5}$$

or equivalently

$$\int_{-\infty}^{\infty} x f_{\mathbf{x}|\mathbf{y}}(x|y)dx = \hat{x}\int_{-\infty}^{\infty} f_{\mathbf{x}|\mathbf{y}}(x|y)dx.$$

Since the last integral of the density function is 1, this implies that the estimator that minimizes the MSE criterion, is given by

$$\hat{x}(\mathbf{y}) = \int_{-\infty}^{\infty} x f_{\mathbf{x}|\mathbf{y}}(x|\mathbf{y})dx = \mathbb{E}(\mathbf{x}|\mathbf{y}). \tag{6.6}$$

This result leads to the observation that

> The mean squared error (MSE) estimator of a random variable $\mathbf{x}$ on the basis of observed random variables $\mathbf{y}$ is given by its conditional mean $\mathbb{E}[\mathbf{x}|\mathbf{y}]$.

Actually a complete characterization of the random variable $\mathbf{x}$ when having knowledge of $\mathbf{y}$ is present in the conditional probability density function $f_{\mathbf{x}|\mathbf{y}}(x|y)$. However in order to construct this conditional pdf knowledge is required of both $f_{\mathbf{x},\mathbf{y}}(x, y)$ and $f_{\mathbf{y}}(y)$. While the latter pdf might be estimated on the basis of measurement data, the first pdf requires detailed knowledge of $\mathbf{x}$. The result of the derivation in this section is that the mean value of $f_{\mathbf{x}|\mathbf{y}}(x|y)$ can be found by minimizing $\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2]$, while no prior knowledge is required on the form of pdf that is considered; the result holds for any pdf.

6.3 Linear MSE estimators - the scalar case

For illustration we consider a simple example that will be approached from two different sides.

Conditional mean for Gaussian distributions

Consider the random variables $\mathbf{x}$ and $\mathbf{y}$ to have a joint Gaussian distribution, with mean value

$$\mu = \begin{pmatrix} \mu_{\mathbf{x}} \\ \mu_{\mathbf{y}} \end{pmatrix}$$

and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_{\mathbf{x}}^2 & \sigma_{\mathbf{xy}} \\ \sigma_{\mathbf{xy}} & \sigma_{\mathbf{y}}^2 \end{pmatrix}$$

and the joint probability density function is given by

$$f_{\mathbf{x},\mathbf{y}}(x,y) = \frac{1}{2\pi\sqrt{det(\Sigma)}} e^{-\frac{1}{2}\begin{bmatrix} x-\mu_{\mathbf{x}} \\ y-\mu_{\mathbf{y}} \end{bmatrix}^T \Sigma^{-1} \begin{bmatrix} x-\mu_{\mathbf{x}} \\ y-\mu_{\mathbf{y}} \end{bmatrix}}$$

while the marginal density function for observed random variable $\mathbf{y}$ is given by

$$f_{\mathbf{y}}(y) = \frac{1}{\sqrt{2\pi}\sigma_{\mathbf{y}}} e^{-\frac{(y-\mu_{\mathbf{y}})^2}{2\sigma_{\mathbf{y}}^2}}.$$

Then by forming the ratio

$$f_{\mathbf{x}|\mathbf{y}}(x|y) = \frac{f_{\mathbf{x},\mathbf{y}}(x,y)}{f_{\mathbf{y}}(y)}$$

and by some tedious manipulation, it can be shown that

$$f_{\mathbf{x}|\mathbf{y}}(x|y) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-m(y))^2/2\sigma^2} \tag{6.7}$$

with

$$m(y) = \mu_{\mathbf{x}} + \sigma_{\mathbf{xy}}\sigma_{\mathbf{y}}^{-2}(y-\mu_{\mathbf{y}}) \tag{6.8}$$

$$\sigma^2 = \sigma_{\mathbf{x}}^2 - \sigma_{\mathbf{xy}}\sigma_{\mathbf{y}}^{-2}\sigma_{\mathbf{xy}}. \tag{6.9}$$

The expressions (6.8) and (6.9) can directly be derived from the more complex multivariate situation that is explained in detail in Appendix C. Alternative formulations are possible employing the correlation coefficient $\rho_{\mathbf{xy}} = \sigma_{\mathbf{xy}}/(\sigma_{\mathbf{x}}\sigma_{\mathbf{y}})$, leading to, e.g., $\sigma^2 = \sigma_{\mathbf{x}}^2(1-\rho_{\mathbf{xy}}^2)$. As a result the conditional mean of $\mathbf{x}$ is given by

$$\hat{x}(y) = m(y) = \mu_{\mathbf{x}} + \sigma_{\mathbf{xy}}\sigma_{\mathbf{y}}^{-2}(y-\mu_{\mathbf{y}}). \tag{6.10}$$

An important observation here is that the estimate depends linearly[1] on y. The optimal MSE estimator therefore is a linear estimator, and this is a direct result of the fact that the joint distribution function is Gaussian.

[1] Actually the function $m(y)$ is not linear but affine, meaning that it is of the form $m(y) = ay + b$. For linearity it needs to satisfy $m(\alpha y_1 + \beta y_2) = \alpha m(y_1) + \beta m(y_2)$, for all α and β and this only holds true when $b = 0$. Nevertheless the y-dependent term appears linearly in $m(y)$.

For Gaussian distributed random variables, the optimal MSE estimator is a linear function of the observed variables.

Note that for calculation of the optimal estimator, information is required on the joint distribution function of $\mathbf{x}$ and $\mathbf{y}$. This is of course a rather restrictive requirement.

Optimal MSE linear estimator

Consider again the scalar random variables $\mathbf{x}$ and $\mathbf{y}$, with $\mathbf{y}$ being observable and $\mathbf{x}$ to be estimated. Now we do not assume to have knowledge of the probability density function but we simply look for the best MSE estimator of the form:

$$\hat{\mathbf{x}} = a + b\mathbf{y} \tag{6.11}$$

where a and b are constants that determine the MSE estimator. The linear structure of the estimator (6.11) now is chosen by the experimenter, and one is looking for the best linear estimator that has this structure.
When minimizing

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2]$$

with $\hat{\mathbf{x}} = a + b\mathbf{y}$, the solution a_{ms}, b_{ms} is obtained by setting the derivative of the criterion equal to zero, i.e.,

$$\begin{pmatrix} \frac{\partial}{\partial a} \\ \frac{\partial}{\partial b} \end{pmatrix} \mathbb{E}[(\mathbf{x} - a - b\mathbf{y})^2] = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad \text{in } (a, b) = (a_{ms}, b_{ms}).$$

As the criterion is quadratic in the unknowns (a, b), this indeed leads to the global minimum of the considered MSE criterion.
Analyzing the above expression shows:

$$\left.\begin{array}{r} -2\mathbb{E}[\mathbf{x} - a_{ms} - b_{ms}\mathbf{y}] = 0 \\ -2\mathbb{E}[(\mathbf{x} - a_{ms} - b_{ms}\mathbf{y})\mathbf{y}] = 0 \end{array}\right\} \Rightarrow \left\{\begin{array}{l} a_{ms} = \mu_{\mathbf{x}} - b_{ms}\mu_{\mathbf{y}} \\ \mathbb{E}[\mathbf{xy}] - a_{ms}\mu_{\mathbf{y}} - b_{ms}\mathbb{E}[\mathbf{y}^2] = 0 \end{array}\right. \tag{6.12}$$

Substitution of the first equation in the second one provides:

$$\mathbb{E}[\mathbf{xy}] - (\mu_{\mathbf{x}} - b_{ms}\mu_{\mathbf{y}})\mu_{\mathbf{y}} - b_{ms}\mathbb{E}[\mathbf{yy}] = 0.$$

Utilizing the definitions for (co)variance:

$$\begin{aligned} \sigma_{\mathbf{y}}^2 &= \mathbb{E}[(\mathbf{y} - \mu_{\mathbf{y}})^2] = \mathbb{E}[\mathbf{y}^2] - \mu_{\mathbf{y}}^2 \\ \sigma_{\mathbf{xy}} &= \mathbb{E}[(\mathbf{x} - \mu_{\mathbf{x}})(\mathbf{y} - \mu_{\mathbf{y}})] = \mathbb{E}[\mathbf{xy}] - \mu_{\mathbf{x}}\mu_{\mathbf{y}} \end{aligned}$$

it follows that $\sigma_{\mathbf{xy}} - b_{ms}\sigma_{\mathbf{y}}^2 = 0$, leading to

$$a_{ms} = \mu_{\mathbf{x}} - \mu_{\mathbf{y}}\frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2} \qquad b_{ms} = \frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2}. \tag{6.13}$$

The linear MSE estimator is obtained by $\hat{\mathbf{x}} = a_{ms} + b_{ms}\mathbf{y}$, which equals

$$\hat{\mathbf{x}} = \mu_{\mathbf{x}} + \frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2}(\mathbf{y} - \mu_{\mathbf{y}}) \tag{6.14}$$

and this optimal estimator is exactly the same as the one that was obtained in (6.10).

Remarks

- Note that in the second situation linearity of the estimator is the result of a choice made by the experimenter, while in the first situation it is a result of the Gaussian probability density function. In the second situation there might exist other estimators (e.g., more complex linear or nonlinear functions $g(\mathbf{y})$) that might lead to smaller values of the MSE criterion.

> The best *linear* MSE estimator is equal to the best MSE estimator under Gaussian assumptions. For other distributions nonlinear estimators might lead to smaller MSE.

- A second remark that can be made is that the procedure for deriving the optimal MSE coefficients of the linear estimators, is very similar to the linear regression techniques analyzed in section 5.3. Indeed the results match very closely in terms of their structure, although the different components in the expressions are slightly different: data-based averages in the linear regression problems, versus covariances and mean values in the linear estimator problems considered in this chapter. Whereas in the linear regression problems the parameter estimators are random variables (with statistical properties), the filter coefficients a_{ms}, b_{ms} derived here are no random variables, but optimal filter coefficients.

In the considered situations, the minimum value of the MSE criterion can be calculated as follows. By substituting the coefficient expressions (6.14) into the estimator $\hat{\mathbf{x}} = a_{ms} + b_{ms}\mathbf{y}$, and calculating

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \mathbb{E}[(\mathbf{x} - \mu_{\mathbf{x}} + \mu_{\mathbf{y}}\frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2} - \frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2}\mathbf{y})^2] \tag{6.15}$$

it follows that

$$\begin{aligned}
\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] &= \mathbb{E}[\left\{(\mathbf{x} - \mu_{\mathbf{x}}) - \frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2}(y - \mu_{\mathbf{y}})\right\}^2] && (6.16)\\
&= \sigma_{\mathbf{x}}^2 + \frac{\sigma_{\mathbf{xy}}^2}{\sigma_{\mathbf{y}}^4}\sigma_{\mathbf{y}}^2 - 2\frac{\sigma_{\mathbf{xy}}^2}{\sigma_{\mathbf{y}}^2} && (6.17)\\
&= \sigma_{\mathbf{x}}^2 - \frac{\sigma_{\mathbf{xy}}^2}{\sigma_{\mathbf{y}}^2} = \sigma_{\mathbf{x}}^2(1 - \rho_{\mathbf{xy}}^2). && (6.18)
\end{aligned}$$

- If $\mathbf{y}$ and $\mathbf{x}$ are strongly correlated, $\rho_{\mathbf{xy}} \sim 1$, and the MSE criterion is close to 0. In this situation $\mathbf{x}$ can very accurately be estimated on the basis of $\mathbf{y}$.

- If $\mathbf{y}$ and $\mathbf{x}$ are only weakly correlated, $\rho_{\mathbf{xy}} \sim 0$, and the MSE criterion is close to the variance of $\mathbf{x}$. Observing $\mathbf{y}$ does not contribute essentially to variance reduction caused by the fact that $\mathbf{y}$ does provide only very limited information on $\mathbf{x}$.

We conclude this section with two -rather simple- examples.

Example 6.1 (MSE Estimation of a random variable with a constant) Let a be a -constant- parameter to estimate a random variable $\mathbf{x}$. Which value of a will lead to an estimate with the smallest MSE?
In the setting of this section, the estimator has the format

$$\hat{\mathbf{x}} = a$$

and so it is not a function of any observable random variable.
By writing

$$\begin{aligned}\mathbb{E}[(\mathbf{x}-a)^2] &= \mathbb{E}[\{(\mathbf{x}-\mu_{\mathbf{x}})+(\mu_{\mathbf{x}}-a)\}^2] \\ &= \mathbb{E}[(\mathbf{x}-\mu_{\mathbf{x}})^2]+\mathbb{E}[(\mu_{\mathbf{x}}-a)^2]+ \\ &\quad +2(\mu_{\mathbf{x}}-a)\underbrace{\mathbb{E}[\mathbf{x}-\mu_{\mathbf{x}}]}_{=0} \\ &= \sigma_{\mathbf{x}}^2+(\mu_{\mathbf{x}}-a)^2\end{aligned}$$

it follows that the MSE is minimal for $a = \mu_{\mathbf{x}}$. This shows that the expected value is the MSE estimate of a random variable, irrespective of its probability density function.

Example 6.2 Let

$$\mathbf{y} = \mathbf{x} + \mathbf{v}$$

with $\mathbf{v}$ a scalar random variable with variance $\sigma_{\mathbf{v}}^2 = 1$, and $\mu_{\mathbf{v}} = 0$, $\mathbf{x}$ a scalar random variable with $\mu_{\mathbf{x}} = 1$, $\sigma_{\mathbf{x}}^2 = 4$, independent of $\mathbf{v}$.
We observe the variable $y = 2$. What is the linear MSE estimate of $\mathbf{x}$ in the format

$$\hat{\mathbf{x}} = a + b \cdot \mathbf{y}?$$

With $\mathbf{y} = \mathbf{x} + \mathbf{v}$ it follows that $\mu_{\mathbf{y}} = 1$.
$\sigma_{\mathbf{y}}^2 = \mathbb{E}[(\mathbf{x}+\mathbf{v}-1)^2] = \mathbb{E}[(\mathbf{x}-\mu_{\mathbf{x}})^2] + \mathbb{E}[\mathbf{v}^2] = 5$.
$\sigma_{\mathbf{xy}} = \mathbb{E}[(\mathbf{x}-\mu_{\mathbf{x}})(\mathbf{y}-\mu_{\mathbf{y}})] = \mathbb{E}[(\mathbf{x}-\mu_{\mathbf{x}})(\mathbf{x}-\mu_{\mathbf{x}}+\mathbf{v})] = \sigma_{\mathbf{x}}^2 = 4$.
Then applying the relations found in the analysis delivers:

$$\begin{aligned}b_{ms} &= \frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2} = \frac{4}{5} \\ a_{ms} &= \mu_{\mathbf{x}} - \mu_{\mathbf{y}}\frac{\sigma_{\mathbf{xy}}}{\sigma_{\mathbf{y}}^2} = \mu_{\mathbf{x}} - \frac{4}{5} = \frac{1}{5}.\end{aligned} \tag{6.19}$$

So the optimal linear MSE estimator is $\hat{\mathbf{x}} = \frac{1}{5} + \frac{4}{5}\mathbf{y}$, and for observation $y = 2$: $\hat{x} = 1/5 + 8/5 = 9/5$.

6.4 Linear MSE estimators - the vector case

This section will consider the situation of having multiple random variables as observed variables, $\mathbf{y}_1, \cdots, \mathbf{y}_n$ and using them in an estimator for a related scalar random variable:

$$\hat{\mathbf{x}} = \hat{x}(\mathbf{y}_1, \cdots, \mathbf{y}_n).$$

The linear case will be considered, where the structure of the estimator is:

$$\hat{\mathbf{x}} = a + \sum_{i=1}^{n} b_i \mathbf{y}_i = a + b^T \mathbf{y}$$

with $b = [b_1 \cdots b_n]^T$ and $\mathbf{y} = [\mathbf{y}_1 \cdots \mathbf{y}_n]^T$.
Substituting this expression into the MSE criterion $\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2]$ and setting the partial derivatives of the criterion with respect to a and b to zero, it follows that

$$\begin{aligned} 2\mathbb{E}[\mathbf{x} - \hat{\mathbf{x}}] &= 0 && (6.20) \\ 2\mathbb{E}[\mathbf{y}(\mathbf{x} - \hat{\mathbf{x}})] &= 0. && (6.21) \end{aligned}$$

The optimal filter, represented by (a_{ms}, b_{ms}), is then found by solving

$$\begin{aligned} 2\mathbb{E}[\mathbf{x} - a_{ms} - b_{ms}^T \mathbf{y}] &= 0 && (6.22) \\ 2\mathbb{E}[\mathbf{y}(\mathbf{x} - a_{ms} - b_{ms}^T \mathbf{y})] &= 0. && (6.23) \end{aligned}$$

This leads to

$$a_{ms} = \mu_{\mathbf{x}} - b_{ms}^T \mu_{\mathbf{y}} = \mu_{\mathbf{x}} - \mu_{\mathbf{y}}^T b_{ms} \tag{6.24}$$

while b_{ms} is determined by

$$\mathbb{E}[\mathbf{yx}] = a_{ms}\mu_{\mathbf{y}} + \mathbb{E}[\mathbf{y}\mathbf{y}^T] b_{ms}.$$

Substituting a_{ms} in this latter equation leads to

$$\mathbb{E}[\mathbf{yx}] - \mu_{\mathbf{x}}\mu_{\mathbf{y}} = (\mathbb{E}[\mathbf{y}\mathbf{y}^T - \mu_{\mathbf{y}}\mu_{\mathbf{y}}^T]) b_{ms}$$

or equivalently:

$$\mathbb{E}[(\mathbf{y} - \mu_{\mathbf{y}})(\mathbf{y} - \mu_{\mathbf{y}})^T] b_{ms} = \mathbb{E}[(\mathbf{y} - \mu_{\mathbf{y}})(\mathbf{x} - \mu_{\mathbf{x}})]$$

which is -according to the notational conventions - denoted as

$$\Sigma_{\mathbf{y}} b_{ms} = \Sigma_{\mathbf{yx}}. \tag{6.25}$$

This leads to the solution:

$$b_{ms} = \Sigma_{\mathbf{y}}^{-1} \cdot \Sigma_{\mathbf{yx}}, \tag{6.26}$$

where

$$\Sigma_{\mathbf{y}} = \begin{pmatrix} \sigma_{\mathbf{y}_1\mathbf{y}_1} & \sigma_{\mathbf{y}_1\mathbf{y}_2} & \cdots & \sigma_{\mathbf{y}_1\mathbf{y}_n} \\ \sigma_{\mathbf{y}_2\mathbf{y}_1} & \sigma_{\mathbf{y}_2\mathbf{y}_2} & \cdots & \sigma_{\mathbf{y}_2\mathbf{y}_n} \\ \vdots & \ddots & \ddots & \vdots \\ \sigma_{\mathbf{y}_n\mathbf{y}_1} & \sigma_{\mathbf{y}_n\mathbf{y}_2} & \cdots & \sigma_{\mathbf{y}_n\mathbf{y}_n} \end{pmatrix} \tag{6.27}$$

and

$$\Sigma_{\mathbf{yx}} = \begin{pmatrix} \sigma_{\mathbf{y}_1\mathbf{x}} \\ \sigma_{\mathbf{y}_2\mathbf{x}} \\ \vdots \\ \sigma_{\mathbf{y}_n\mathbf{x}} \end{pmatrix}. \tag{6.28}$$

Substituting the expression for the estimator coefficients into the MSE criterion, it follows that

$$\begin{aligned} \mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] &= \mathbb{E}[(\mathbf{x} - \mu_{\mathbf{x}} - b_{ms}^T(\mathbf{y} - \mu_{\mathbf{y}}))^2] && (6.29) \\ &= \sigma_{\mathbf{x}}^2 - 2\Sigma_{\mathbf{xy}} b_{ms} + b_{ms}^T \Sigma_{\mathbf{y}} b_{ms} && (6.30) \\ &= \sigma_{\mathbf{x}}^2 - \Sigma_{\mathbf{xy}} \Sigma_{\mathbf{y}}^{-1} \Sigma_{\mathbf{yx}}. && (6.31) \end{aligned}$$

Again, as in the scalar case, it can be observed that the MSE criterion reaches a smaller value the higher the correlation between $\mathbf{x}$ and the components of $\mathbf{y}$.
Summarizing one can state that

- The estimator based on the conditional expectation is always the optimal MSE estimator.
- In general the conditional expectation is hard to determine since it requires knowledge of the joint pdf's.
- The linear MSE estimator is simply determined on the basis of first and second order moments $(\Sigma_{\mathbf{xy}}, \sigma_{\mathbf{x}}, \Sigma_{\mathbf{y}})$
- In the situation of Gaussian joint pdf's, the estimator based on conditional expectation is a linear MSE estimator.

The orthogonality principle in linear MSE estimation

There is one important principle in linear MSE estimation that is applicable in many different filtering problems. It is called the *orthogonality principle*, and it states that the linear MSE estimator

$$\hat{\mathbf{x}} = b_{ms}^T \mathbf{y}$$

satisfies (and actually is determined by) the orthogonality property that

$$\mathbb{E}[\boldsymbol{\epsilon}\mathbf{y}] = 0 \quad \text{with } \boldsymbol{\epsilon} = \mathbf{x} - \hat{\mathbf{x}}. \tag{6.32}$$

In other words: the linear MSE estimator is constructed in such a way that the estimation error is orthogonal to the measured (regression) variables. Orthogonality of two random variables $\mathbf{v}$ and $\mathbf{w}$ has to be interpreted here on the basis of the inner product $< \mathbf{v}, \mathbf{w} >:= \mathbb{E}[\mathbf{v}\mathbf{w}^*]$. The situation is graphically illustrated in Figure 6.6.
The result is a direct consequence of the optimality conditions for the linear MSE estimator. When considering the MSE criterion for any value of the parameter vector b

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \mathbb{E}[(\mathbf{x} - b^T\mathbf{y})^2]$$

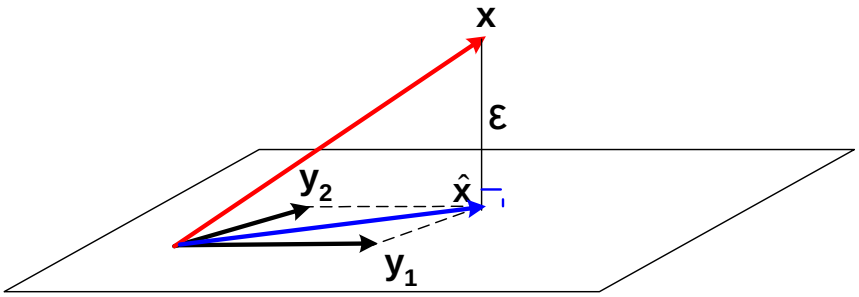

Figure 6.6: MSE estimation error is orthogonal to the measured variables $\mathbf{y}_i$.

and setting the partial derivative with respect to b to zero in $b = b_{ms}$, it follows that

$$-2\mathbb{E}[\mathbf{y}(\mathbf{x} - b_{ms}^T\mathbf{y})] = 0,$$

which directly shows the result.
Note that in this analysis we have not taken account of a constant term a in the linear estimator. Discarding this term actually comes down to restricting attention to the situation that both $\mathbf{x}$ and $\mathbf{y}$ have mean value zero, or have been preprocessed to have zero-mean. Note that in absence of a constant a in the estimator, the expression for b_{ms} becomes

$$b_{ms} = (\mathbb{E}[\mathbf{y}\mathbf{y}^T])^{-1}\mathbb{E}[\mathbf{y}\mathbf{x}] \tag{6.33}$$

with

$$\mathbb{E}[\mathbf{y}\mathbf{y}^T] = R_{\mathbf{y}\mathbf{y}} = \begin{pmatrix} r_{\mathbf{y}_1\mathbf{y}_1} & r_{\mathbf{y}_1\mathbf{y}_2} & \cdots & r_{\mathbf{y}_1\mathbf{y}_n} \\ r_{\mathbf{y}_2\mathbf{y}_1} & r_{\mathbf{y}_2\mathbf{y}_2} & \cdots & r_{\mathbf{y}_2\mathbf{y}_n} \\ \vdots & \ddots & \ddots & \vdots \\ r_{\mathbf{y}_n\mathbf{y}_1} & r_{\mathbf{y}_n\mathbf{y}_2} & \cdots & r_{\mathbf{y}_n\mathbf{y}_n} \end{pmatrix} \text{ and } \mathbb{E}[\mathbf{y}\mathbf{x}] = R_{\mathbf{y}\mathbf{x}} = \begin{pmatrix} r_{\mathbf{y}_1\mathbf{x}} \\ r_{\mathbf{y}_2\mathbf{x}} \\ \vdots \\ r_{\mathbf{y}_n\mathbf{x}} \end{pmatrix}.$$

The minimum value of the MSE criterion now becomes:

$$\begin{aligned} \mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] &= \mathbb{E}[(\mathbf{x} - b_{ms}^T\mathbf{y})^2] = & (6.34) \\ &= r_{\mathbf{x}\mathbf{x}} - 2b_{ms}^T R_{\mathbf{y}\mathbf{x}} + b_{ms}^T R_{\mathbf{y}\mathbf{y}} b_{ms} & (6.35) \\ &= r_{\mathbf{x}\mathbf{x}} - R_{\mathbf{y}\mathbf{x}}^T R_{\mathbf{y}}^{-1} R_{\mathbf{y}\mathbf{x}}. & (6.36) \end{aligned}$$

As a general expression for the minimal MSE one can also write:

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \mathbb{E}[\boldsymbol{\epsilon}(\mathbf{x} - b_{ms}\mathbf{y})] \tag{6.37}$$

and since $\mathbb{E}[\boldsymbol{\epsilon}\mathbf{y}] = 0$ this simplifies to

$$\mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})^2] = \mathbb{E}[\boldsymbol{\epsilon}\mathbf{x}]. \tag{6.38}$$

6.5 Optimal MSE filtering

6.5.1 Introduction

The results on MSE estimators will now be applied to the situation of (random) signals that are a sequenced in time[2]. This implies that the measured random variables $\mathbf{y}_1, \cdots \mathbf{y}_n$

[2] As before, the sequencing is considered in a time domain but can equivalently be handled in other domains, such as the spatial domain when considering, e.g., optical or image processing type of problems.

are part of a stochastic process, and therefore indexed in the time domain:

$$\mathbf{y}_1, \cdots, \mathbf{y}_n \rightarrow \mathbf{y}(1), \cdots \mathbf{y}(n).$$

The estimator to be determined can then be interpreted as a (dynamical) filter that operates on the stochastic process $\mathbf{y}$ in order to construct an estimator $\hat{\mathbf{x}}$ that is intended to match a related random variable $\mathbf{x}$ as close as possible, in MSE sense.

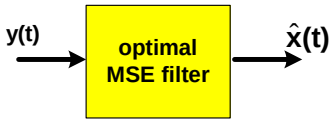

Figure 6.7: Optimal MSE filter operating on stochastic process $\mathbf{y}$.

First we will consider the most simple form of the problem, namely the linear prediction problem. Subsequently we will move to more general estimation and related filtering problems.

6.5.2 Linear prediction

In linear prediction the random variable to be estimated is $\mathbf{y}(t)$ while the available data is given by the measured random variables $\mathbf{y}(t-1), \cdots \mathbf{y}(t-n)$, $n \geq 1$. So it is assumed that past data is available, and that it is desired to estimate the current value $y(t)$. This occurs in many different (prediction) problems, as , e.g., occurring in tracking the position of a target; if $\mathbf{y}(t)$ represents the target position, a sensor tracking the target must be able to predict the target's next likely position before it arrives.
In the framework of the previously handled MSE estimation problem we simply substitute $\mathbf{x} = \mathbf{y}(t)$ and $\mathbf{y} = [\mathbf{y}(t-1) \cdots \mathbf{y}(t-n)]$. A linear estimator is then constructed of the form

$$\hat{\mathbf{y}}(t) = \sum_{i=1}^{n} h_i \mathbf{y}(t-i).$$

For optimality in MSE sense, the optimal filter, determined by the filter coefficients

$$h_{ms} = [h_1 \cdots h_n]^T$$

is then determined by the normal equations resulting from the orthogonality principle:

$$\mathbb{E}[(\mathbf{y}(t) - \sum_{i=1}^{n} h_i \mathbf{y}(t-i))\mathbf{y}(t-k)] = 0 \quad \text{for } k = 1, \cdots n$$

or equivalently

$$R_{\mathbf{y}}(k) - \sum_{i=1}^{n} h_i R_{\mathbf{y}}(k-i) = 0 \quad \text{for } k = 1, \cdots n$$

with $R_{\mathbf{y}}(k)$ the autocorrelation function of the stochastic process $\mathbf{y}$ at time shift k.
The resulting description is a set of n equations with n unknowns, being the filter coefficients h_i, $i = 1, \cdots n$. The equations can be written in a matrix form as:

$$\begin{bmatrix} R_{\mathbf{y}}(0) & R_{\mathbf{y}}(1) & \cdots & R_{\mathbf{y}}(n-1) \\ R_{\mathbf{y}}(1) & R_{\mathbf{y}}(0) & \cdots & R_{\mathbf{y}}(n-2) \\ \vdots & \ddots & & \vdots \\ R_{\mathbf{y}}(n-1) & R_{\mathbf{y}}(n-2) & \cdots & R_{\mathbf{y}}(0) \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \\ \vdots \\ h_n \end{bmatrix} = \begin{bmatrix} R_{\mathbf{y}}(1) \\ R_{\mathbf{y}}(2) \\ \vdots \\ R_{\mathbf{y}}(n) \end{bmatrix}. \tag{6.39}$$

Note that this result is simply a special case of the optimal linear MSE result (6.33). The matrix on the left hand side of the equation is a so-called correlation matrix (see also section 3.7.2), and because of its particular structure having equal elements on all (sub)diagonals it is a Toeplitz matrix. Inversion of this matrix is necessary for obtaining the optimal filter coefficients.
The resulting MSE of this estimator

$$\mathbb{E}[\boldsymbol{\epsilon}(t)^2] \quad \text{with } \boldsymbol{\epsilon}(t) = \mathbf{y}(t) - \sum_{i=1}^{n} h_i \mathbf{y}(t-i)$$

then becomes, using (6.38),

$$\mathbb{E}[\boldsymbol{\epsilon}(t)^2] = \mathbb{E}[\boldsymbol{\epsilon}(t)\mathbf{y}(t)] = R_{\mathbf{y}}(0) - \sum_{i=1}^{n} h_i R_{\mathbf{y}}(i).$$

6.5.3 Optimal filtering - the FIR case

In the most simple general filtering problem, we consider a variable to be estimated that is different from the measurable stochastic process $\mathbf{y}$. In this situation a linear estimator

$$\hat{\mathbf{x}}(t) = \sum_{i=1}^{n} h_i \mathbf{y}(t-i)$$

is constructed for a stochastic process $\mathbf{x}$ that is supposed to be related to measurements $\mathbf{y}$. This linear estimator is called an FIR (finite impulse response) estimator since h_i reflects the finite pulse response sequence of the dynamical filter with $\mathbf{y}$ as input and $\hat{\mathbf{x}}$ as output. The analysis of this situation is very similar to the previous case. The normal equations resulting from the orthogonality principle now become:

$$\mathbb{E}[(\mathbf{x}(t) - \sum_{i=1}^{n} h_i \mathbf{y}(t-i))\mathbf{y}(t-k)] = 0 \quad \text{for } k = 1, \cdots n$$

or equivalently

$$R_{\mathbf{xy}}(k) - \sum_{i=1}^{n} h_i R_{\mathbf{y}}(k-i) = 0 \quad \text{for } k = 1, \cdots n$$

leading to the set of equations

$$\begin{bmatrix} R_{\mathbf{y}}(0) & R_{\mathbf{y}}(1) & \cdots & R_{\mathbf{y}}(n-1) \\ R_{\mathbf{y}}(1) & R_{\mathbf{y}}(0) & \cdots & R_{\mathbf{y}}(n-2) \\ \vdots & \ddots & & \vdots \\ R_{\mathbf{y}}(n-1) & R_{\mathbf{y}}(n-2) & \cdots & R_{\mathbf{y}}(0) \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \\ \vdots \\ h_n \end{bmatrix} = \begin{bmatrix} R_{\mathbf{xy}}(1) \\ R_{\mathbf{xy}}(2) \\ \vdots \\ R_{\mathbf{xy}}(n) \end{bmatrix}. \tag{6.40}$$

The resulting MSE of this estimator then becomes, using (6.38),

$$\mathbb{E}[\boldsymbol{\epsilon}(t)^2] = \mathbb{E}[\boldsymbol{\epsilon}(t)\mathbf{x}(t)] = R_{\mathbf{x}}(0) - \sum_{i=1}^{n} h_i R_{\mathbf{xy}}(i).$$

6.6 Wiener filter

6.6.1 Introduction

We consider the optimal filter problem in its most basic form, as depicted in Figure 6.8, where $\mathbf{x}$ and $\mathbf{v}$ are assumed to be uncorrelated stochastic processes.

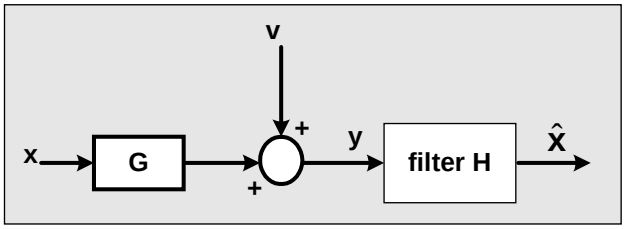

Figure 6.8: Wiener filter problem

The filter problem concerns the construction of optimal filters:

$$\hat{\mathbf{x}}(t) = \sum_{k=-\infty}^{\infty} h(k)\mathbf{y}(t-k) \quad \text{IIR-noncausal}$$

or

$$\hat{\mathbf{x}}(t) = \sum_{k=0}^{\infty} h(k)\mathbf{y}(t-k) \quad \text{IIR-causal}$$

such that $\hat{\mathbf{x}}(t)$ is the optimal MSE estimate of $\mathbf{x}(t)$.
Note that the filters in the two cases are described by $H(z) = \sum_{k=-\infty}^{\infty} h(k)z^{-k}$ (non-causal filter), and $H(z) = \sum_{k=0}^{\infty} h(k)z^{-k}$ (causal filter).

If the filter has a finite impulse response (FIR): $\{h(k)\}_{k=-M_1,\cdots,M_2}$ then the optimal solution is available from the previous section, and determined by the equation (6.40).
However, for filters with an infinite-length pulse response a new analysis will appear to be necessary. For simplicity, in the sequel it will be assumed that $\mathbf{x}$ and $\mathbf{v}$ are zero-mean stationary stochastic processes. Additionally it is, for the moment, assumed that the spectral densities of $\mathbf{x}$ and $\mathbf{v}$ are known, as well as the dynamical transfer function G. We will reconsider these assumptions at the end of this section. Furthermore, it will be analyzed in Chapter 7 how spectral density information can be obtained (i.e. estimated) from practical data.

6.6.2 Non-causal Wiener filter

In the situation of a non-causal Wiener filter, the estimator for $\hat{\mathbf{x}}(t)$ is parameterized as

$$\hat{\mathbf{x}}(t) = \sum_{k=-\infty}^{\infty} h(k)\mathbf{y}(t-k),$$

where $h(k)$ are the pulse response coefficients of the filter to be determined.
The orthogonality conditions for the linear MSE estimate as formulated in the orthogonality principle (6.32), show that in the optimum, $\mathbf{x}(t) - \hat{\mathbf{x}}(t)$ is orthogonal to the components $\{\mathbf{y}(t-k)\}_{k=-\infty,\cdots,\infty}$. Consequently

$$\mathbb{E}[(\mathbf{x}(t) - \hat{\mathbf{x}}(t))\mathbf{y}(t-i)] = 0 \qquad i = -\infty, \cdots, \infty$$

As a result: $R_{\mathbf{xy}}(i) = \sum_{k=-\infty}^{\infty} h(k) R_{\mathbf{y}}(i-k) \quad \forall i.$
When applying z-transform to both sides of this equation, it follows that $\Phi_{\mathbf{xy}}(z) = H(z) \cdot \Phi_{\mathbf{y}}(z)$, leading to the following simple expression for the optimal filter:

$$H(z) = \frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}(z)} \tag{6.41}$$

Remarks

- The filter is simply determined from $\Phi_{\mathbf{y}}(z)$ and $\Phi_{\mathbf{xy}}(z)$. Since $\mathbf{y}$ is measurable, the spectral density $\Phi_{\mathbf{y}}$ can be estimated from data (see the subsequent chapter). For constructing $\Phi_{\mathbf{xy}}$, additional knowledge of $\Phi_{\mathbf{x}}$ is required.
- The filter is not necessarily causal, and therefore it is likely not implementable in an on-line form. I.e., future values of $\mathbf{y}(t)$ are required for estimating $\mathbf{x}(t)$.
- If $\mathbf{v} \equiv 0$, then $\mathbf{y} = G\mathbf{x}$ and $H(z) = 1/G(z)$ (the inverse filter of G is the optimal estimate).

Resulting MSE of estimated signal

The MSE of the resulting optimal estimate can be written as

$$\begin{aligned} P &:= \mathbb{E}[\{\mathbf{x}(t) - \hat{\mathbf{x}}(t)\}^2] = \mathbb{E}[\{\mathbf{x}(t) - \hat{\mathbf{x}}(t)\}\mathbf{x}(t)] \\ &= R_{\mathbf{x}}(0) - \sum_{k=-\infty}^{\infty} h(k) R_{\mathbf{yx}}(-k). \end{aligned}$$

In order to find a frequency domain expression for this MSE, define

$$\bar{P}_m = R_{\mathbf{x}}(m) - \sum_{k=-\infty}^{\infty} h(k) R_{\mathbf{yx}}(m-k).$$

Note that $P = \bar{P}_0$. When considering $\bar{P}_m$ as a sequence, it can be Fourier transformed to obtain

$$\bar{P}(e^{i\omega}) = \Phi_{\mathbf{x}}(\omega) - H(e^{i\omega})\Phi_{\mathbf{yx}}(\omega).$$

The relation of Parseval implies that $P = \bar{P}_0 = \frac{1}{2\pi}\int_{-\pi}^{\pi} \bar{P}(e^{i\omega})d\omega$, showing that

$$P = \frac{1}{2\pi}\int_{-\pi}^{\pi} [\Phi_{\mathbf{x}}(\omega) - H(e^{i\omega})\Phi_{\mathbf{yx}}(\omega)]d\omega. \tag{6.42}$$

Example 6.3 Show that for the Wiener filter problem

$$\mathbf{y} = G\mathbf{x} + \mathbf{v}$$

with $\mathbf{x}$ and $\mathbf{v}$ uncorrelated, the optimal noncausal Wiener filter has a MSE given by

$$\mathbb{E}[(\mathbf{x}(t) - \hat{\mathbf{x}}(t))^2] = \frac{1}{2\pi}\int_{-\pi}^{\pi} \frac{\Phi_{\mathbf{x}}\Phi_{\mathbf{v}}}{|G|^2\Phi_{\mathbf{x}} + \Phi_{\mathbf{v}}} d\omega.$$

As a result it is clear that the MSE error is 0 whenever the two spectral density functions $\Phi_{\mathbf{x}}$ and $\Phi_{\mathbf{v}}$ do not overlap in the frequency domain (see figure 6.4).

Solution:
Substituting the expression (6.41) for $H(e^{i\omega})$ into (6.42), the integrand of this expression becomes

$$\Phi_{\mathbf{x}}(\omega) - \frac{\Phi_{\mathbf{xy}}(\omega)}{\Phi_{\mathbf{y}}(\omega)}\Phi_{\mathbf{yx}}(\omega).$$

With the expression for $\mathbf{y}$, it follows that $\Phi_{\mathbf{xy}}\Phi_{\mathbf{yx}} = |G|^2\Phi_{\mathbf{x}}^2$. With $\Phi_{\mathbf{y}} = |G|^2\Phi_{\mathbf{x}} + \Phi_{\mathbf{v}}$ the integrand becomes

$$\Phi_{\mathbf{x}} - \frac{|G|^2\Phi_{\mathbf{x}}^2}{|G|^2\Phi_{\mathbf{x}} + \Phi_{\mathbf{v}}} = \frac{\Phi_{\mathbf{x}}\Phi_{\mathbf{v}}}{|G|^2\Phi_{\mathbf{x}} + \Phi_{\mathbf{v}}}.$$

If $\Phi_{\mathbf{x}}$ and $\Phi_{\mathbf{v}}$ do not overlap in the frequency domain, their product equals 0 for all frequencies and consequently the MSE is 0.

6.6.3 Causal Wiener filter

When restricting the optimal filter to be causal (and therefore also on-line implementable), the filter takes the form

$$\hat{x}(t) = \sum_{k=0}^{\infty} h(k)\mathbf{y}(t-k).$$

The orthogonality conditions for the minimum MSE estimate now show that

$$R_{\mathbf{xy}}(j) = \sum_{k=0}^{\infty} h(k)R_{\mathbf{y}}(j-k) \qquad j = 0, 1, 2, \cdots \tag{6.43}$$

and this set of equations is referred as the (discrete-time) Wiener-Hopf equation.
Because of the fact that the summation intervals are not doubly infinite, but bounded below by 0, this set of equations can not simply be solved by z- or Fourier transform techniques, as was done in the non-causal case.[3]

As an alternative the following strategy is used for solving the set of equations. First it will be assumed that $\mathbf{y}$ is a stationary white noise stochastic process, and a solution to the problem will be constructed for this particular situation. Then in a second stage this solution is generalized to the situation of non-white processes $\mathbf{y}$.

[3]Note that an attempt to transform the equations to the frequency domain leads to $\sum_{\ell=0}^{\infty} R_{\mathbf{xy}}(\ell)e^{-i\omega\ell} = \sum_{k=0}^{\infty} h(k)e^{-i\omega k}\sum_{\ell=0}^{\infty} R_{\mathbf{y}}(\ell-k)e^{-i\omega(\ell-k)}$. Now the second summation in the right-hand side can not be made independent of k.

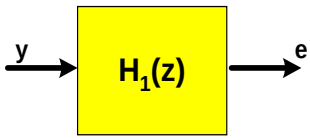

Figure 6.9: White-shaping filter for process $\mathbf{y}$.

For $\mathbf{y}$ being a stationary white noise process with variance $\sigma^2_{\mathbf{y}}$, and consequently $R_{\mathbf{y}}(j) = \sigma^2_{\mathbf{y}}\delta(j)$, the optimality conditions (6.43) reduce to

$$R_{\mathbf{xy}}(j) = \sum_{k=0}^{\infty} h(k) \cdot \sigma^2_{\mathbf{y}} \cdot \delta(j-k) = h(j) \cdot \sigma^2_{\mathbf{y}} \qquad j \geq 0.$$

As a result the optimal filter coefficients are given by

$$h(j) = \begin{cases} \frac{1}{\sigma^2_{\mathbf{y}}} \cdot R_{\mathbf{xy}}(j) & j \geq 0 \\ 0 & j < 0 \end{cases}$$

and consequently the optimal causal Wiener filter can be written as

$$H(z) = \frac{1}{\sigma^2_{\mathbf{y}}} \cdot [\Phi_{\mathbf{xy}}(z)]_+ = \left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}(z)}\right]_+ \tag{6.44}$$

where the operation $[\cdot]_+$ (see also (2.21)) refers to taking the causal part of the filter, in the series expansion that converges on the unit circle; i.e., if a filter $H(z)$ is written as

$$H(z) = \sum_{k=-\infty}^{\infty} h(k) z^{-k}$$

with $\sum_{k=-\infty}^{\infty} |h(k)|$ bounded, then

$$[H(z)]_+ = \sum_{k=0}^{\infty} h(k) z^{-k}.$$

In case the measurable process $\mathbf{y}$ is non-white, a two-step procedure can be followed to find the optimal causal Wiener filter, according to the following strategy:

- Step 1: Make $\mathbf{y}$ white by applying a dynamic "white-shaping" filter H_1 that is causal, and such that $\mathbf{e} = H_1\mathbf{y}$ is a white noise process.
- Step 2: Determine the optimal causal filter for white noise process $\mathbf{e}$.

First the optimal filter will be derived according to this two-step procedure. Subsequently it will be motivated that it indeed leads to an overall optimal causal Wiener filter.

Step 1

For a given stochastic process $\mathbf{y}$ with (rational) power spectral density $\Phi_{\mathbf{y}}$, there exists a minimum-phase system represented by H_m, such that

$$\Phi_{\mathbf{y}}(z) = H_m(z) \cdot H_m(1/z), \tag{6.45}$$

(with $H_m(z)$ a rational function in z). This is formulated in the Spectral Factorization Theorem 4.4, where without loss of generality the scalar constant σ^2 can be chosen to be 1. The minimum-phase factor H_m satisfies: $\Phi_{\mathbf{y}}(\omega) = |H_m(e^{i\omega})|^2$ (on the unit circle). Additionally, $\mathbf{y}$ can be written as $\mathbf{y}(t) = H_m(q)\mathbf{e}(t)$ with $\mathbf{e}$ a unit variance white noise process, and so alternatively

$$\mathbf{e} = \frac{1}{H_m}\mathbf{y}$$

where $1/H_m$ is again a causal and stable filter operation.

Summarizing, let

$$\Phi_{\mathbf{y}}(z) = \Phi_{\mathbf{y}}^{+}(z) \cdot \Phi_{\mathbf{y}}^{-}(z)$$

be the spectral factorization of $\Phi_{\mathbf{y}}$ according to (6.45), with $\Phi_{\mathbf{y}}^{+}(z)$ the minimum-phase spectral factor $H_m(z)$. Then the filter operation

$$H_1(z) = \frac{1}{\Phi_{\mathbf{y}}^{+}(z)}$$

is causal and stable, and transfers the process $\mathbf{y}$ into a unit variance white noise.

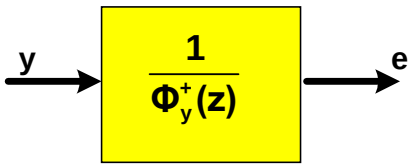

Figure 6.10: White-shaping filter for process $\mathbf{y}$.

Step 2

Using the result (6.44) for white noise processes, it follows now that the optimal causal Wiener filter for $\mathbf{e}$ is given by

$$H_2(z) = \left[\frac{\Phi_{\mathbf{xe}}(z)}{\Phi_{\mathbf{e}}(z)}\right]_+.$$

However since $\mathbf{e}$ is only an intermediate (artificially constructed) signal, an expression for the filter in terms of $\mathbf{y}$ is preferred, and so one has to substitute the relation $\mathbf{e} = H_1\mathbf{y}$. As a result

$$H_2(z) = \left[\frac{H_1(1/z)\Phi_{\mathbf{xy}}(z)}{H_1(z)H_1(1/z)\Phi_{\mathbf{y}}(z)}\right]_+ \tag{6.46}$$

$$= \left[\frac{\Phi_{\mathbf{xy}}(z)}{H_1(z)\Phi_{\mathbf{y}}(z)}\right]_+. \tag{6.47}$$

Since $\Phi_{\mathbf{y}}(z) = \Phi_{\mathbf{y}}^{+}(z) \cdot \Phi_{\mathbf{y}}^{-}(z)$ and $H_1(z) = 1/\Phi_{\mathbf{y}}^{+}(z)$ this becomes:

$$H_2(z) = \left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)}\right]_+. \tag{6.48}$$

The optimal filter for $\mathbf{e}$ is sketched in Figure 6.11.

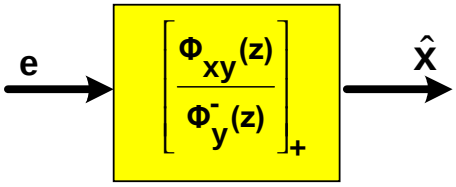

Figure 6.11: Optimal causal Wiener filter for white noise **e**.

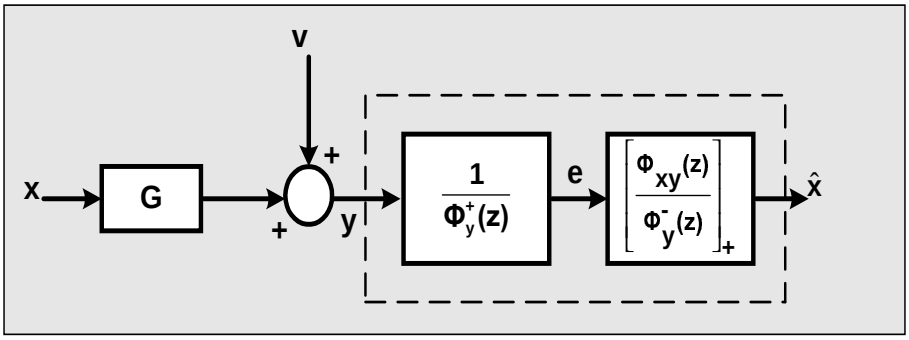

Figure 6.12: Optimal causal Wiener filter construction

When combining the two steps of the procedure, the causal Wiener filter is specified by

$$\hat{\mathbf{x}}(t) = \frac{1}{\Phi_{\mathbf{y}}^{+}(z)} \cdot \left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)}\right]_{+} \cdot \mathbf{y}(t). \tag{6.49}$$

The result for the optimal causal Wiener filter is now depicted schematically in Figure 6.12.

Note that in the diagram of Figure 6.12 also the optimal non-causal Wiener filter appears. This is visualized by removing the $[\cdot]_+$ operation, and by realizing that $\Phi_{\mathbf{y}}^{+}(z) \cdot \Phi_{\mathbf{y}}^{-}(z) = \Phi_{\mathbf{y}}(z)$.

Example 6.4 **x** is a so-called AR(1) process:

$$\mathbf{x}(t) - 0.8 \cdot \mathbf{x}(t-1) = 0.6\mathbf{w}(t)$$

$$\mathbf{y}(t) = \mathbf{x}(t) + v(t)$$

w and **v** unit variance white noise (uncorrelated with each other).
The filter that transforms $\mathbf{w}(t)$ into $\mathbf{x}(t)$ is given by $H(z) = 0.6/(1 - 0.8z^{-1})$. As a result:

$$\begin{aligned}
\Phi_{\mathbf{x}}(z) &= H(z)H(z^{-1})\sigma_{\mathbf{w}}^2 \\
&= \frac{0.6}{1-0.8z^{-1}} \cdot \frac{0.6}{1-0.8z} = \frac{0.36}{(1-0.8z^{-1})(1-0.8z)} \\
&= \frac{-0.45z}{(z-0.8)(z-1.25)} \\
\Phi_{\mathbf{xy}}(z) &= \Phi_{\mathbf{x}}(z) \\
\Phi_{\mathbf{y}}(z) &= \Phi_{\mathbf{x}}(z) + 1 = \frac{(z-0.5)(z-2)}{(z-0.8)(z-1.25)}
\end{aligned}$$

Non-causal Wiener filter:

$$\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}(z)} = \frac{-0.45z}{(z-2)(z-0.5)}$$

For constructing the causal Wiener filter:

$$\Phi_{\mathbf{y}}(z) = \frac{(z-0.5)(z-2)}{(z-0.8)(z-1.25)} = \underbrace{\frac{(z-0.5)}{(z-0.8)}}_{\Phi_{\mathbf{y}}^{+}} \cdot \underbrace{\frac{(z-2)}{(z-1.25)}}_{\Phi_{\mathbf{y}}^{-}}$$

$$\begin{aligned}\left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)}\right]_{+} &= \left[\frac{-0.45z}{(z-0.8)(z-1.25)} \cdot \frac{(z-1.25)}{(z-2)}\right]_{+} \\ &= \left[\frac{-0.45z}{(z-0.8)(z-2)}\right]_{+}\end{aligned}$$

In order to decompose the rational function into a causal and an anti-causal expansion, we need to write the function as a linear combination of basic terms (see Appendix B.5):

$$\begin{aligned}&\frac{z}{z-a}, \quad |a|<1 \qquad \text{for the causal part} \\ &\frac{z}{1-cz}, \quad |c|<1 \qquad \text{for the anti-causal part.}\end{aligned}$$

Since

$$\frac{-0.45z}{(z-0.8)(z-2)} = \frac{0.375z}{z-0.8} + \frac{0.1875z}{1-0.5z}$$

it follows that

$$\left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)}\right]_{+} = \frac{0.375z}{z-0.8}.$$

The causal Wiener filter is now specified as follows:

$$\begin{aligned}H_1(z)\cdot H_2(z) &= \frac{1}{\Phi_{\mathbf{y}}^{+}(z)} \cdot \left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)}\right]_{+} = \frac{z-0.8}{z-0.5} \cdot \frac{0.375z}{z-0.8} \\ &= \frac{0.375z}{z-0.5}\end{aligned}$$

Time realizations of $\mathbf{y}$ and $\mathbf{x}$, as well as of the filtered signals (causal and non-causal) are plotted in figures 6.13 and 6.14. Note that the several signals are shifted over a constant amplitude, in order to display the differences more clearly.

MSE values of the several estimates are for $\mathbf{x}-\mathbf{y}$: 0.9288; $\mathbf{x}-\hat{\mathbf{x}}_c$: 0.3572; $\mathbf{x}-\hat{\mathbf{x}}_{nc}$: 0.2839. The non-causal filter leads to a smaller MSE value as it allows more flexibility in the filter than the causal one.

Optimality of the two-step procedure

It is not straightforward that the solution of the causal Wiener filter in the two-step procedure as described is the optimal solution. The estimator $\hat{\mathbf{x}}$ should be the optimal MSE

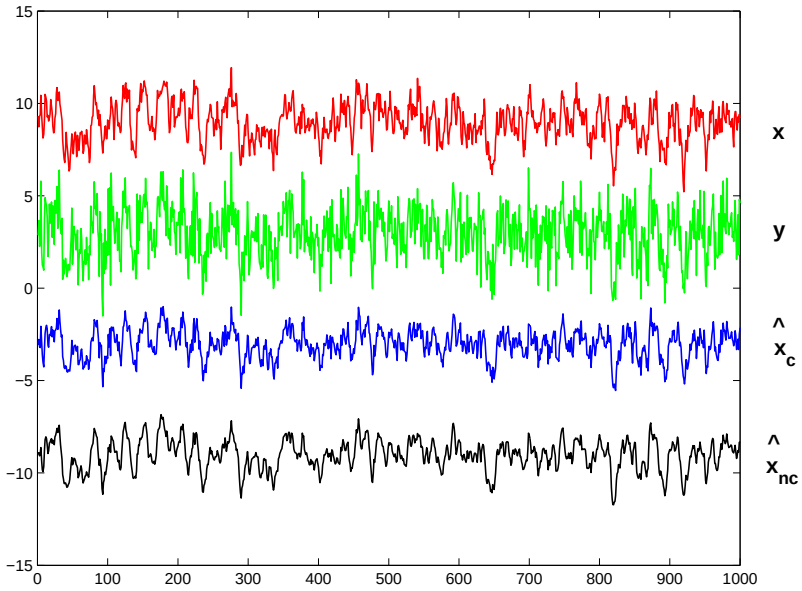

Figure 6.13: Time realizations of $\mathbf{x}$, $\mathbf{y}$, $\hat{\mathbf{x}}_c$ (causal Wiener filter) and $\hat{\mathbf{x}}_{nc}$ (non-causal Wiener filter).

estimator based on the past information $\mathbf{y}(t-1), \mathbf{y}(t-2), \cdots$. However in step 2 the optimal estimator is constructed on the basis of past "information" $\mathbf{e}(t-1), \mathbf{e}(t-2), \cdots$. Because of the fact that $\mathbf{y}$ and $\mathbf{e}$ are related through a minimum phase system (causal, and having a causal inverse), past information on $\mathbf{y}$ is equivalent to past information of $\mathbf{e}$, and this justifies the MSE optimality of the two-step procedure that is considered here.

Summary Wiener filters

The optimal Wiener filter (causal or non-causal) optimizes

$$\mathbb{E}[\boldsymbol{\epsilon}(t)^2] = \mathbb{E}[(\mathbf{x}(t) - \hat{\mathbf{x}}(t))^2]$$

where the estimation error is schematically depicted in Figure 6.15.

Which information is required to design a Wiener filter?

In order to design a Wiener filter knowledge is required of the spectral densities $\Phi_{\mathbf{xy}}$ and $\Phi_{\mathbf{y}}$. If there is no explicit information on $\Phi_{\mathbf{xy}}$, the required knowledge is reflected in

- Transfer function G (when present), and
- Two out of the three spectral density functions $\Phi_{\mathbf{y}}, \Phi_{\mathbf{v}}, \Phi_{\mathbf{x}}$.

This can be understood by notifying that $\Phi_{\mathbf{xy}} = G^* \Phi_{\mathbf{x}}$ and $\Phi_{\mathbf{y}} = |G|^2 \Phi_{\mathbf{x}} + \Phi_{\mathbf{v}}$.
If G is known, then knowledge of two of the three spectral densities that occur in the second equation, implies knowledge of the third function also. This gives full information on both $\Phi_{\mathbf{xy}}$ and $\Phi_{\mathbf{y}}$. Therefore, next to the spectral density of the measurable signal $\mathbf{y}$, information on the spectral density of either $\mathbf{x}$ or $\mathbf{v}$ should be available.

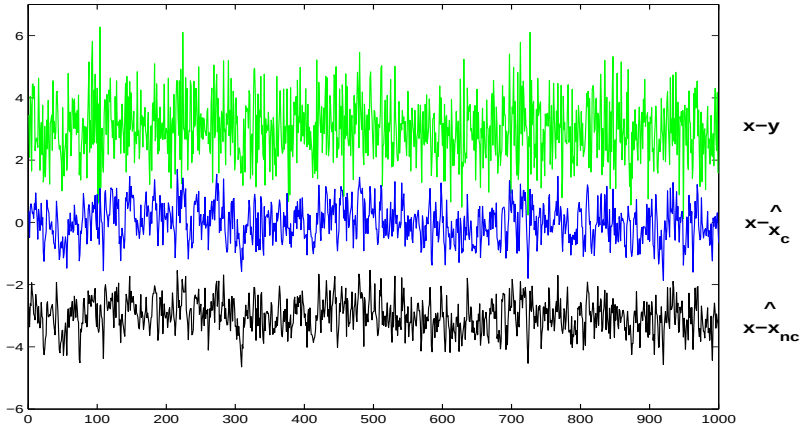

Figure 6.14: Time realizations of $\mathbf{x} - \mathbf{y}$, $\mathbf{x} - \hat{\mathbf{x}}_c$ and $\mathbf{x} - \hat{\mathbf{x}}_{nc}$

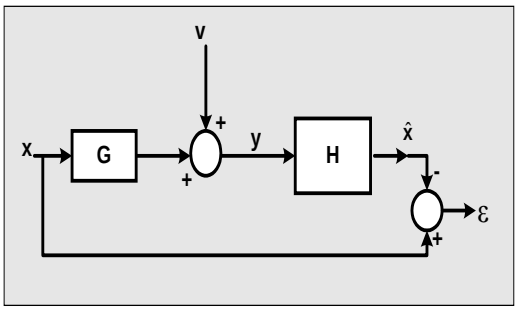

Figure 6.15: Diagram of optimal filter and estimation error.

6.7 Examples

6.7.1 Noise cancellation

In problems of communication it is of interest to remove noise from a communication channel. Consider, e.g., the radio communication between an airplane pilot and the traffic control tower. Background noise of the engine, as well as wind noise within the cockpit can severely hinder comprehensible communication.
Again the problem is to estimate an information carrying signal $\mathbf{x}(t)$ from noisy observations

$$\mathbf{y}(t) = \mathbf{x}(t) + \mathbf{v}_1(t).$$

The problem is illustrated in Figure 6.16. The - standard - Wiener filter is in this case given by:

$$H(z) = \frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}(z)}, \tag{6.50}$$

with

$$\Phi_{\mathbf{xy}}(z) = \Phi_{\mathbf{x}}(z) + \Phi_{\mathbf{xv}_1}(z), \tag{6.51}$$
$$\Phi_{\mathbf{y}}(z) = \Phi_{\mathbf{x}}(z) + \Phi_{\mathbf{v}_1}(z). \tag{6.52}$$

If we assume that $\mathbf{x}$ and $\mathbf{v}_1$ are uncorrelated with each other, $\Phi_{\mathbf{xy}}(z) = 0$ and (6.51) reduces to $\Phi_{\mathbf{xy}}(z) = \Phi_{\mathbf{x}}(z)$. Hence, the Wiener filter is given by

$$H(z) = \frac{\Phi_{\mathbf{x}}(z)}{\Phi_{\mathbf{x}}(z) + \Phi_{\mathbf{v}_1}(z)}. \tag{6.53}$$

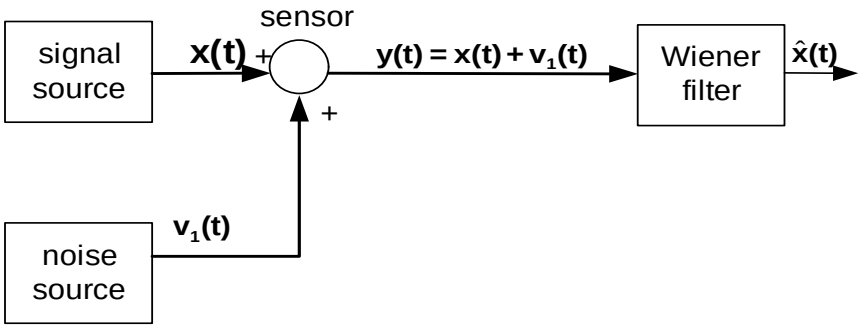

Figure 6.16: Noise cancellation.

For applying the -standard- Wiener filter one thus needs to know the spectral density functions $\Phi_{\mathbf{x}}(\omega)$ and $\Phi_{\mathbf{v}_1}(\omega)$ or $\Phi_{\mathbf{x}}(\omega)$ and $\Phi_{\mathbf{y}}(\omega)$. Whereas $\Phi_{\mathbf{y}}(\omega)$ can probably be estimated from measurable data $\mathbf{y}$, the spectral density functions $\Phi_{\mathbf{x}}(\omega)$ and $\Phi_{\mathbf{v}_1}(\omega)$ require knowledge of the statistical properties of either $\mathbf{x}$ and/or $\mathbf{v}_1$. If these properties are not known (and this will generally be the case in the considered application), a second sensor measurement can be helpful, as indicated in Figure 6.17.

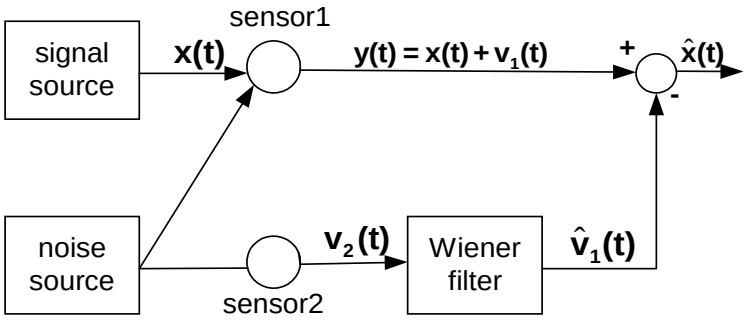

Figure 6.17: Noise cancellation by using a second (noise) sensor.

With sensor 1 the source signal $\mathbf{x}(t)$ is measured while being corrupted with the noise $\mathbf{v}_1(t)$. With the extra sensor 2 only the noise signal is measured, e.g., by positioning a sensor somewhere else in the cockpit. However as the position of both sensors is different, sensor 2 will measure a noise signal $\mathbf{v}_2(t)$ that is different from $\mathbf{v}_1(t)$. For both noise signals, the source is likely to be the same; therefore $\mathbf{v}_2$ and $\mathbf{v}_1$ will be highly correlated.

The Wiener filter problem now is formulated as follows:
Given the observed signal $\mathbf{v}_2(t) = G(q)\mathbf{v}_1(t)$ with $G(q)$ unknown. Estimate $v_1(t)$ on the basis of $v_2(t)$.
The non-causal Wiener filter solution for this problem is

$$H(z) = \frac{\Phi_{\mathbf{v}_1\mathbf{v}_2}(z)}{\Phi_{\mathbf{v}_2}(z)}$$

By writing $\mathbf{v}_1(t) = \mathbf{y}(t) - \mathbf{x}(t)$ it follows that $\Phi_{\mathbf{v}_1\mathbf{v}_2}(z) = \Phi_{(\mathbf{y}-\mathbf{x})\mathbf{v}_2}(z)$. If $\mathbf{v}_2$ can be assumed to be independent of $\mathbf{x}$, then $\Phi_{\mathbf{v}_1\mathbf{v}_2}(z) = \Phi_{\mathbf{y}\mathbf{v}_2}(z)$ and as a result

$$H(z) = \frac{\Phi_{\mathbf{y}\mathbf{v}_2}(z)}{\Phi_{\mathbf{v}_2}(z)}.$$

The expression for this filter is dependent on spectral densities of measurable signals $\mathbf{y}, \mathbf{v}_2$ only. So, if these spectral densities can be estimated from measurement data, the optimal Wiener filter can be constructed without requiring further prior know of, e.g., the spectral density of $\mathbf{x}$.

Exercise Show that the MSE error of the estimate $\hat{\mathbf{x}}(t)$ is 0 when sensor 2 can measure the noise source without measurement error.

Solution
With $\mathbf{v}_2(t) = G(q)\mathbf{v}_1(t) + \mathbf{w}(t)$, and $\mathbf{w}(t)$ a measurement (disturbance) error, it follows that the Wiener filter for constructing $\hat{\mathbf{v}}_1(t)$ is given by

$$H(z) = \frac{\Phi_{\mathbf{v}_1\mathbf{v}_2}(z)}{\Phi_{\mathbf{v}_2}(z)} = \frac{G(1/z)\Phi_{\mathbf{v}_1}(z)}{|G(z)|^2\Phi_{\mathbf{v}_1}(z) + \Phi_{\mathbf{w}}(z)}.$$

Whenever $\Phi_{\mathbf{w}}(z) = 0$, it follows that $H(z) = 1/G(z)$.
Then

$$\begin{aligned}
\hat{\mathbf{x}}(t) &= \mathbf{y}(t) - \hat{\mathbf{v}}_1(t) = && (6.54)\\
&= \mathbf{y}(t) - 1/G(q) \cdot \mathbf{v}_2(t) = && (6.55)\\
&= \mathbf{y}(t) - 1/G(q)[G(q)\mathbf{v}_1(t) + \mathbf{w}(t)] && (6.56)\\
&= \mathbf{y}(t) - 1/G(q)[G(q)(\mathbf{y}(t) - \mathbf{x}(t)) + \mathbf{w}(t)] && (6.57)\\
&= \mathbf{x}(t) - G^{-1}(q)\mathbf{w}(t). && (6.58)
\end{aligned}$$

Consequently the MSE estimate is exact when $\mathbf{w} \equiv 0$.

6.7.2 Deconvolution

Many measurement problems in optics come down to reconstructing an (optical) source signal on the basis of an optical measurement (lens) system that influences the original source through a point-spread function (psf). A point source will then be measured as a function that is spread out over a certain are, while this function is characterized by the (im)pulse response of the measurement system. This situation is covered by the setup that is sketched in Figure 6.18, where the dynamical transfer G reflects the operation of the point-spread function which can be interpreted in a (two-dimensional) spatial domain.

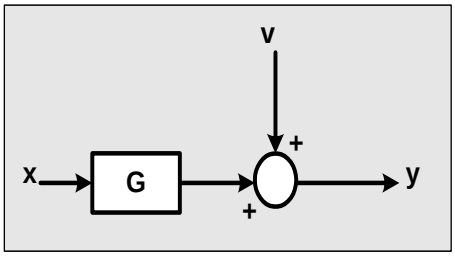

Figure 6.18: Optical system with point spread function characterized by G.

Since the operation G is actually a convolution of $\mathbf{x}$ with the point spread function of G, the reconstruction of $\mathbf{x}$ on the basis of $G\mathbf{x}$ is actually a deconvolution.
The most straightforward suggestion for reconstructing $\mathbf{x}$ on the basis of measurements $\mathbf{y}$ is by processing $\mathbf{y}$ through the inverse of G. This is generally referred to as *inverse filtering*.
Note that under influence of a disturbance process $\mathbf{v}$ the resulting reconstructed signal will be:

$$G^{-1}\mathbf{y} = G^{-1}(G\mathbf{x} + \mathbf{v}) = \mathbf{x} + G^{-1}\mathbf{v}.$$

Consequently the reconstructed signal has an "error" component $G^{-1}\mathbf{v}$ which can become excessive, dependent on properties of G and $\mathbf{v}$. If G is very small in (frequency) areas where $\mathbf{v}$ is present, this noise disturbance will be strongly amplified in the reconstruction of $\mathbf{x}$ and therefore will be dominantly present.
As a result this straightforward inverse filter is often a very poor choice, and outperformed by a properly designed Wiener filter that adequately takes account of the disturbance term $\mathbf{v}$.

6.7.3 One-step-ahead prediction

The optimal IIR causal Wiener filter for a one-step-ahead prediction problem can be derived as follows.
Suppose there is a measurable stochastic process $\mathbf{x}$ with spectral density $\Phi_{\mathbf{x}}$ that is modelled as

$$\mathbf{x}(t) = W(q)\mathbf{e}(t)$$

with $\mathbf{e}$ a white noise process with variance σ^2, and W a linear filter that is minimum phase and monic, i.e., $W(z) = 1 + \sum_{k=1}^{\infty} w_k z^{-k}$. Consequently, $\Phi_{\mathbf{x}}(\omega) = |W(e^{i\omega})|^2 \cdot \sigma^2$.
In a prediction problem we observe $\mathbf{y} = \mathbf{x}$ without noise, but we only observe the signal up to time instant $t-1$ while we would like to predict $\mathbf{x}(t)$ (one-step-ahead prediction). This situation is characterized by the requirement that the Wiener filter:

$$\hat{\mathbf{x}}(t) = H(q)\mathbf{y}(t)$$

with $\mathbf{y} = \mathbf{x}$, should contain at least one step of time delay such that $\hat{\mathbf{x}}(t)$ is not dependent on $\mathbf{x}(t)$ and its past, but only on $\mathbf{x}(t-1)$ and its past.
Applying the formulas for the IIR Wiener filter, it follows that the optimal filter is given by

$$H(z) = \frac{1}{\Phi_{\mathbf{y}}^{+}(z)} \left[\frac{\Phi_{\mathbf{xy}}(z)}{\Phi_{\mathbf{y}}^{-}(z)} \right]_{++} \tag{6.59}$$

where the operation $[\cdot]_{++}$ now has to be interpreted as taking the strictly proper (or strictly causal) part of the function. This means that a constant term z^0 is not allowed, and all powers of z are required to be negative. As a result there is at least one time delay in the considered transfer, and this is necessary to have a (one-step-ahead) prediction.
Substituting the corresponding expressions for the spectral densities leads to

$$\begin{aligned} H(z) &= \frac{1}{W(z)\sigma} \cdot \left[\frac{W(z)W(1/z)\sigma^2}{W(1/z)\sigma} \right]_{++} && (6.60) \\ &= \frac{1}{W(z)\sigma} \cdot [\sigma W(z)]_{++} && (6.61) \\ &= \frac{1}{W(z)\sigma} \cdot \sigma[W(z) - 1] && (6.62) \\ &= 1 - W(z)^{-1}. && (6.63) \end{aligned}$$

Because of the fact that $W(z)$ is monic and minimum-phase, its inverse will also be monic and minimum phase, and consequently $H(z)$ will have a series expansion of the form $H(z) = h_1 z^{-1} + h_2 z^{-2} + \cdots$, therefore being a stricly causal system that predicts $\mathbf{x}(t)$ on the basis of measurements up to $\mathbf{x}(t-1)$.

6.8 Discussion

In this chapter we have discussed and presented optimal filters for reconstructing information-carrying data on the basis of corrupted measurements. The optimal filters are widely used in many parts of science, communication networks, optics, image and signal processing, and form the basis of the area which can be denoted as model-based measurement. The optimal Wiener filters can also be rewritten into a time-recursive form, where estimates at time instant t are written as a function (an update) of the estimates at time $t-1$. This form leads to the celebrated Kalman filter.
The optimal filters presented in this chapter require knowledge of spectral densities and cross-spectral densities of measurable signals. The estimation of these density functions on the basis of measurable data is the topic of the next chapter.

Chapter 7

Estimation of correlation functions and power spectra

The estimation of auto- and cross-correlation functions, as well as the related power spectral density functions, is the subject of this chapter. An analysis is made of basic nonparametric estimation methods.

7.1 Introduction

Estimating the power spectral density functions of stationary stochastic processes on the basis of measurement data is one of the crucial topics in signal analysis. Besides the intrinsic importance of characterizing the spectral properties of stochastic processes, it has been shown in the previous chapter that knowledge of auto- and cross spectral density functions is essential in the design of optimal (Wiener) filters. In this chapter spectral properties of stochastic processes will be estimated from data, where it is typically assumed hat one single data sequence is available for estimation. It is also assumed that the underlying stochastic process is stationary, implying that the statistical properties do not change with time. As a result (very) long data sequences can be used to estimate the statistical properties of the stochastic process. If this assumption is not justified, i.e. if the process has properties that may change with time, these changes have to be detected and/or relatively short data sequences can only be used for "stationary" estimation.

There are several options for estimation of power spectral density functions. In the parametric or model-based approach, parameterized models are used to estimate smooth spectra by determining a limited number of (unknown) parameters from the data. In this chapter attention is restricted to the non-parametric (or model-free) approach, where spectral density functions are considered graphical curves that, as such, are to be estimated from data.

Since power spectral density functions are Fourier transformed versions of correlation functions, the estimation of both concepts is very closely related to each other. For that reason, the estimation of correlation and covariance functions will be discussed in this chapter also.

7.2 Mean value estimation

Estimation of the mean value of a stationary stochastic process has been discussed in Example 5.1. There it was shown that for a stationary stochastic process $\mathbf{x}(t)$ the sample average:

$$\bar{\mathbf{x}}_N := \frac{1}{N}\sum_{i=0}^{N-1}\mathbf{x}(i) \tag{7.1}$$

is an unbiased estimator of the mean value $\mu_{\mathbf{x}}$, having a variance:

$$var(\bar{\mathbf{x}}_N) = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} C_{\mathbf{x}}(|i-j|). \tag{7.2}$$

As a result the sample average is a consistent estimator of the mean value of $\mathbf{x}$ under the condition that $C_{\mathbf{x}}(k)$ converges to 0 sufficiently fast for growing values of k.

7.3 Estimation of correlation functions

7.3.1 The sample autocorrelation function - unbiased version

On the basis of an available data sequence $\{\mathbf{x}(t)\}_{t=0,\cdots N-1}$ a straightforward choice for an estimator of the autocorrelation function of $\mathbf{x}$ is the sample autocorrelation

$$\breve{R}_{\mathbf{x}}(k) = \frac{1}{N-k}\sum_{t=k}^{N-1}\mathbf{x}(t)\mathbf{x}(t-k),\ k=0,1,\cdots N-1. \tag{7.3}$$

For $k \geq N$ the sample autocorrelation is defined to be equal to 0.
For negative values of k, use is made of the even property of autocorrelation functions, i.e. $R_{\mathbf{x}}(k) = R_{\mathbf{x}}(-k)$, and therefore for $k<0$, $\breve{R}_{\mathbf{x}}(k)$ is defined by $\breve{R}_{\mathbf{x}}(k) = \breve{R}_{\mathbf{x}}(-k)$.
When analyzing the properties of $\breve{R}_{\mathbf{x}}(k)$ attention will be restricted to the situation $0 \leq k \leq N-1$. This is only for ease of notation.

Bias of $\breve{R}_{\mathbf{x}}(k)$

Since

$$\mathbb{E}[\breve{R}_{\mathbf{x}}(k)] = \frac{1}{N-k}\sum_{t=k}^{N-1}\mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-k)] = R_{\mathbf{x}}(k)$$

the estimator is unbiased for any finite value of N.

Variance of $\breve{R}_{\mathbf{x}}(k)$

The variance of the estimator can be written as

$$var(\breve{R}_{\mathbf{x}}(k)) = \mathbb{E}[(\breve{R}_{\mathbf{x}}(k) - R_{\mathbf{x}}(k))^2] \tag{7.4}$$

which under assumption of a Gaussian distribution function of $\mathbf{x}$ can be shown to be given by (see Appendix for a derivation):

$$var(\breve{R}_{\mathbf{x}}(k)) = \frac{1}{N-|k|}\sum_{m=-N+|k|+1}^{N-1-|k|}(1-\frac{|m|}{N-|k|})\left[R_{\mathbf{x}}(m-k)R_{\mathbf{x}}(m+k)+R_{\mathbf{x}}^2(m)\right] \tag{7.5}$$

for any $|k| < N$. For non-Gaussian processes this expression is a good approximation to the variance for $|k| << N$ (Bartlett, 1946).
For a Gaussian white noise process with variance $\sigma_{\mathbf{x}}^2$ the expression reduces to

$$var(\breve{R}_{\mathbf{x}}(k)) = \begin{cases} \frac{2}{N}\sigma_{\mathbf{x}}^4 & k = 0 \\ \frac{1}{N-|k|}\sigma_{\mathbf{x}}^4 & k \neq 0 \end{cases} \tag{7.6}$$

Discussion

The considered estimator is unbiased for any value of N. Additionally it follows from (7.5) that for any fixed value of k, the variance of $\breve{R}_{\mathbf{x}}(k)$ tends to 0 for $N \to \infty$. Therefore the estimator is consistent for any value of k. On the other hand, the variance becomes bigger for increasing values of k, leading to really large values if k approaches N. This is understandable by noting that the number of terms in the summation (7.3) decreases for increasing values of k, as illustrated in Figure 7.1.

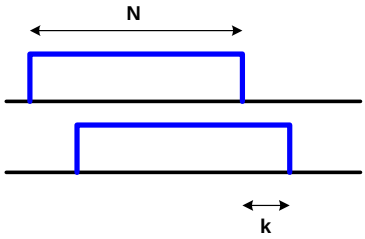

Figure 7.1: Number of terms in calculation of $\hat{R}_{\mathbf{x}}(k)$ decreases for increasing values of k.

A second remark that has to be made is that it is often not very attractive to calculate the convolutional time domain expression (7.3) directly. This will be commented upon later.
It has been analyzed in section 3.7.2 that not every sequence of numbers can be an autocorrelation function. Autocorrelation functions are characterized by the fact that they lead to a positive semi-definite correlation matrix. The considered estimator $\breve{R}_{\mathbf{x}}(k)$ has one major problem in the sense that this estimator is not guaranteed to lead to an appropriate autocorrelation function.
An example illustrating this effect can be considered by evaluating the 3-point measurement data $x(k)$, $k = 0, \cdots 2$ as indicated in Figure 7.2(left), leading to a sample autocorrelation function indicated in Figure 7.2(right).

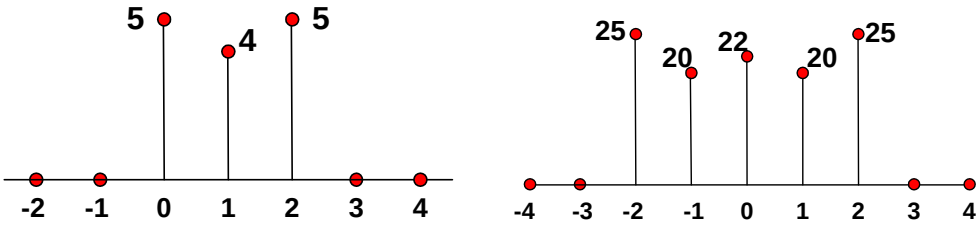

Figure 7.2: Time sequence $x(k)$ (left) and related sample autocorrelation function $\breve{R}_{\mathbf{x}}(k)$ (right).

Since $\breve{R}_{\mathbf{x}}(0) < \breve{R}_{\mathbf{x}}(3)$, the sequence $\{\breve{R}_{\mathbf{x}}(k)\}$ does not satisfy the basic requirements of an autocorrelation function, as it violates condition 3 in section 3.7.2. A direct consequence of this is that the related spectral density function estimator:

$$\breve{\Phi}_{\mathbf{x}}(\omega) := \sum_{k=-\infty}^{\infty} \breve{R}_{\mathbf{x}}(k) e^{-i\omega k} \tag{7.7}$$

is not guaranteed to be positive for all frequencies.

7.3.2 The sample autocorrelation function - biased version

As an alternative to the estimator considered in the previous subsection, we consider the following expression:

$$\hat{R}_{\mathbf{x}}(k) = \frac{1}{N} \sum_{t=k}^{N-1} \mathbf{x}(t)\mathbf{x}(t-k) \tag{7.8}$$

Following a similar analysis as in the previous section, it follows that

$$\mathbb{E}[\hat{R}_{\mathbf{x}}(k)] = \frac{N-|k|}{N} R_{\mathbf{x}}(k) \tag{7.9}$$

$$var\{\hat{R}_{\mathbf{x}}(k)\} = \frac{1}{N} \sum_{m=-N+|k|+1}^{N-1-|k|} (1 - \frac{|k|+|m|}{N}) \left[R_{\mathbf{x}}(m-k)R_{\mathbf{x}}(m+k) + R_{\mathbf{x}}^2(m)\right] \tag{7.10}$$

where a Gaussian distribution is assumed for the variance expression. For white noise processes the variance expression simplifies to

$$var\{\hat{R}_{\mathbf{x}}(k)\} = \begin{cases} \frac{2}{N}\sigma_{\mathbf{x}}^4 & k = 0 \\ \frac{1}{N}(1 - \frac{|k|}{N})\sigma_{\mathbf{x}}^4 & k \neq 0. \end{cases} \tag{7.11}$$

This alternative estimator is not unbiased. From (7.9) it follows that the bias disappears only for $N \to \infty$. Therefore the estimator is asymptotically unbiased for a fixed value of k. The variance is clearly different from the previous situation; now the variance of the estimator does not increase for increasing values of k, but it essentially becomes dependent on $1/N$.

This alternative estimator will appear to be attractive when considering it as a basis for the estimation of the spectral density function; it is guaranteed to deliver a positive semi-definite correlation matrix, so that consequently the Fourier transform of $\hat{R}_{\mathbf{x}}(k)$ will be a power spectral density function estimate that is non-negative for all frequencies. Denoting this Fourier Transform as

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \sum_{k} \hat{R}_{\mathbf{x}}(k) e^{-i\omega k}$$

we can use the result of Lemma 2A.3 to show that

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \frac{1}{N} |\mathbf{X}_N(\omega)|^2 \tag{7.12}$$

i.e., $\hat{R}_{\mathbf{x}}(k)$ constitutes a Fourier transform pair with $\frac{1}{N}|\mathbf{X}_N(\omega)|^2$, with $\mathbf{X}_N(\omega)$ the DFT of the stochastic process $\mathbf{x}$ over the time interval $[0, N-1]$:

$$\mathbf{X}_N(\omega) := \sum_{k=0}^{N-1} \mathbf{x}(k) e^{-i\omega k}. \tag{7.13}$$

Consequently $\frac{1}{N}|\mathbf{X}_N(\omega)|^2$, being the Fourier transform of $\hat{R}_{\mathbf{x}}$, will act as the related power spectral density estimator which is non-negative by construction.

Example 7.1 For a first order autoregressive process:

$$\mathbf{x}(t) = a \cdot \mathbf{x}(t-1) + \mathbf{e}(t) \tag{7.14}$$

with $\mathbf{e}$ a unit variance white noise process, autocorrelation function estimates $\hat{R}_{\mathbf{x}}(k)$ are calculated for 50 different experiments, each of length N. Results are depicted in Figure 7.3 for $a = 0.8$ and for $N = 128$ (left plot) and $N = 1024$ (right plot).
It can be observed that even for large data sets ($N = 1024$) the variance of the estimator is still substantial. Although 'on average' the estimate is quite close to the exact value, the large variance induces a serious lack of reliability if one has only one data sequence available. Note also that for higher values of k, as e.g., $k > 20$, all contributions in $\hat{R}_{\mathbf{x}}(k)$ are actually due to a variance contribution and not to any structural phenomenon, as the contribution of the exact autocorrelation is negligible in this region.

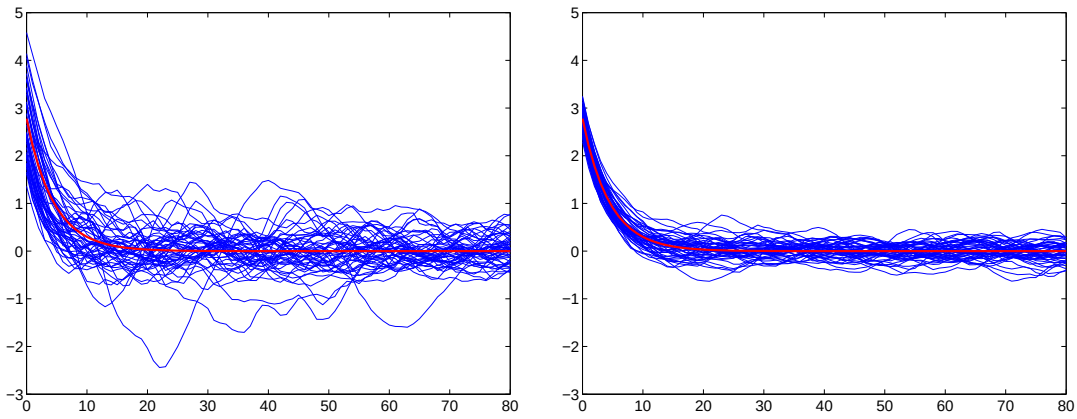

Figure 7.3: Estimates $\hat{R}_{\mathbf{x}}(k)$ of autocorrelation function for 50 realizations of data resulting from a first order AR process (blue tiny curves), and exact autocorrelation function (red solid curve) for $N = 128$ (left) and $N = 1024$ (right).

7.3.3 The sample cross-correlation function

For estimating cross-correlation functions the analysis follows similar lines as for estimating autocorrelation functions. The sample cross-correlation function is given by

$$\hat{R}_{\mathbf{yx}}(k) = \frac{1}{N}\sum_{t=k}^{N-1} \mathbf{y}(t)\mathbf{x}(t-k) \quad 0 \le k \le N-1, \tag{7.15}$$

$$= \frac{1}{N}\sum_{t=0}^{N-1-|k|} \mathbf{y}(t)\mathbf{x}(t-k) \quad -(N-1) \le k \le 0. \tag{7.16}$$

Note that because the cross-correlation function is not symmetric, the terms for either negative and positive values of k have to be calculated separately.

The sample cross-correlation function considered here is a biased estimator of the cross-correlation function, since

$$\mathbb{E}[\hat{R}_{\mathbf{yx}}(k)] = \frac{N-|k|}{N} R_{\mathbf{yx}}(k) \tag{7.17}$$

Its variance properties are similar to those of the autocorrelation function estimator, and are, for jointly Gaussian and stationary processes $\mathbf{y}$ and $\mathbf{x}$, expressed by

$$var\{\hat{R}_{\mathbf{yx}}(k)\} = \frac{1}{N} \sum_{m=-N+|k|+1}^{N-1-|k|} (1 - \frac{|k|+|m|}{N}) \left[R_{\mathbf{xy}}(m-k)R_{\mathbf{yx}}(m+k) + R_{\mathbf{x}}(m)R_{\mathbf{y}}(m)\right] \tag{7.18}$$

and will therefore be essentially inversely proportional to N for small values of k. Expression (7.18) is derived in the appendix of this chapter.

Where the sample autocorrelation $\hat{R}_{\mathbf{x}}(k)$ can be calculated as the inverse DFT of $\frac{1}{N}|\mathbf{X}_N(\omega)|^2$, the sample cross-correlation $\hat{R}_{\mathbf{yx}}(k)$ constitutes the Fourier transform pair:

$$\hat{R}_{\mathbf{yx}}(k) \leftrightarrow \frac{1}{N}\mathbf{Y}_N(\omega) \cdot \mathbf{X}_N(\omega)^* \tag{7.19}$$

and therefore the sample cross-correlation function can be calculated as the inverse DFT of the righthand expression. Since Fourier transforms can be effectively calculated using Fast Fourier Transform (FFT) algorithms, this is an effective computational procedure.

7.4 Spectral estimation

7.4.1 Introduction

The power spectral density function constitutes the distribution of the power of a stochastic process over frequencies. It satisfies

$$\mathbb{E}[\mathbf{x}^2(t)] = \frac{1}{2\pi}\int_{-\pi}^{\pi} \Phi_{\mathbf{x}}(\omega)d\omega.$$

Since $\Phi_{\mathbf{x}}(\omega)$ is the Fourier transform of $R_{\mathbf{x}}(k)$ it is a natural choice to construct an estimator $\hat{\Phi}_{\mathbf{x}}(\omega)$ of $\Phi_{\mathbf{x}}(\omega)$ by the Fourier transform of an estimator $\breve{R}_{\mathbf{x}}(k)$ or $\hat{R}_{\mathbf{x}}(k)$ of $R_{\mathbf{x}}(k)$.
As motivated in the previous section the unbiased autocorrelation function estimator $\breve{R}_{\mathbf{x}}(k)$ does not lead to a Fourier transform that is nonnegative for all ω. Therefore, when considering estimators of the power spectral density function, attention will be focussed on the Fourier transform of the biased autocorrelation function estimator $\hat{R}_{\mathbf{x}}(k)$.

7.4.2 Periodogram estimator of spectral density function

Periodogram estimator

The autocorrelation function estimator discussed in section 7.3.2 generates a related estimator of the spectral density function by

$$\hat{\Phi}_{\mathbf{x}}(\omega) := \sum_k \hat{R}_{\mathbf{x}}(k)e^{-i\omega k} \tag{7.20}$$

$$= \frac{1}{N}|\mathbf{X}_N(\omega)|^2. \tag{7.21}$$

The resulting estimator is referred to as the *periodogram* estimator of the spectral density function. It is simply obtained by a finite length DFT of the time sequence $\mathbf{x}(t)$.

Bias of the periodogram

The autocorrelation function estimator $\hat{R}_{\mathbf{x}}(k)$ is a biased estimator of $R_{\mathbf{x}}(k)$. This of course has consequences for the related spectral density function estimator.
For an analysis of the bias properties of $\hat{\Phi}_{\mathbf{x}}(\omega)$ we write

$$\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] = \mathbb{E}\left\{ \sum_{k=-\infty}^{\infty} \hat{R}_{\mathbf{x}}(k) e^{-i\omega k} \right\}$$

In this expression one can substitute

$$\hat{R}_{\mathbf{x}}(k) = \frac{1}{N} \sum_{t=-\infty}^{\infty} d(t)\mathbf{x}(t)d(t-k)\mathbf{x}(t-k) \tag{7.22}$$

where $d(t)$ is defined as the function: $d(t) = 1$ for $t \in [0, N-1]$ and $d(t) = 0$ elsewhere. It selects those components of $\mathbf{x}(t)$ that are part of the measurement interval $[0, N-1]$. As a result, (7.22) is a valid expression for $\hat{R}_{\mathbf{x}}(k)$ for both positive and negative values of k.
Then

$$\begin{aligned} \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] &= \sum_{k=-\infty}^{\infty} R_{\mathbf{x}}(k) e^{-i\omega k} \cdot \underbrace{\sum_{t=-\infty}^{\infty} \frac{1}{N} d(t)d(t-k)}_{window\ q_N(k)} \\ &= \sum_{k=-\infty}^{\infty} q_N(k) R_{\mathbf{x}}(k) e^{-i\omega k} \end{aligned} \tag{7.23}$$

with

$$q_N(k) = \begin{cases} 1 - \frac{|k|}{N} & |k| < N \\ 0 & \text{elsewhere} \end{cases}$$

The particular window function $q_N(k)$ is referred to as the Bartlett window and is sketched in figure 7.4 (left plot).
According to (7.23) the autocorrelation function $R_{\mathbf{x}}(k)$ is multiplied by $q_N(k)$ before being transformed to the Fourier domain. Since multiplication in the time domain induces a convolution in the Fourier domain, it follows that

$$\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] = \frac{1}{2\pi} \cdot Q_N(\omega) \star \Phi_{\mathbf{x}}(\omega) \tag{7.24}$$

or equivalently

$$\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{\mathbf{x}}(\xi) Q_N(\omega - \xi) d\xi \tag{7.25}$$

with $Q_N(\omega)$ the Fourier transform of the Bartlett window function $q_N(k)$, i.e.

$$Q_N(\omega) = \frac{1}{N} \left[\frac{sin(\omega N/2)}{sin(\omega/2)} \right]^2 \tag{7.26}$$

which is being sketched in the right plot of Figure 7.4.

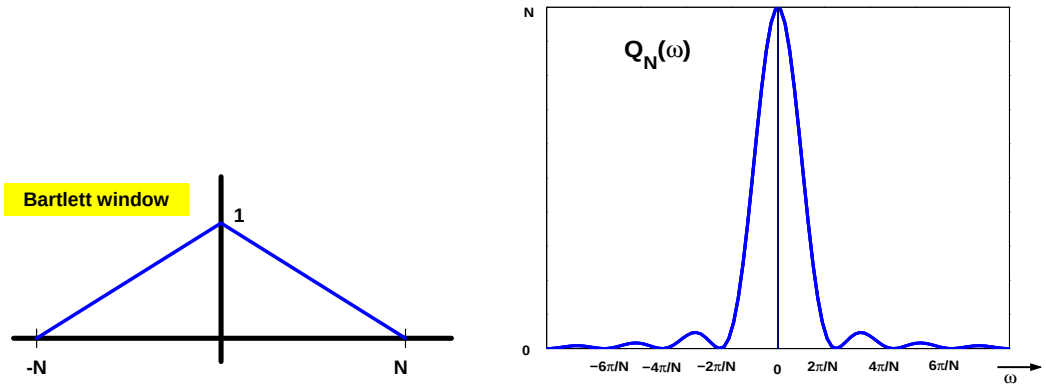

Figure 7.4: Bartlett window $q_N(k)$ (left) and its Fourier transform $Q_N(\omega)$ (right).

> The mean value of the periodogram estimator $\hat{\Phi}_{\mathbf{x}}(\omega)$ of a spectral density function $\Phi_{\mathbf{x}}(\omega)$ is given by the convolution (in the frequency domain) of the exact spectral density function $\Phi_{\mathbf{x}}(\omega)$ with the window function $Q_N(\omega)$. As a result the mean value is a smoothed version of the original density.

Convolution of $\Phi_{\mathbf{x}}(\omega)$ with $Q_N(\omega)$ will have a "smoothing" effect on the estimate. Sharp peaks that may be present in $\Phi_{\mathbf{x}}(\omega)$ will be smeared out over adjacent frequencies, thus leading to an effect which is often referred to as *spectral leakage*.
The leakage effect is an effect of the finite sample size N. If N increases the window function $Q_N(\omega)$ will become more narrow and the bias will disappear. Eventually, for increasing values of N, $Q_N(\omega)$ will approach a Dirac function, leading to $\lim_{N\to\infty} \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] = \Phi_{\mathbf{x}}(\omega)$, i.e. an asymptotically unbiased estimator.

Note that the bias effect considered here is also related to the resolution that can be achieved in the frequency domain. The sequence of N data points available in the time domain is transformed to the frequency domain, and delivers there also a sequence of N -uniquely determined - frequency points that are the basis for the back (inverse Fourier) transformation. The N independent points in the frequency domain are actually given by the equidistant grid $\omega = k\pi/N$ with $k = 1, \cdots N$. Therefore π/N is the "smallest" frequency difference that can be observed on the basis of the N data points, and therefore determines a limit on the resolution that results from observing N data points. If two sinusoids are present in the signal (two delta functions in the spectral density) that are within a frequency distance of π/N, one can not expect to be able to distinguish the two separate sinusoids from N observations in time without additional a priori information.

As an example we consider the situation of a single sinusoid that is observed under an additive white noise disturbance:

$$\mathbf{x}(t) = Acos(\omega_0 t + \phi) + \mathbf{v}(t) \tag{7.27}$$

with ϕ a random phase, uniformly distributed in the interval $[0, 2\pi]$, and $\mathbf{v}$ a stationary white noise process with variance $\sigma^2_{\mathbf{v}}$, while ϕ and $\mathbf{v}(t)$ are independent for all t. Under these conditions the process $\mathbf{x}(t)$ is stationary.

The power spectral density function of $\mathbf{x}$ is given by

$$\Phi_{\mathbf{x}}(\omega) = \frac{\pi}{2} A^2[\delta(\omega - \omega_0) + \delta(\omega + \omega_0)] + \sigma_{\mathbf{v}}^2 \tag{7.28}$$

This spectral density is sketched in figure 7.5(left) for $\omega_0 = \pi/4$.

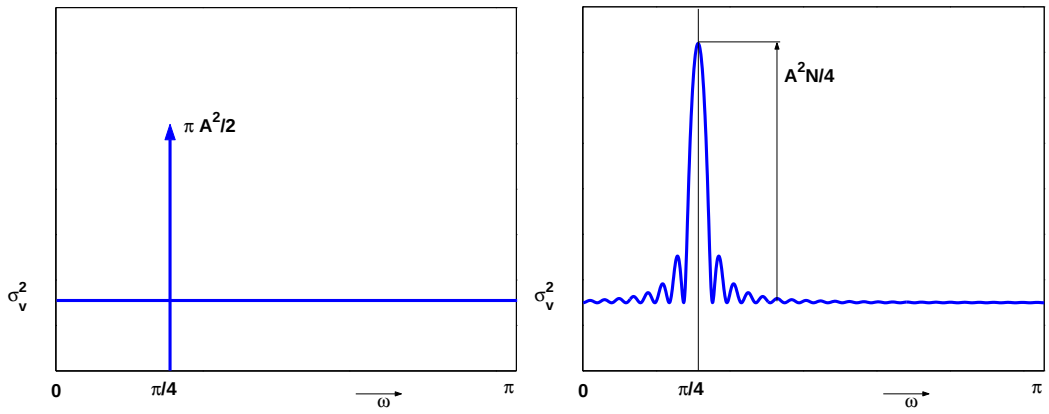

Figure 7.5: Spectral density function of the two components of $\Phi_{\mathbf{x}}(\omega)$ of sinusoid in white noise (left) and mean value of the periodogram estimate $\hat{\Phi}_{\mathbf{x}}(\omega)$ on the basis of $N = 64$ data points.

The mean value of the periodogram estimator is

$$\begin{aligned} \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] &= \frac{1}{2\pi} \cdot \Phi_{\mathbf{x}}(\omega) \star Q_N(\omega) && (7.29) \\ &= \sigma_{\mathbf{v}}^2 + \frac{1}{4} A^2[Q_N(\omega - \omega_0) + Q_N(\omega + \omega_0)] && (7.30) \end{aligned}$$

and this is sketched for in figure 7.5(right). The smoothing effect that the delta function of the sinusoid undergoes due to the Bartlett window is clearly visible in the periodogram.

As a second example we consider 100 periodogram estimates of the considered process $\mathbf{x}$, based on 100 different realizations of the noise, and the average value of the 100 estimates. Results for this test are depicted in figure 7.6 for data sets of length $N = 64$ (upper pictures) and $N = 256$ (lower pictures).
The important difference between the situations of $N = 64$ and $N = 256$ is that in the latter case ($N = 256$) the frequency resolution is higher, while in the former situation ($N = 64$) the smoothing effect is stronger.

Variance of the periodogram

For analysis of the variance properties of the periodogram estimator of the spectral density function first an analysis is made for a Gaussian distributed white noise process. In a second stage the expressions can then be generalized to the situation of Gaussian non-white processes.
The main result for Gaussian white noise processes is formulated first:

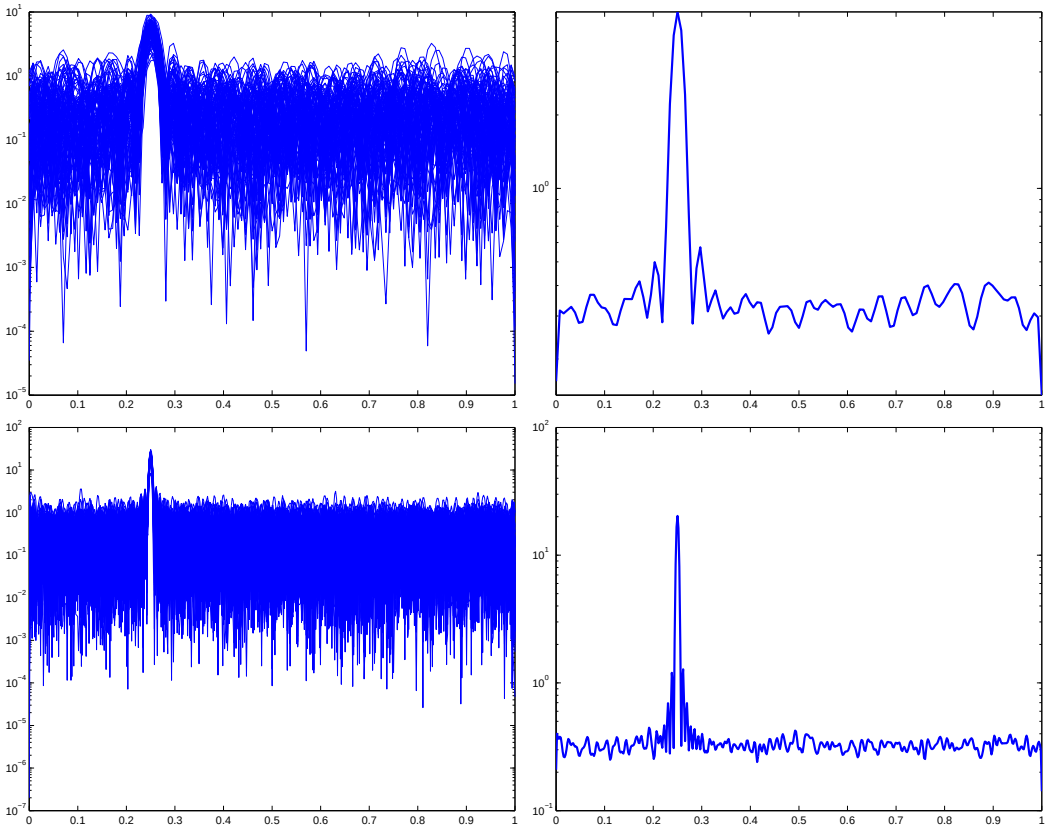

Figure 7.6: Periodogram estimates of a sinusoid in white noise. Overlay of 100 periodograms using $N = 64$ data (left upper), and the periodogram average (right upper). Overlay of 100 periodograms using $N = 256$ data (left lower), and the periodogram average (right lower). Frequency axes are in π rad/sec.

For a Gaussian white noise process $\mathbf{x}$ with variance $\sigma_{\mathbf{x}}^2$ the periodogram estimator $\hat{\Phi}_{\mathbf{x}}(\omega)$ satisfies

$$\begin{aligned} var\{\hat{\Phi}_{\mathbf{x}}(\omega)\} &= \sigma_{\mathbf{x}}^4 \qquad \omega = \frac{2\pi k}{N},\ k = 1, \cdots N-1 \\ &= 2\sigma_{\mathbf{x}}^4 \qquad \omega = 0, \pi \end{aligned}$$

$$Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} = 0, \qquad \omega_j = \frac{2\pi k_j}{N}, j = 1, 2; \omega_1 \neq \pm\omega_2$$

Since the variance does not tend to zero for increasing N, the periodogram estimator is not consistent.

The derivation of these expressions is added in the appendix of this chapter.
Basic result is that the variance of the periodogram is equal to the squared spectral density of $\mathbf{x}$. It does not tend to zero for increasing values of N, showing that the periodogram estimator is not consistent.

Extension to non-white Gaussian processes

For a non-white process $\mathbf{x}$ it follows that there always exists a minimum-phase filter $H(z)$

such that

$$\Phi_{\mathbf{x}}(\omega) = |H(e^{i\omega})|^2 \cdot \Phi_{\mathbf{e}}(\omega)$$

with $\mathbf{e}$ a unit variance white noise process.
If N is large compared to the length of the pulse response of $H(z)$, then it will hold approximately that

$$\mathbf{X}_N(\omega) \sim H(e^{i\omega}) \cdot \mathbf{E}_N(\omega)$$

and consequently

$$|\mathbf{X}_N(\omega)|^2 \sim |H(e^{i\omega})|^2 \cdot |\mathbf{E}_N(\omega)|^2$$

so that

$$\begin{aligned} Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} &= |H(e^{i\omega_1})|^2 |H(e^{i\omega_2})|^2 \cdot Cov\{\hat{\Phi}_{\mathbf{e}}(\omega_1), \hat{\Phi}_{\mathbf{e}}(\omega_2)\} \\ &= \Phi_{\mathbf{x}}(\omega_1)\Phi_{\mathbf{x}}(\omega_2) \cdot Cov\{\hat{\Phi}_{\mathbf{e}}(\omega_1), \hat{\Phi}_{\mathbf{e}}(\omega_2)\} \end{aligned} \tag{7.31}$$

This implies that the variance and covariance of a Gaussian process can simply be obtained from the earlier expressions for white processes by premultiplication of the covariance expression by $\Phi_{\mathbf{x}}(\omega_1)\Phi_{\mathbf{x}}(\omega_2)$, and replacing $\sigma_{\mathbf{x}}^2$ by $\sigma_{\mathbf{e}}^2 = 1$. The following result then appears.

For a Gaussian stochastic process $\mathbf{x}$ the periodogram estimator $\hat{\Phi}_{\mathbf{x}}(\omega)$ satisfies, for sufficiently large N,

$$\begin{aligned} var\{\hat{\Phi}_{\mathbf{x}}(\omega)\} &= \Phi_{\mathbf{x}}(\omega)^2 \qquad \omega = \frac{2\pi k}{N},\ k = 1, \cdots N-1 \\ &= 2\Phi_{\mathbf{x}}(\omega)^2 \qquad \omega = 0, \pi \\ Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} &= 0, \qquad \omega_j = \frac{2\pi k_j}{N}, j = 1, 2; \omega_1 \neq \pm\omega_2 \end{aligned}$$

Since the variance does not tend to zero for increasing N, the periodogram estimator is not consistent.

So the variance of the periodogram estimator at a particular frequency is equal to the squared valued of the spectral density itself.

Example 7.2 The same stochastic process is considered as in Example 7.1, i.e. a first order AR process. Five data sets of each N data points are used to generate 5 periodogram estimates. Figure 7.7 depicts these estimates for data sets of length $N = 256$ (upper pictures) and $N = 512$ (lower pictures), together with the exact (calculated) power spectral density function. Note that the variance of the estimator is substantial. In the right plots of Figure 7.7 the average of the 5 periodogram estimates is compared with the exact density function.

The characteristic phenomenon, as illustrated in Figure 7.7, is that periodogram estimates generally are erratic and non-smooth functions. If the number of data N increases this erratic character is not reduced.

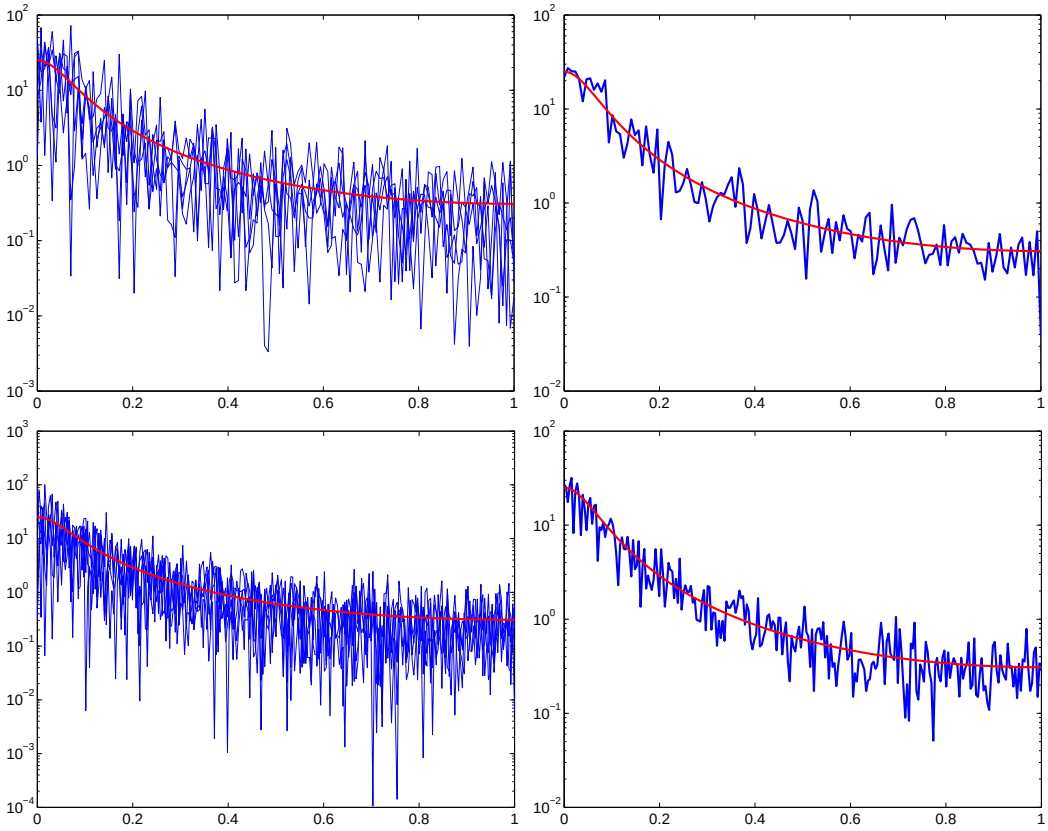

Figure 7.7: Periodogram estimate of 5 realizations of $N = 256$ (upper plots) and $N = 512$ (lower plots) datapoints from an AR(1) process, compared with the exact power spectral density (left plots); average of the 5 periodogram estimates compared with the exact power spectral density (right plots); frequency axes are in π rad/sec.

> The periodogram estimator of a spectral density function is generally an erratic function of frequency. This erratic character does not diminish if the number of data N increases; this is reflected in a non-vanishing variance for each frequency. Rather the frequency resolution increases with increasing N.

Several options for improving the periodogram estimators will be discussed in subsection 7.4.4.

7.4.3 Extension to cross-spectral densities

For a cross-spectrum $\Phi_{\mathbf{xy}}$ where

$$\mathbf{x}(t) = H_x(q)\mathbf{e}(t) \tag{7.32}$$
$$\mathbf{y}(t) = H_y(q)\mathbf{e}(t) \tag{7.33}$$

and $\mathbf{e}$ a unit-variance white noise process, it follows that $\hat{\Phi}_{\mathbf{yx}}(\omega) = \frac{1}{N}\mathbf{Y}_N(\omega)\mathbf{X}_N(\omega)^*$, while $\Phi_{\mathbf{yx}}(\omega) = H_y(e^{i\omega})H_x(e^{i\omega})^*$.

With

$$\mathbf{X}_N(\omega) \sim H_x(e^{i\omega})\mathbf{E}_N(\omega) \tag{7.34}$$
$$\mathbf{Y}_N(\omega) \sim H_y(e^{i\omega})\mathbf{E}_N(\omega) \tag{7.35}$$

it follows that

$$\mathbf{Y}_N(\omega)\mathbf{X}_N(\omega)^* \sim H_y(e^{i\omega})H_x(e^{i\omega})^*|\mathbf{E}_N(\omega)|^2$$

and so, similar as before[1]:

$$\begin{aligned} &Cov\{\hat{\Phi}_{\mathbf{yx}}(\omega_1), \hat{\Phi}_{\mathbf{yx}}(\omega_2)\} = \\ &= H_y(e^{i\omega_1})H_x(e^{i\omega_1})^* \cdot H_y(e^{i\omega_2})^* H_x(e^{i\omega_2}) \cdot Cov\{\Phi_{\mathbf{e}}(e^{i\omega_1}), \Phi_{\mathbf{e}}(e^{i\omega_2})\} \\ &= \Phi_{\mathbf{yx}}(\omega_1)\Phi_{\mathbf{yx}}(\omega_2)^* \cdot Cov\{\Phi_{\mathbf{e}}(\omega_1), \Phi_{\mathbf{e}}(\omega_2)\} \end{aligned}$$

and for the variance:

$$var\{\hat{\Phi}_{\mathbf{yx}}(\omega)\} = |\Phi_{\mathbf{yx}}(\omega)|^2,$$

subject to the same frequency-dependent conditions as mentioned before:

$$\begin{aligned} var\{\hat{\Phi}_{\mathbf{yx}}(\omega)\} &= |\Phi_{\mathbf{yx}}(\omega)|^2 \quad \text{for } \omega = 2\pi k/N, k \neq 0, N/2 \\ &= 2|\Phi_{\mathbf{yx}}(\omega)|^2 \quad \text{for } \omega = 0, \pi \end{aligned} \tag{7.36}$$

7.4.4 Smoothed periodogram estimators

As shown in the previous sections the periodogram estimator for spectral density functions is not really ideal. Although it is asymptotically unbiased at every frequency, its variance is substantial and does not decay to zero for increasing number of data.

Lag windows and frequency windows

The fact that periodogram estimators suffer from this high variance can be understood by considering that the periodogram can be interpreted as the Fourier transform of the sample autocorrelation function (7.8). The sample autocorrelation function contains a number of N terms with finite variance that all are taken into account when constructing the Fourier transform:

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \sum_{k=-\infty}^{\infty} \hat{R}_{\mathbf{x}}(k)e^{-i\omega k}.$$

This implies that for increasing N the number of to-be-estimated parameters increases, and that an increasing number of terms is involved in the construction of $\hat{\Phi}_{\mathbf{x}}(\omega)$.
We can circumvent this mechanism in the spectral estimate, by applying a so-called *lag window* to the sample autocorrelation function before applying the Fourier transform:

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \sum_{k=-\infty}^{\infty} w_\gamma(k) \cdot \hat{R}_{\mathbf{x}}(k)e^{-i\omega k}. \tag{7.37}$$

Here, $w_\gamma(k)$ is a positive real-valued window function satisfying

$$w_\gamma(k) = 0 \quad |k| > \gamma \tag{7.38}$$

[1]Note that since the cross-spectral density is a complex-valued function its covariance involves a complex conjugate operation: $cov(a, b) = \mathbb{E}(a - \mathbb{E}a)(b - \mathbb{E}b)^*$

showing that $\gamma > 0$ is a variable that determines the width of the window. This lag window causes $w_\gamma(k)\hat{R}_{\mathbf{x}}(k)$ to be regularized to zero. The smaller the value of γ, the bigger the part of the sample autocorrelation function that is "smoothed out". The higher the value of γ, the less smoothing is taking place. A typical choice for a lag window is e.g. a rectangular form:

$$w_\gamma(k) = 1 \quad 0 \leq |k| \leq \gamma \tag{7.39}$$
$$= 0 \quad |k| > \gamma. \tag{7.40}$$

However more general choices are also possible, having a more smooth decay of the window towards zero. Three popular choices of windows are sketched in Figure 7.8 and characterized in Table 7.1. For an extensive list of windows the reader is referred to Jenkins and Watts (1968), Brillinger (1981) and Priestley (1981).

The three lag-windows are also sketched in Figure 7.8.

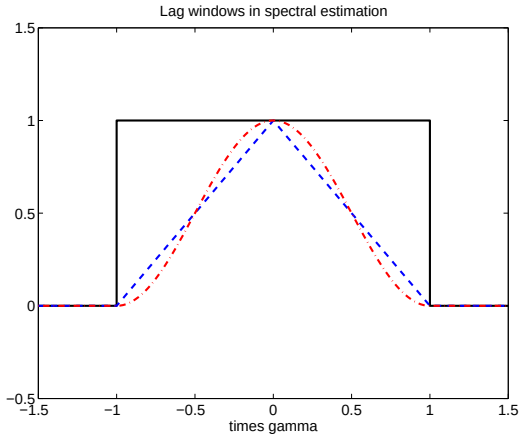

Figure 7.8: Lag-windows $w_\gamma(k)$; rectangular window (solid), Bartlett window (dashed), and Hamming window (dash-dotted).

Design rules for the choice of γ are:

- γ should be small compared to the number of data N, in order to guarantee enough smoothing operation;
- $|R_{\mathbf{x}}(k)| << \hat{R}_{\mathbf{x}}(0)$ for $k \geq \gamma$ in order to guarantee that interesting dynamics is not smoothed out.

The first point refers to a sufficient reduction of variance, whereas the second point refers to the avoidance of substantial bias.

The application of a lag window in the time domain, has a direct interpretation as a smoothing operation in the frequency domain. Using the fact that $\hat{R}_{\mathbf{x}}$ and $\hat{\Phi}_{\mathbf{x}}$ are related through Fourier transform, and using the fact that multiplication in the time-domain relates

to convolution in the frequency domain, the following derivation becomes trivial

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \sum_{k=-\infty}^{\infty} w_\gamma(k) \cdot \hat{R}_{\mathbf{x}}(k) e^{-i\omega k} \tag{7.41}$$

$$= \mathcal{F}\{w_\gamma(k) \cdot \hat{R}_{\mathbf{x}}(k)\} \tag{7.42}$$

$$= W_\gamma(\omega) \star \frac{1}{N} |\mathbf{X}_N(\omega)|^2 \tag{7.43}$$

$$= \frac{1}{2\pi} \int_{-\pi}^{\pi} W_\gamma(\xi - \omega) \frac{1}{N} |\mathbf{X}_N(\xi)|^2 d\xi \tag{7.44}$$

where the *frequency window* $W_\gamma(\omega)$ is the Fourier transform of the lag window $w_\gamma(k)$, i.e.

$$w_\gamma(k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} W_\gamma(\xi) e^{i\xi k} d\xi, \tag{7.45}$$

and where the $\star$ in (7.43) is the convolution operator.

	$2\pi \cdot W_\gamma(\omega)$	$w_\gamma(k),\ 0 \leq \lvert k \rvert \leq \gamma$
Rectangular	$\frac{sin(\gamma+\frac{1}{2})\omega}{sin(\omega/2)}$	1
Bartlett	$\frac{1}{\gamma}\left(\frac{sin\ \gamma\omega/2}{sin\ \omega/2}\right)^2$	$1 - \frac{k}{\gamma}$
Hamming	$\frac{1}{2}D_\gamma(\omega) + \frac{1}{4}D_\gamma(\omega - \frac{\pi}{\gamma})$ $+\frac{1}{4}D_\gamma(\omega + \frac{\pi}{\gamma})$, where $D_\gamma(\omega) = \frac{sin(\gamma+\frac{1}{2})\omega}{sin(\omega/2)}$	$\frac{1}{2}(1 + cos\frac{\pi k}{\gamma})$

Table 7.1: Windows for spectral estimation in time domain (right column) and frequency domain (left column).

As a result the introduction of a smoothing window could equally well have been motivated from the frequency domain perspective. The statistical properties of the periodogram estimator show that it has a nonvanishing variance at each frequency, but also that for different frequencies the covariance asymptotically tends to zero. This observation supports the introduction of an averaging operation in the frequency domain, where at each frequency ω the periodogram estimate $\hat{\Phi}_{\mathbf{x}}(\omega)$ is replaces by an (weighted) averaged value over estimates $\hat{\Phi}_{\mathbf{x}}(\xi)$ in the neighborhood of ω. The frequency range over which this average should take place should be dictated by the question whether the underlying real spectrum $\Phi_{\mathbf{x}}$ is constant over the considered range. This averaging operation in the frequency domain is -in mathematical terms- represented by a convolution operation operation with window $W_\gamma(\omega)$ where the width of W_γ determines the frequency range that is involved in the averaging. If the frequency window is chosen too narrow then hardly any smoothing will take place. If it is chosen too wide, then essential dynamic properties of the spectrum are averaged out, resulting in a substantial bias of the spectrum estimator.
The frequency-windows $W_\gamma(\omega)$, equivalent to the previously presented lag windows, are given in Table 7.1 also, and their characteristics are sketched in Figure 7.9. In this plot the

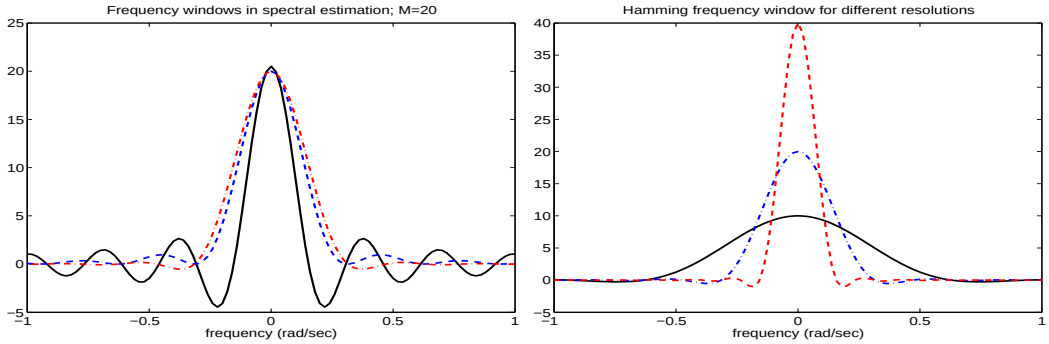

Figure 7.9: Frequency-windows in spectral estimation. Left: $W_\gamma(\xi)$, $\gamma = 20$ for rectangular window (solid), Bartlett window (dashed) and Hamming window (dash-dotted). Right: $W_\gamma(\xi)$ of Hamming window for $\gamma = 10$ (solid), $\gamma = 20$ (dash-dotted) and $\gamma = 40$ (dashed).

frequency window for the rectangular lag-window has been scaled by 0.5 in order to give it a similar amplitude as the other two windows. It is clearly illustrated in these figures how the several frequency-windows perform their smoothing operation. The higher the value of γ the narrower the window, and the less smoothing operation is performed when applying the window to the periodogram estimator, as in (7.44). If γ is chosen small, then the spectrum estimate is a smoothed version of the periodogram, where a wide frequency region is averaged out. Finding a correct choice of γ is sometimes a difficult task. Choosing γ too small may result in a spectral estimate in which relatively sharp peaks in the real spectrum have been smoothed out, and choosing γ too big will result in an erratic spectral estimate, having too high variance.

The smoothing operation of a Hamming window is illustrated in Figure 7.10, where Hamming windows of different widths are applied to data from the first order process of Examples 7.1 and 7.2.

The results for $\gamma = 8$ (first row of plots) show that this smoothing operation is too strong. A serious bias is introduced, in particular in the low frequency range. The smoothing with $\gamma = 16$ (second row of plots) is very accurate; the variance of the estimator is drastically reduced and there does not seem to be an apparent bias. The smoothing with $\gamma = 64$ leaves a considerable amount of variance left in the estimator.

Note that in this simulation example the best choice of smoothing operation can be evaluated by comparison with the exact spectral density function. This latter function is of course generally unknown in real-life problems, making the choice of appropriate smoothing operations rather subjective. Since the estimate obtained is very sensitive to the choice of smoothing operation, one should be cautious in drawing conclusions from the graphical results. In general it will be a trial and error procedure to arrive at an appropriate smoothing.

Periodogram averaging - Bartlett's method

Another approach to the estimation of signal spectra is by averaging periodogram estimates over several data segments. In this method one decomposes a signal of length N into r

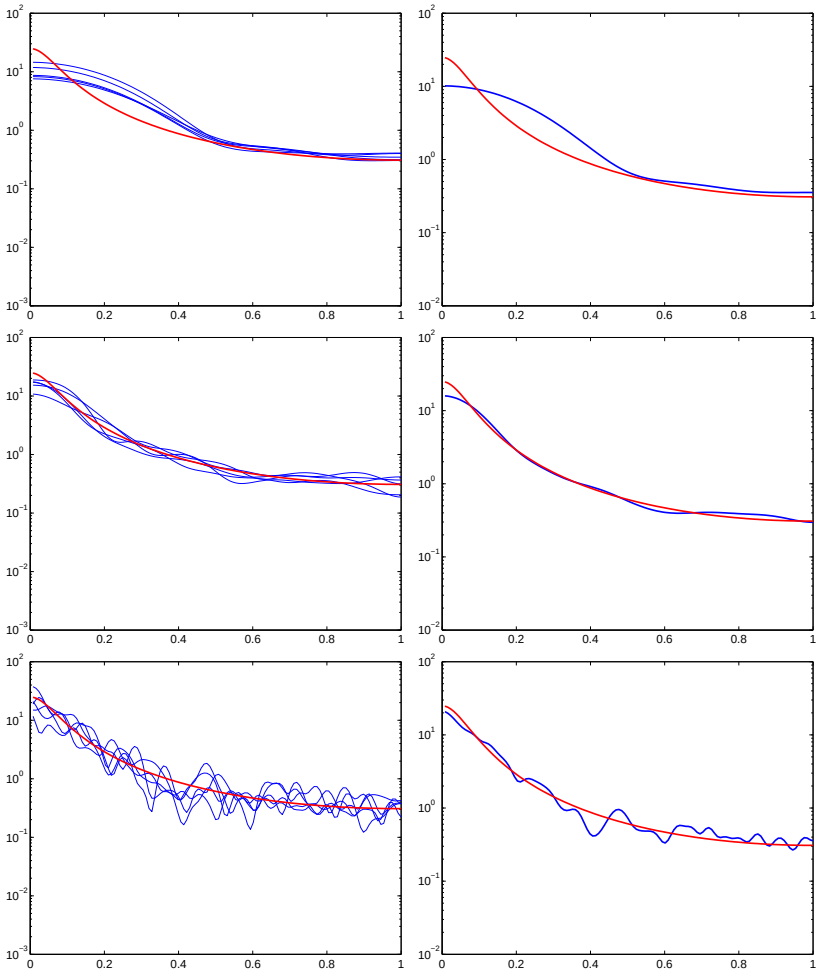

Figure 7.10: Smoothed periodogram estimates on the basis of $N = 256$ data points of a first order AR process. The left plots show 5 realizations of estimates together with the exact spectrum; the right plots show the average of the 5 realizations together with the exact spectrum. First row: $\gamma = 8$; second row: $\gamma = 16$; third row: $\gamma = 64$.

sequences of length N_0 and constructs a periodogram of each sequence:

$$\hat{\Phi}_{\mathbf{x},j}(\omega) = \frac{1}{N_0}|\mathbf{X}_{N_0,j}(\omega)|^2 \tag{7.46}$$

where j denotes the different sequences, $j = 1, \cdots r$.
A spectral estimate of the signal $\mathbf{x}$ is then obtained by averaging over the several periodograms:

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \frac{1}{r}\sum_{j=1}^{r}\hat{\Phi}_{\mathbf{x},j}(\omega). \tag{7.47}$$

This idea refers to the classical way of reducing the variance of an estimate by taking averages of several independent estimates. Independence of the several estimates will only be possible when the several data segments do not overlap. However independence is not guaranteed when choosing non-overlapping segments, due to possibly present dynamics in

the underlying process. In general the data segments will be chosen with a length N_0 that is a power of 2, in order to facilitate efficient calculation of the periodograms through Fast Fourier Transform.
Also in this method the conflicting aspects of reducing the variance and obtaining a high frequency resolution are present. Variance reduction is related to the number of averages r that is achieved, while the frequency resolution is related to the number of data samples in a data segment. In finding a satisfactory compromise between these choices, the use of overlapping data segments is also possible.

This method of periodogram smoothing is applied to the example process of Examples 7.1 and 7.2. In Figure 7.11 results are shown for data sets of length $N = 256$ that are split into 8 non-overlapping datasets of $N_0 = 32$ samples.

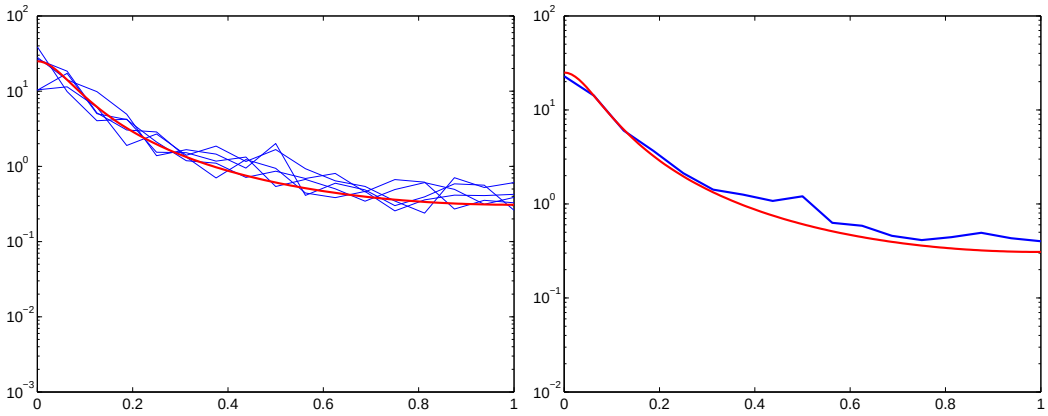

Figure 7.11: Periodogram estimates obtained by Bartlett's procedure of periodogram averaging applied to 8 non-overlapping data segments of $N_0 = 32$ of the AR(1) process (7.14). Periodogram estimates of 5 realizations of data of length $N = 256$, compared with the real spectral density function (left plot); average of 5 realizations compared with the real spectral density function (right plot). Frequency axes are in π rad/sec.

This method of periodogram averaging has some relationship with a windowing operation related to the windows discussed before. This can be observed from the following derivation. Let $\mathbf{x}_j(t)$ denote the signal in the j-th data segment, then

$$\frac{1}{N_0}|\mathbf{X}_{N_0,j}(\omega)|^2 = \sum_{k=-\infty}^{\infty} \hat{R}_{\mathbf{x}_j}(k)e^{-i\omega k} = \frac{1}{N_0}\sum_{k=-\infty}^{\infty}\sum_{t=0}^{N_0-1} \mathbf{x}_j(t)\mathbf{x}_j(t-k)e^{-i\omega k}, \tag{7.48}$$

using the convention that $\mathbf{x}_j(t) := 0$ outside the interval $[0, N_0 - 1]$.
As a result we can write

$$\frac{1}{r}\sum_{j=1}^{r}\frac{1}{N_0}|\mathbf{X}_{N_0,j}(\omega)|^2 = \frac{1}{r}\sum_{j=1}^{r}\sum_{k=-N_0}^{N_0}\frac{1}{N_0}\sum_{t=0}^{N_0-1} \mathbf{x}_j(t)\mathbf{x}_j(t-k)e^{-i\omega k} \tag{7.49}$$

where the interval for k has been limited to its allowed values: $[-N_0, N_0]$. When taking the expected value of this expression, and noting that

$$\mathbb{E}[\sum_{t=0}^{N_0-1} \mathbf{x}_j(t)\mathbf{x}_j(t-k)] = (N_0 - |k|)R_{\mathbf{x}_j}(k)$$

it follows that

$$\frac{1}{r}\sum_{j=1}^{r}\mathbb{E}[\hat{\Phi}_{\mathbf{x}_j}(\omega)] = \frac{1}{r}\sum_{j=1}^{r}\sum_{k=-N_0}^{N_0}\frac{N_0-|k|}{N_0}R_{\mathbf{x}_j}(k)e^{-i\omega k} \tag{7.50}$$

which with the stationarity of $\mathbf{x}$ leads to

$$\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega)] = \sum_{k=-N_0}^{N_0}\frac{N_0-|k|}{N_0}R_{\mathbf{x}}(k)e^{-i\omega k}, \tag{7.51}$$

where $R_{\mathbf{x}}(k) := \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-k)] = \mathbb{E}[\mathbf{x}_j(t)\mathbf{x}_j(t-k)]$.
Through averaging of periodograms, apparently a lag-window is applied to the covariance function of the signal, being the same as the (triangular) Bartlett-window discussed in the previous section. Note the difference however that here the lag window is applied to the correlation function $R_{\mathbf{x}}$, while in spectral estimation the lag window is applied to the sample correlation $\hat{R}_{\mathbf{x}}$. In the frequency domain the effect of the averaging of periodograms naturally relates to the application of the Bartlett frequency window, convoluted with the spectral density $\Phi_{\mathbf{x}}$.
For this reason, the procedure of periodogram averaging is also known as the Bartlett-procedure. When the averaging mechanism is combined with a data window it is referred to as the Welch-method, see also Therrien (1992).

> Smoothing the periodogram estimator $\hat{\Phi}_{\mathbf{x}}(\omega)$ of a spectral density function $\Phi_{\mathbf{x}}(\omega)$ by either a window operation or an averaging will generally reduce the variance at the cost of an introduction of bias.

7.5 Non-parametric and parametric models

The (non-parametric) estimators for correlation functions and spectral density functions as presented in this chapter, provide first and rather rough tools for estimation of these characteristic properties of stochastic processes. The estimators generally suffer from a substantial variance, that can be reduced with some engineering tools as the smoothing and averaging operations discussed in the previous section. Nevertheless, in particular for the first order AR process:

$$\mathbf{x}(t) = a\mathbf{x}(t-1) + \mathbf{e}(t) \tag{7.52}$$

with $\mathbf{e}$ a unit-variance white noise, that is used as a running example in this chapter, it seems not very convincing that estimators of spectral density suffer from a large variance, while the basic properties can actually be described by only one real-valued parameter: a. In spectral estimation, the number of unknowns that is estimated from data actually equals the number of different frequencies that is considered. When phrasing the estimation problem as a parameter estimation problem, the correlation functions and spectral densities can be parameterized by one single parameter, which can be estimated from data. In the situation of a first order AR process this is reflected in

$$R_{\mathbf{x}}(k) = a^{|k|}\sigma_{\mathbf{x}}^2 = \frac{a^{|k|}}{1-a^2}\sigma_{\mathbf{e}}^2 \tag{7.53}$$

$$\Phi_{\mathbf{x}}(\omega) = \left|\frac{1}{1-ae^{-i\omega}}\right|^2\sigma_{\mathbf{e}}^2 = \frac{\sigma_{\mathbf{e}}^2}{1+a^2-2a\cos(\omega)}, \tag{7.54}$$

as was derived in section 4.3.2. Note that by parametrizing the estimation problem by the parameter a, the estimated correlation and spectral density functions are enforced to satisfy the structure as represented in (7.53)-(7.54) and therefore automatically satisfy a smoothness condition. In other words, smoothness of the estimators is enforced by the underlying model structure.

To illustrate the potential of parametric spectral estimation, let's consider once more the AR(1) process (7.52) that has been used as running example in this chapter. Given an available data sequence $\mathbf{x}(t)_{\{t=0,\cdots,N-1\}}$ the problem is to estimate the unknown parameter a. Note the similarity between (7.52) with the classical linear regression model (5.26), in the sense that (7.52) expresses $\mathbf{x}(t)$ as a linear function of the regressor $\mathbf{x}(t-1)$ with a acting as the (unknown) regressor coefficient and $\mathbf{e}(t)$ as the error term (also known as the "residual"). Of course, (7.52) is not identical to the classical regression model since, as t increases from 0 to $N-1$, $\mathbf{x}(t)$ simultaneously plays the role of both the *dependent* and *independent* variables. However, since the values of $\mathbf{x}(t)$ and $\mathbf{x}(t-1)$ are all observed there is nothing to prevent us from applying the same least squares procedure used in linear regression (see Chapter 5). Thus, if we adopt a least squares approach, we estimate a by finding the value that minimizes

$$V(a) = \sum_{t=1}^{N-1} e^2(t) = \sum_{t=1}^{N-1} \left(x(t) - ax(t-1)\right)^2. \tag{7.55}$$

Note that we cannot include the term $e^2(0)$ in (7.55) since it cannot be computed in terms of the observed $x(t)$. This is an unavoidable feature of the AR(1) model, but if N is large compared with 1, the effect of ignoring the first value of $e(t)$ will only be small, and (7.55) may be regarded as a good approximation of the "full" sum of squares $\sum_{t=0}^{N-1} e^2(t)$ (Priestly, 1981). The least squares criterion can (as usual) be justified from a maximum likelihood approach under the assumption that the residuals $\mathbf{e}(t)$ are zero mean Gaussian distributed. Straightforward calculation (see section 5.3) shows that the least squares estimator of the parameter a is given by:

$$\hat{\mathbf{a}}_{LS} = \left[\sum_{t=1}^{N-1} \mathbf{x}^2(t-1)\right]^{-1} \left[\sum_{t=1}^{N-1} \mathbf{x}(t)\mathbf{x}(t-1)\right]. \tag{7.56}$$

Note that (7.56) can be rewritten as

$$\hat{\mathbf{a}}_{LS} = \frac{\hat{R}_{\mathbf{x}}(1)}{\hat{R}_{\mathbf{x}}(0)}, \tag{7.57}$$

where

$$\hat{R}_{\mathbf{x}}(1) = \frac{1}{N-1} \sum_{t=1}^{N-1} \mathbf{x}(t)\mathbf{x}(t-1), \tag{7.58}$$

and

$$\hat{R}_{\mathbf{x}}(0) = \frac{1}{N-1} \sum_{t=1}^{N-1} \mathbf{x}^2(t-1) \tag{7.59}$$

are estimators of the values $R_{\mathbf{x}}(1)$ and $R_{\mathbf{x}}(0)$ of the autocorrelation function $R_{\mathbf{x}}(k)$.

If the noise variance $\sigma_{\mathbf{e}}^2$ is also unknown, an approximately unbiased estimator of $\sigma_{\mathbf{e}}^2$ can be obtained by (Priestley, 1981):

$$\hat{\sigma}_{\mathbf{e}}^2 = \frac{\sum_{t=1}^{N-1}(\mathbf{x}(t) - \hat{\mathbf{a}}_{LS}\mathbf{x}(t-1))^2}{N-2} = \frac{N-1}{N-2}\left(\hat{R}_{\mathbf{x}}(0) - \hat{\mathbf{a}}_{LS}\hat{R}_{\mathbf{x}}(1)\right). \tag{7.60}$$

If we substitute (7.56) and (7.60) for, respectively, a and $\sigma_{\mathbf{e}}^2$ in (7.53) and (7.54), (smooth) estimates of the correlation and spectral density function of the AR(1) process are obtained. The analysis performed can easily be extended to include higher order AR models (Priestly, 1981).
An alternative way to derive an estimator of the parameter a of the AR(1) process (7.52) is to follow a prediction error estimation approach. In this approach, we look for that value of a that minimizes the mean squared error criterion

$$\mathbb{E}[(\mathbf{x}(t) - \hat{\mathbf{x}}(t))^2] \tag{7.61}$$

with

$$\hat{\mathbf{x}}(t) = \mathbb{E}[\mathbf{x}(t)|\mathbf{x}(t-1)] = a\mathbf{x}(t-1). \tag{7.62}$$

Straightforward minimization then yields:

$$\hat{\mathbf{a}} = \arg\min_a \mathbb{E}[(\mathbf{x}(t) - a\mathbf{x}(t-1))^2] = \frac{R_{\mathbf{x}}(1)}{R_{\mathbf{x}}(0)}. \tag{7.63}$$

The similarity between the expressions (7.63) and (7.57) is obvious. If we replace the expectation operator $\mathbb{E}[\cdot]$ in (7.63) by the time average operator $\frac{1}{N-1}\sum_{t=1}^{N-1}[\cdot]$, the terms $R_{\mathbf{x}}(0)$ and $R_{\mathbf{x}}(1)$ in (7.63) will be replaced by their estimators (7.59) and (7.58), respectively, and (7.63) will become equal to (7.57).
A more comprehensive treatment of the parametric spectral estimation problem is beyond the current scope of this manuscript, and will be incorporated in a future chapter to be added.

7.6 Summary

Simple sample-correlation estimators for auto- and cross-correlation functions of stochastic processes have been analyzed. It has been shown that an asymptotically unbiased estimator can be constructed for which the variance tends to zero for increasing number of data. By transforming the estimator to the frequency domain, an estimator for the power spectral density is obtained. However without any additional precautions these frequency domain estimators suffer from a non-vanishing variance for increasing data. Special -rather heuristic- smoothing windows can be introduced to reduce the variance however at the cost of an introduction of bias.

Appendix

Proof of variance expression for $\breve{R}_{\mathbf{x}}(k)$ (7.5)
The variance of the estimator can be analyzed by considering

$$\begin{aligned} var(\breve{R}_{\mathbf{x}}(k)) &= \mathbb{E}[\breve{R}_{\mathbf{x}}(k) - R_{\mathbf{x}}(k)]^2 && \text{(7A.1)} \\ &= \mathbb{E}[\breve{R}^2_{\mathbf{x}}(k)] - R^2_{\mathbf{x}}(k). && \text{(7A.2)} \end{aligned}$$

The first term of this expression is given by (for $k \geq 0$):

$$\mathbb{E}[\breve{R}^2_{\mathbf{x}}(k)] = \frac{1}{(N-k)^2} \sum_{t=k}^{N-1} \sum_{j=k}^{N-1} \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t-k)\mathbf{x}(j)\mathbf{x}(j-k)]. \quad \text{(7A.3)}$$

The fourth order moment terms $\mathbb{E}[\{\mathbf{x}(t)\mathbf{x}(t-k)\mathbf{x}(j)\mathbf{x}(j-k)\}]$ are hard to analyze in the general situation. Only in the Gaussian case this analysis becomes tractable, by using the result that for real-valued Gaussian random variables (Shanmugan and Breipohl, 1988, p. 55, eq. 2.69):

$$\mathbb{E}\{\mathbf{x}_1\mathbf{x}_2\mathbf{x}_3\mathbf{x}_4\} = \mathbb{E}\{\mathbf{x}_1\mathbf{x}_2\}\mathbb{E}\{\mathbf{x}_3\mathbf{x}_4\} + \mathbb{E}\{\mathbf{x}_2\mathbf{x}_3\}\mathbb{E}\{\mathbf{x}_1\mathbf{x}_4\} + \mathbb{E}\{\mathbf{x}_2\mathbf{x}_4\}\mathbb{E}\{\mathbf{x}_1\mathbf{x}_3\}. \quad \text{(7A.4)}$$

Applying this to the above expression for $\mathbb{E}\breve{R}^2_{\mathbf{x}}(k)$ leads to (restricting attention to the situation $k \geq 0$):

$$\mathbb{E}[\breve{R}^2_{\mathbf{x}}(k)] = \frac{1}{(N-k)^2} \sum_{t=k}^{N-1} \sum_{j=k}^{N-1} R^2_{\mathbf{x}}(k) + R_{\mathbf{x}}(t-k-j)R_{\mathbf{x}}(t+k-j) + R^2_{\mathbf{x}}(t-j). \quad \text{(7A.5)}$$

Substituting this result into (7A.2) shows that

$$var(\breve{R}_{\mathbf{x}}(k)) = \frac{1}{(N-k)^2} \sum_{t=k}^{N-1} \sum_{j=k}^{N-1} R_{\mathbf{x}}(t-k-j)R_{\mathbf{x}}(t+k-j) + R^2_{\mathbf{x}}(t-j). \quad \text{(7A.6)}$$

Through variable substitution $t - j = m$ this leads to

$$\begin{aligned} var(\breve{R}_{\mathbf{x}}(k)) &= \frac{1}{(N-k)^2} \sum_{m=-N+k+1}^{N-1-k} (N-k-|m|)[R_{\mathbf{x}}(m-k)R_{\mathbf{x}}(m+k) + R^2_{\mathbf{x}}(m)] \\ &= \frac{1}{N-k} \sum_{m=-N+k+1}^{N-1-k} (1 - \frac{|m|}{N-k})[R_{\mathbf{x}}(m-k)R_{\mathbf{x}}(m+k) + R^2_{\mathbf{x}}(m)]. \quad \text{(7A.7)} \end{aligned}$$

And since the variance will be the same for either positive or negative values of k it follows that for any $|k| \leq N$:

$$var(\breve{R}_{\mathbf{x}}(k)) = \frac{1}{N-|k|} \sum_{m=-N+|k|+1}^{N-1-|k|} (1 - \frac{|m|}{N-|k|}) \left[R_{\mathbf{x}}(m-k)R_{\mathbf{x}}(m+k) + R^2_{\mathbf{x}}(m)\right]. \quad \text{(7A.8)}$$

Proof of variance expression for $\hat{R}_{\mathbf{yx}}(k)$

The variance of the estimator can be analyzed by considering

$$var(\hat{R}_{\mathbf{yx}}(k)) = \mathbb{E}[\hat{R}_{\mathbf{yx}}(k) - R_{\mathbf{yx}}(k)]^2 = \mathbb{E}[\hat{R}^2_{\mathbf{yx}}(k)] - [\mathbb{E}R_{\mathbf{yx}}(k)]^2. \tag{7A.9}$$

First considering the situation $k \geq 0$ it follows with

$$\hat{R}_{\mathbf{yx}}(k) = \frac{1}{N}\sum_{t=k}^{N-1} y(t)x(t-k)$$

that

$$\mathbb{E}[\hat{R}^2_{\mathbf{yx}}(k)] = \frac{1}{N^2}\sum_{t=k}^{N-1}\sum_{j=k}^{N-1} \mathbb{E}[\mathbf{y}(t)\mathbf{x}(t-k)\mathbf{y}(j)\mathbf{x}(j-k)]. \tag{7A.10}$$

With (7A.4) this is written as

$$\mathbb{E}[\hat{R}^2_{\mathbf{yx}}(k)] = \frac{1}{N^2}\sum_{t=k}^{N-1}\sum_{j=k}^{N-1} R^2_{\mathbf{yx}}(k)+R_{\mathbf{xy}}(t-j-k)R_{\mathbf{yx}}(t+k-j)+R_{\mathbf{x}}(t-j)R_{\mathbf{y}}(t-j). \tag{7A.11}$$

Then it follows that

$$var(\hat{R}_{\mathbf{yx}}(k)) = \frac{1}{N^2}\sum_{t=k}^{N-1}\sum_{j=k}^{N-1} R_{\mathbf{xy}}(t-j-k)R_{\mathbf{yx}}(t+k-j) + R_{\mathbf{x}}(t-j)R_{\mathbf{y}}(t-j). \tag{7A.12}$$

Through variable substitution $t - j = m$ this leads to

$$var\{\hat{R}_{\mathbf{yx}}(k)\} = \frac{1}{N}\sum_{m=-N+|k|+1}^{N-1-|k|} (1 - \frac{|k|+|m|}{N}) \left[R_{\mathbf{xy}}(m-k)R_{\mathbf{yx}}(m+k) + R_{\mathbf{x}}(m)R_{\mathbf{y}}(m)\right]. \tag{7A.13}$$

In this expression $k = |k|$ has been substituted, which is justified by the fact that $k \geq 0$. Analysis for $k \leq 0$ and using the related expression

$$\hat{R}_{\mathbf{yx}}(k) = \frac{1}{N}\sum_{t=0}^{N-1-|k|} y(t)x(t-k)$$

shows that exactly the same expression (7A.13) results, which proves that it is valid for all values of k.

Proof of (co)variance expressions of periodogram estimator for white noise x

The covariance of $\hat{\Phi}_{\mathbf{x}}(\omega_1)$ and $\hat{\Phi}_{\mathbf{x}}(\omega_2)$ is given by

$$Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} = \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)\hat{\Phi}_{\mathbf{x}}(\omega_2)] - \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)]\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_2)] \tag{7A.14}$$

First an analysis will be made of the first term on the right hand side.
With

$$\hat{\Phi}_{\mathbf{x}}(\omega) = \frac{1}{N}\left\{\sum_{k=0}^{N-1} \mathbf{x}(k)e^{-i\omega k}\right\}\left\{\sum_{\ell=0}^{N-1} \mathbf{x}(\ell)e^{i\omega \ell}\right\}$$

it follows that

$$\begin{aligned}\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)\hat{\Phi}_{\mathbf{x}}(\omega_2)] = \\ = \frac{1}{N^2}\sum_{k=0}^{N-1}\sum_{\ell=0}^{N-1}\sum_{m=0}^{N-1}\sum_{n=0}^{N-1}\mathbb{E}\{\mathbf{x}(k)\mathbf{x}(\ell)\mathbf{x}(m)\mathbf{x}(n)\}e^{-i(k-\ell)\omega_1}e^{-i(m-n)\omega_2}\end{aligned} \tag{7A.15}$$

Using again (7A.4):

$$\mathbb{E}\{\mathbf{x}_1\mathbf{x}_2\mathbf{x}_3\mathbf{x}_4\} = \mathbb{E}\{\mathbf{x}_1\mathbf{x}_2\}\mathbb{E}\{\mathbf{x}_3\mathbf{x}_4\} + \mathbb{E}\{\mathbf{x}_2\mathbf{x}_3\}\mathbb{E}\{\mathbf{x}_1\mathbf{x}_4\} + \mathbb{E}\{\mathbf{x}_2\mathbf{x}_4\}\mathbb{E}\{\mathbf{x}_1\mathbf{x}_3\} \tag{7A.16}$$

expression (7A.15) leads to three separate terms that will be analyzed separately.

- The first term of expression (7A.15) leads to a contribution for $k = \ell$ and $m = n$, leading to

$$\frac{1}{N^2}\sum_{k=0}^{N-1}\sum_{m=0}^{N-1}\sigma_x^4 = \sigma_x^4$$

- The second term of (7A.16) leads to a contribution for $k = n$ and $\ell = m$, leading to

$$\begin{aligned}&\frac{1}{N^2}\sum_{k=0}^{N-1}\sum_{\ell=0}^{N-1}\sigma_x^4 e^{-i(k-\ell)\omega_1}e^{i(k-\ell)\omega_2} = \\ &= \frac{\sigma_x^4}{N^2}\sum_{k=0}^{N-1}e^{-ik(\omega_1-\omega_2)}\sum_{\ell=0}^{N-1}e^{i\ell(\omega_1-\omega_2)} \\ &= \frac{\sigma_x^4}{N^2}\left[\frac{1-e^{-iN(\omega_1-\omega_2)}}{1-e^{-i(\omega_1-\omega_2)}}\right]\left[\frac{1-e^{iN(\omega_1-\omega_2)}}{1-e^{i(\omega_1-\omega_2)}}\right] \\ &= \frac{\sigma_x^4}{N^2}\frac{e^{iN(\omega_1-\omega_2)/2}\cdot[1-e^{-iN(\omega_1-\omega_2)}]}{e^{i(\omega_1-\omega_2)/2}\cdot[1-e^{-i(\omega_1-\omega_2)}]}\cdot\frac{e^{-iN(\omega_1-\omega_2)/2}\cdot[1-e^{iN(\omega_1-\omega_2)}]}{e^{-i(\omega_1-\omega_2)/2}\cdot[1-e^{i(\omega_1-\omega_2)}]} \\ &= \sigma_x^4\left[\frac{sin[N(\omega_1-\omega_2)/2]}{N\cdot sin(\omega_1-\omega_2)/2}\right]^2\end{aligned} \tag{7A.17}$$

- The third term in (7A.16) leads to a contribution for $k = m$ and $\ell = n$, leading to

$$\begin{aligned}&\frac{1}{N^2}\sum_{k=0}^{N-1}\sum_{\ell=0}^{N-1}\sigma_x^4 e^{-i(k-\ell)\omega_1}e^{-i(k-\ell)\omega_2} = \\ &= \frac{\sigma_x^4}{N^2}\sum_{k=0}^{N-1}e^{-ik(\omega_1+\omega_2)}\sum_{\ell=0}^{N-1}e^{i\ell(\omega_1+\omega_2)}\end{aligned} \tag{7A.18}$$

which with a similar analysis as before will lead to

$$\sigma_x^4\left[\frac{sin[N(\omega_1+\omega_2)/2]}{N\cdot sin(\omega_1+\omega_2)/2}\right]^2$$

Combining the three terms now result in

$$\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)\hat{\Phi}_{\mathbf{x}}(\omega_2)] = \sigma_x^4 \left\{ 1 + \left[\frac{sin[N(\omega_1 - \omega_2)/2]}{N \cdot sin(\omega_1 - \omega_2)/2} \right]^2 + \left[\frac{sin[N(\omega_1 + \omega_2)/2]}{N \cdot sin(\omega_1 + \omega_2)/2} \right]^2 \right\}$$

Since

$$Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} = \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)\hat{\Phi}_{\mathbf{x}}(\omega_2)] - \mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_1)]\mathbb{E}[\hat{\Phi}_{\mathbf{x}}(\omega_2)]$$

and $\mathbb{E}\hat{\Phi}_{\mathbf{x}}(\omega) = \sigma_{\mathbf{x}}^2$ for all ω, it follows that

$$Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} = \sigma_x^4 \left\{ \left[\frac{sin[N(\omega_1 - \omega_2)/2]}{N \cdot sin(\omega_1 - \omega_2)/2} \right]^2 + \left[\frac{sin[N(\omega_1 + \omega_2)/2]}{N \cdot sin(\omega_1 + \omega_2)/2} \right]^2 \right\} \quad \text{(7A.19)}$$

The variance expression can now be derived, by considering $\omega_1 = \omega_2 = 2\pi k/N$ with k integer, and using the fact that $sin\alpha \to \alpha$ for $\alpha \to 0$, it follows that

$$var\{\hat{\Phi}_{\mathbf{x}}(\omega)\}_{\omega=2\pi k/N} = \sigma_x^4 \{1 + \left(\frac{sin2\pi k}{N sin(2\pi k/N)} \right)^2 \}$$

leading to

$$var\{\hat{\Phi}_{\mathbf{x}}(f)\}_{\omega=2\pi k/N} = \begin{cases} 2 \cdot \sigma_{\mathbf{x}}^4 & k = 0, N/2 \\ \sigma_{\mathbf{x}}^4 & k \neq 0, N/2. \end{cases} \quad \text{(7A.20)}$$

For the covariance, it follows from (7A.19) that

$$Cov\{\hat{\Phi}_{\mathbf{x}}(\omega_1), \hat{\Phi}_{\mathbf{x}}(\omega_2)\} = 0, \quad \text{for } \omega_1 = 2\pi k_1/N; \omega_2 = 2\pi k_2/N; \omega_1 \neq \pm\omega_2$$

Appendix A

Probability Theory

A brief overview is given of the basic definitions and properties that are involved in basic probability theory. For a more extensive discussion and background material see e.g. Papoulis (1991).

A.1 Random Experiment

A *random experiment* or *random model* is defined as a triplet $(\mathbb{S}, S, P)$ where

- $\mathbb{S}$ is a *class of events* to which a probability can be assigned
- S is the *sample space*, i.e. the set of possible outcomes of an experiment
- P a *probability measure*, applied to $\mathbb{S}$

The sample space S is the set of possible outcomes of an experiment. For the experiment of tossing a die and counting the values, this set is simply $\{1, 2, 3, 4, 5, 6\}$. The class of events $\mathbb{S}$ is the set of all events to which a probability can be assigned. It is composed of elements or subsets of S. For the example of tossing the die, it incorporates elements as {odd numbers smaller than 6}, $\{1, 2, 3\}$, etcetera. $\mathbb{S}$ is a family of subsets of S that is required to be completely additive, i.e.

- The whole sample space S is an allowed event, i.e. $S \subset \mathbb{S}$;
- Any union of events is also an event, i.e. if $A_k \subset \mathbb{S}$ for $k = 1, 2, 3, \cdots$, then $\bigcup_{k=1}^{n} A_k \subset \mathbb{S}$, for $n = 1, 2, 3, \cdots$;
- If A is an event, then *not* A is also an event, i.e. if $A \subset \mathbb{S}$, then $\bar{A} \subset \mathbb{S}$, where $\bar{A}$ is the complement [1] of A in $\mathbb{S}$.

[1] Complement is here considered in a set-theoretic sense.

A.2 Random Variables

Scalar real random variables

A *random variable* is a function that connects a real-valued number to each of the elements of the sample space of a random model. In other words, every outcome of an experiment is represented by a real-valued number. It is a function $S \to \mathbb{R}$, such that for all $\lambda \in S$, $\mathbf{x}(\lambda) \in \mathbb{R}$ satisfies

- The set $\{\lambda : \mathbf{x}(\lambda) \leq x\}$ is an event $\forall x \in \mathbb{R}$
- $P\{\lambda : \mathbf{x}(\lambda) = \infty\} = 0$; $P\{\lambda : \mathbf{x}(\lambda) = -\infty\} = 0$.

The random variable is denoted with the sans-serif symbol $\mathbf{x}$, while a particular outcome, also referred to as a *realization* is indicated by the italic symbol x.
The first condition above guarantees the validity of the expression $P(\mathbf{x} \leq x)$, the probability that a random variable takes a value that is smaller than a particular real-valued number x. This function is called the *distribution function*, defined as

$$\text{Distribution function:} \qquad F_{\mathbf{x}}(x) := P(\mathbf{x} \leq x).$$

Typical properties of distribution functions are that they are non-decreasing functions of x. Their characteristic properties are:

- $F_{\mathbf{x}}(-\infty) = 0$
- $F_{\mathbf{x}}(\infty) = 1$
- $F_{\mathbf{x}}(x_1) \leq F_{\mathbf{x}}(x_2)$ if $x_1 < x_2$
- $P[x_1 < \mathbf{x} \leq x_2] = F_{\mathbf{x}}(x_2) - F_{\mathbf{x}}(x_1)$.

For a continuous random variable ($\mathbf{x}$ can take values in a continuous interval on the real line) the *probability density function (pdf)* is defined as:

$$f_{\mathbf{x}}(x) := \frac{dF_{\mathbf{x}}(x)}{dx} \qquad \geq 0 \ \forall \ x \tag{A.1}$$

$$P(a \leq \mathbf{x} \leq b) = \int_a^b f_{\mathbf{x}}(x)dx. \tag{A.2}$$

For a discrete random variable ($\mathbf{x}$ can take values in a discrete interval on the real line) the *probability mass function (pmf)* is defined as:

$$P(\mathbf{x} = x_i) \quad \text{for } i = 1, \cdots n \tag{A.3}$$

where x_i ranges over the allowed (discrete) values of $\mathbf{x}$.
A random variable maps the sample space S onto a subset $S_{\mathbf{x}}$ of $\mathbb{R}$. This image set $S_{\mathbf{x}}$ is the set of possible real-valued outcomes of a random experiment, and is also referred to as the *ensemble*.

Mean and Variance of scalar random variables

The mean of a continuous random variable is defined as

$$\text{Mean:} \qquad \mu_{\mathbf{x}} := \mathbb{E}\{\mathbf{x}\} = \int_{-\infty}^{\infty} x f_{\mathbf{x}}(x)dx$$

and the variance:

$$\text{Variance:} \qquad \sigma_{\mathbf{x}}^2 := \mathbb{E}\{(\mathbf{x} - \mu_{\mathbf{x}})^2\} = \int_{-\infty}^{\infty} (x - \mu_{\mathbf{x}})^2 f_{\mathbf{x}}(x) dx$$

The square root of the variance is referred to as the *standard deviation* $\sigma_{\mathbf{x}}$.

For discrete random variables the corresponding notions are

$$\mu_{\mathbf{x}} := \mathbb{E}\{\mathbf{x}\} = \sum_{i=1}^{n} x_i P(\mathbf{x} = x_i)$$

and the variance:

$$\sigma_{\mathbf{x}}^2 = \mathbb{E}\{(\mathbf{x} - \mu_{\mathbf{x}})^2\} = \sum_{i=1}^{n} (x_i - \mu_{\mathbf{x}})^2 P(\mathbf{x} = x_i).$$

Functions of two continuous random variables

For two continuous random variables the joint distribution function is determined by the joint occurrence of two events, that is

$$F_{\mathbf{x},\mathbf{y}}(x, y) = P[\mathbf{x} \leq x, \mathbf{y} \leq y].$$

The corresponding *joint probability density function* is

$$f_{\mathbf{x},\mathbf{y}}(x, y) := \frac{\partial^2 F_{\mathbf{x},\mathbf{y}}(x, y)}{\partial x \partial y}$$

and as a result

$$P(\mathbf{x} \leq a, \mathbf{y} \leq b) = \int_{-\infty}^{a} \int_{-\infty}^{b} f_{\mathbf{x},\mathbf{y}}(x, y) dx dy.$$

The mean value of any function g of the two random variables $\mathbf{x}$,$\mathbf{y}$ is then given by:

$$\mathbb{E}[g(\mathbf{x}, \mathbf{y})] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) f_{\mathbf{x},\mathbf{y}}(x, y) dx dy.$$

Since $F_{\mathbf{x},\mathbf{y}}(\infty, \infty) = 1$ it follows that

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{\mathbf{x},\mathbf{y}}(x, y) dx dy = 1.$$

Additionally, from $F_{\mathbf{x},\mathbf{y}}(x, \infty) = F_{\mathbf{x}}(x)$ it follows that for any fixed valued x:

$$f_{\mathbf{x}}(x) = \int_{-\infty}^{\infty} f_{\mathbf{x},\mathbf{y}}(x, y) dy. \tag{A.4}$$

By dual reasoning it follows that for any fixed value y, $f_{\mathbf{y}}(y) = \int_{-\infty}^{\infty} f_{\mathbf{x},\mathbf{y}}(x, y) dx$.

Conditional densities and Bayes' rule

Based on *Bayes' rule* for conditional probabilities:

$$P(A|B) = \frac{P(A, B)}{P(B)}$$

the *conditional probability density function* for two random variables x, y is defined by

$$f_{\mathbf{y}|\mathbf{x}}(y|x) := \frac{f_{\mathbf{x},\mathbf{y}}(x,y)}{f_{\mathbf{x}}(x)} \tag{A.5}$$

and thus $f_{\mathbf{x},\mathbf{y}}(x,y) = f_{\mathbf{y}|\mathbf{x}}(y|x) \cdot f_{\mathbf{x}}(x)$.
With (A.4) it is noted that

$$\int_{-\infty}^{\infty} f_{\mathbf{y}|\mathbf{x}}(y|x)dy = \int_{-\infty}^{\infty} \frac{f_{\mathbf{x},\mathbf{y}}(x,y)}{f_{\mathbf{x}}(x)}dy = \frac{f_{\mathbf{x}}(x)}{f_{\mathbf{x}}(x)} = 1$$

for any value of x.
The *conditional mean* $\mathbb{E}_{\mathbf{y}|\mathbf{x}}[g(\mathbf{x},\mathbf{y})|\mathbf{x} = x]$ is defined by

$$\mathbb{E}_{\mathbf{y}|\mathbf{x}}[g(\mathbf{x},\mathbf{y})|\mathbf{x} = x] = \int_{-\infty}^{\infty} g(x,y) f_{\mathbf{y}|\mathbf{x}}(y|x)dy.$$

From (A.5) then follows that

$$\mathbb{E}\{g(\mathbf{x},\mathbf{y})\} = \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\mathbf{y}|\mathbf{x}}[g(\mathbf{x},\mathbf{y})|\mathbf{x}].$$

Moments of scalar real random variables The moments of order (p,r) of two scalar random variables $\mathbf{x}$, $\mathbf{y}$ are defined as

$$m(p,r) = \mathbb{E}[\mathbf{x}]^p[\mathbf{y}]^r,$$

with $p+r$ referred to as the order of the moment. First order moments are the mean of $\mathbf{x}$ (corresponding to $p=1$, $r=0$), and the mean of $\mathbf{y}$ ($p=0$, $r=1$). Second order moments include the correlation between $\mathbf{x}$ and $\mathbf{y}$:

$$\text{Correlation:} \quad r_{\mathbf{xy}} := \mathbb{E}(\mathbf{xy}). \tag{A.6}$$

The *central* moments of order (p,r) are defined as

$$m_c(p,r) = \mathbb{E}[\mathbf{x} - \mu_{\mathbf{x}}]^p[\mathbf{y} - \mu_{\mathbf{y}}]^r.$$

Second order central moments include the variance of $\mathbf{x}$ and $\mathbf{y}$, and the *covariance*:

$$\text{Covariance:} \quad \sigma_{\mathbf{xy}} := \mathbb{E}\{(\mathbf{x} - \mu_{\mathbf{x}})(\mathbf{y} - \mu_{\mathbf{y}})\}. \tag{A.7}$$

A scaled version of this covariance, is denoted as the correlation coefficient between random variables $\mathbf{x}$ and $\mathbf{y}$:
Correlation coefficient:

$$\rho_{\mathbf{xy}} = \frac{\mathbb{E}\{(\mathbf{x} - \mu_{\mathbf{x}})(\mathbf{y} - \mu_{\mathbf{y}})\}}{\sigma_{\mathbf{x}}\sigma_{\mathbf{y}}} \qquad -1 \leq \rho_{\mathbf{xy}} \leq 1.$$

The correlation coefficient directly reflects the statistical relation between the two random variables $\mathbf{x}$ and $\mathbf{y}$. Because of its scaling it takes on real values between -1 en $+1$. This latter property can simply be shown by using the Schwartz inequality (A.12).
Two random variables $\mathbf{x}$ and $\mathbf{y}$ are defined to be

- *independent* if $f_{\mathbf{x},\mathbf{y}}(x,y) = f_{\mathbf{x}}(x) \cdot f_{\mathbf{y}}(y)$;
- *uncorrelated* if $\sigma_{\mathbf{xy}} = 0$ or equivalently $\mathbb{E}\{\mathbf{xy}\} = \mu_{\mathbf{x}} \cdot \mu_{\mathbf{y}}$;
- *orthogonal* if $r_{\mathbf{xy}} = 0$ or equivalently $\mathbb{E}\{\mathbf{xy}\} = 0$.

Note that two independent random variables are always uncorrelated; however two random variables that are uncorrelated can be statistically dependent. Therefore independency is a stronger notion than uncorrelation.

Moments of vector complex random variables

If the random variables are vectors and complex-valued, the above notions simply generalize. Consider the n-dimensional complex-valued random variable $\mathbf{x}$, denoted by

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \cdot \\ \mathbf{x}_n \end{bmatrix}; \quad \mathbf{x}^T = [\mathbf{x}_1 \ \cdots \ \mathbf{x}_n].$$

The mean of $\mathbf{x}$ is denoted by

$$\mu_{\mathbf{x}} = \mathbb{E}\{\mathbf{x}\} = \mathbb{E} \begin{bmatrix} \mathbf{x}_1 \\ \cdot \\ \mathbf{x}_n \end{bmatrix}.$$

For an n-dimensional random variable $\mathbf{x}$ and an m-dimensional random variable $\mathbf{y}$, the *correlation matrix* becomes

$$\text{correlation matrix} \qquad R_{\mathbf{xy}} = \mathbb{E}\{\mathbf{x}\mathbf{y}^*\} \tag{A.8}$$

and the *covariance matrix* is denoted as:

$$\text{covariance matrix} \qquad \Sigma_{\mathbf{xy}} := \mathbb{E}\{(\mathbf{x} - \mu_{\mathbf{x}})(\mathbf{y} - \mu_{\mathbf{y}})^*\}, \tag{A.9}$$

being matrices with dimensions $n \times m$. The covariance matrix is structured as

$$\Sigma_{\mathbf{xy}} = \begin{bmatrix} \sigma_{\mathbf{x}_1\mathbf{y}_1} & \sigma_{\mathbf{x}_1\mathbf{y}_2} & \cdots & \sigma_{\mathbf{x}_1\mathbf{y}_m} \\ \sigma_{\mathbf{x}_2\mathbf{y}_1} & \sigma_{\mathbf{x}_2\mathbf{y}_2} & \cdots & \sigma_{\mathbf{x}_2\mathbf{y}_m} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{\mathbf{x}_n\mathbf{y}_1} & \sigma_{\mathbf{x}_n\mathbf{y}_2} & \cdots & \sigma_{\mathbf{x}_n\mathbf{y}_m} \end{bmatrix}.$$

In line with the above notation, for the n-dimensional random variable $\mathbf{x}$, the $n \times n$ covariance matrix is denoted as $\Sigma_{\mathbf{xx}}$, also often abbreviated to $\Sigma_{\mathbf{x}}$.
For real-valued random variables, the moments are real-valued also. For complex-valued random variables, the mean and covariance matrix are complex-valued.
Similar to the scalar case, two vector random variables $\mathbf{x}$ and $\mathbf{y}$ are defined to be

- independent, if $f_{\mathbf{x},\mathbf{y}}(x,y) = f_{\mathbf{x}}(x) f_{\mathbf{y}}(y)$;
- uncorrelated, if $\Sigma_{\mathbf{xy}} = 0$ or equivalently $\mathbb{E}\{\mathbf{x}\mathbf{y}^*\} = \mu_{\mathbf{x}} \cdot \mu_{\mathbf{y}}^*$;
- orthogonal, if $\mathbb{E}\{\mathbf{x}\mathbf{y}^*\} = 0$.

For vector-valued real random variables, the complex conjugate transpose operation $(\cdot)^*$ simply reduces to a transpose $(\cdot)^T$.
Note that for complex-valued random variables, expressions as $F_{\mathbf{x}}(x) = P(\mathbf{x} \leq x)$ have to be interpreted for complex valued $\mathbf{x}$ and x. In this situation $P(\mathbf{x} \leq x)$ has to be read as $P(\mathbf{x}_r \leq x_r, \mathbf{x}_i \leq x_i)$ which leads to a notation and interpretation which is fully consistent with the framework presented above. The probability density function $f_{\mathbf{x}}(x)$ is then interpreted in a similar way.

A.3 Gaussian distribution

A.3.1 Gaussian probability density function (pdf)

The *Gaussian*[2] or *Normal distribution* has a probability density function that is given by

$$f_{\mathbf{x}}(x) = \frac{1}{\sigma_{\mathbf{x}}\sqrt{2\pi}}\, e^{-\frac{(x-\mu_{\mathbf{x}})^2}{2\sigma_{\mathbf{x}}^2}}.$$

Its mean value equals $\mu_{\mathbf{x}}$ and its variance is $\sigma_{\mathbf{x}}^2$. A sketch of a Gaussian distribution is given in Figure A.3.1.

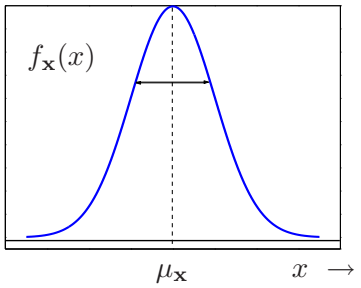

Figure A.1: Gaussian probability density function $f_{\mathbf{x}}(x)$.

In the multivariate (vector-valued) situation, where $\mathbf{x}$ is an n-dimensional random variable, the distribution function generalizes to

$$\begin{aligned} f_{\mathbf{x}_1,\mathbf{x}_2,\cdots\mathbf{x}_n}(x_1, x_2, \cdots x_n) &= f_{\mathbf{x}}(x) \\ = \frac{1}{(2\pi)^{n/2}\sqrt{det(\Sigma_{\mathbf{x}})}}\, exp\left(-\tfrac{1}{2}(x-\mu_{\mathbf{x}})^T\Sigma_{\mathbf{x}}^{-1}(x-\mu_{\mathbf{x}})\right). \end{aligned} \tag{A.10}$$

A.3.2 Properties of multivariate Gaussian random variables

Linear combinations of Gaussian random variables

If $\mathbf{x}$ is an n-dimensional Gaussian random variable with p.d.f

$$f_{\mathbf{x}}(x) = \frac{1}{(2\pi)^{n/2}\sqrt{det(\Sigma_{\mathbf{x}})}} e^{-\frac{1}{2}(\mathbf{x}-\mu_{\mathbf{x}})^T\Sigma_{\mathbf{x}}^{-1}(\mathbf{x}-\mu_{\mathbf{x}})}$$

and the m-dimensional random variable $\mathbf{y}$ is constructed from $\mathbf{x}$ by

$$\mathbf{y} = A\mathbf{x}$$

[2] Johann Carl Friedrich Gauss (1777-1855) was a German mathematician.

with A an $m \times n$ matrix having rank m, then $\mathbf{y}$ is also Gaussian with p.d.f.

$$f_{\mathbf{y}}(y) = \frac{1}{(2\pi)^{m/2}\sqrt{det(\Sigma_{\mathbf{y}})}} e^{-\frac{1}{2}(\mathbf{y}-\mu_{\mathbf{y}})^T \Sigma_{\mathbf{y}}^{-1}(\mathbf{y}-\mu_{\mathbf{y}})}$$

with $\mu_{\mathbf{y}} = A\mu_{\mathbf{x}}$ and $\Sigma_{\mathbf{y}} = A\Sigma_{\mathbf{x}}A^T$.

Partitioning and conditional expectation of Gaussian random variables
Let $\mathbf{x}$ have an n-dimensional multivariate Gaussian distribution. If $\mathbf{x}$ is partitioned as

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} \quad \mu_{\mathbf{x}} = \begin{bmatrix} \mu_{\mathbf{x}_1} \\ \mu_{\mathbf{x}_2} \end{bmatrix} \quad \Sigma_{\mathbf{x}} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$$

with $dim(\mathbf{x}_1) = k$, then

- $\mathbf{x}_1$ is a k-dimensional Gaussian random variable with mean $\mu_{\mathbf{x}_1}$ and covariance matrix Σ_{11};
- The conditional mean of $\mathbf{x}_1$ given $\mathbf{x}_2$ is given by

  $$\mathbb{E}[\mathbf{x}_1|\mathbf{x}_2 = x_2] = \mu_{\mathbf{x}_1} + \Sigma_{12}\Sigma_{22}^{-1}(x_2 - \mu_{\mathbf{x}_2})$$

  being a linear function of x_2.

A.3.3 Higher order moments of Gaussian random variables

For a zero-mean Gaussian random variable with variance $\sigma_{\mathbf{x}}^2$, it follows that

$$\mathbb{E}\mathbf{x}^4 = 3\sigma_{\mathbf{x}}^4. \tag{A.11}$$

This result can be shown as follows:

$$\begin{aligned}
\mathbb{E}\mathbf{x}^4 &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_{\mathbf{x}}} e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} x^4 dx = \frac{2}{\sqrt{2\pi}\sigma_{\mathbf{x}}} \int_0^{\infty} x^4 e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} dx \\
&= \frac{-2 \cdot 2\sigma_{\mathbf{x}}^2}{\sqrt{2\pi}\sigma_{\mathbf{x}}} \int_0^{\infty} x^4 \frac{1}{2x} de^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} = \frac{-2\sigma_{\mathbf{x}}}{\sqrt{2\pi}} \int_0^{\infty} x^3 de^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} \\
&= \frac{-2\sigma_{\mathbf{x}}}{\sqrt{2\pi}} \left\{ x^3 e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} \Big|_0^{\infty} - \int_0^{\infty} e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} 3x^2 dx \right\} \\
&= \frac{2\sigma_{\mathbf{x}}}{\sqrt{2\pi}} 3 \int_0^{\infty} x^2 e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} dx = \sigma_{\mathbf{x}}^2 \cdot 3 \cdot \int_{-\infty}^{\infty} x^2 \frac{1}{\sqrt{2\pi}\sigma_{\mathbf{x}}} e^{-\frac{x^2}{2\sigma_{\mathbf{x}}^2}} dx = 3\sigma_{\mathbf{x}}^4.
\end{aligned}$$

A.4 Schwartz[3] inequality

Consider two random variables $\mathbf{x}$ and $\mathbf{y}$. Then

$$(\mathbb{E}\mathbf{xy})^2 \leq \mathbb{E}(\mathbf{x}^2) \cdot \mathbb{E}(\mathbf{y}^2). \tag{A.12}$$

[3] H. Amandus Schwarz (1843-1921) was a German mathematician.

The proof of this equation is obtained by considering any $a \in \mathbb{R}$, and by verifying that

$$\mathbb{E}(a\mathbf{x} - \mathbf{y})^2 = a^2\mathbb{E}(\mathbf{x}^2) - 2a\mathbb{E}(\mathbf{xy}) + \mathbb{E}(\mathbf{y}^2) \geq 0 \quad \forall a \in \mathbb{R}.$$

This actually is a quadratic form in a, and in order to satisfy the inequality for all a its discriminant should be ≤ 0 so as to avoid that the quadratic form has at maximum 1 point of a where the equality holds. As a result,

$$4\mathbb{E}^2(\mathbf{xy}) - 4\mathbb{E}(\mathbf{x}^2)\mathbb{E}(\mathbf{y}^2) \leq 0.$$

A.5 Convergence of random variables

Mean-square convergence

A sequence $\{\mathbf{x}_n\}_{n=1,2,\cdots}$ of random variables is said to converge in mean square sense if there exists a random variable $\mathbf{x}$ such that

$$\mathbb{E}[(\mathbf{x}_n - \mathbf{x})^2] \to 0 \quad \text{as } n \to \infty.$$

If this equation is satisfied, then the random variable $\mathbf{x}$ is called the mean square limit of the sequence $\mathbf{x}_n$, denoted by

$$l.i.m.\ \mathbf{x}_n = \mathbf{x}.$$

Convergence in probability

A sequence $\{\mathbf{x}_n\}_{n=1,2,\cdots}$ of random variables is said to converge in probability to the random variable $\mathbf{x}$ if the probability $P\{|\mathbf{x}_n - \mathbf{x}| > \varepsilon\}$ of the event $|\mathbf{x}_n - \mathbf{x}| > \varepsilon$ tends to zero as $n \to \infty$ for any $\varepsilon > 0$.

Convergence in mean square sense implies convergence in probability.

Convergence almost everywhere or convergence with probability 1

A sequence $\{\mathbf{x}_n\}_{n=1,2,\cdots}$ of random variables is said to converge almost everywhere (a.e.) or with probability 1 (w.p. 1) if

$$P\{\lambda : \lim_{N\to\infty} \mathbf{x}_n(\lambda) = \mathbf{x}(\lambda)\} = 1$$

also denoted as $P\{\mathbf{x}_n \to \mathbf{x}\} = 1$.
Here λ reflects one outcome of the experiment, inducing a sequence of numbers $\mathbf{x}_1(\lambda), \mathbf{x}_2(\lambda), \cdots \mathbf{x}_n(\lambda)$. For (almost) every outcome λ, the resulting sequence will converge to $\mathbf{x}(\lambda)$.

Convergence almost everywhere implies convergence in probability.

Central limit theorem

Let $\mathbf{x}_1, \mathbf{x}_2, \cdots \mathbf{x}_n$ be a sequence of identically distributed independent random variables, with mean value μ and variance σ^2. Then

$$\mathbf{z}_n = \sum_{i=1}^{n} (\mathbf{x}_i - \mu)/\sqrt{n\sigma^2}$$

converges for $n \to \infty$ in distribution to a Gaussian pdf with mean value 0 and variance 1.

Appendix B

Linear algebra and system theory

B.1 Vectors, matrices and norms

Identity matrix:

$$I_r = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix} \in \mathbb{R}^{r \times r}$$

Vector norms: For a vector $x = [x_1 \; x_2 \; \cdots x_n]^T \in \mathbb{R}^n$ the 2-norm is given by

$$\|x\|_2 = \sqrt{x^T x} = \sqrt{x_1^2 + \cdots + x_n^2}.$$

Matrix norms: Let $A \in \mathbb{R}^{n \times m}$ then

- $\|A\|_2 = \max_{x \neq 0} \dfrac{\|Ax\|_2}{\|x\|_2}$
- $\|A\|_F = \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{m} |a_{ij}|^2} = \sqrt{tr(AA^T)}$

For a square matrix $A \in \mathbb{R}^{n \times n}$ with elements a_{ij} its trace is defined by

$$tr(A) = \sum_{i=1}^{n} a_{ii}.$$

A square matrix $A \in \mathbb{R}^{n \times n}$ is *positive definite* if for all $x \in \mathbb{R}^n$, $x^T Ax > 0$. A is *positive semi-definite* if for all $x \in \mathbb{R}^n$, $x^T Ax \geq 0$.

B.2 Singular Value Decomposition

Definition B.1 (Singular Value Decomposition) For every finite matrix $P \in \mathbb{R}^{q \times r}$ there exist unitary matrices $U \in \mathbb{R}^{q \times q}$, $V \in \mathbb{R}^{r \times r}$, i.e.

$$U^T U = I_q \tag{B.1}$$
$$V^T V = I_r \tag{B.2}$$

such that

$$P = U\Sigma V^T$$

with Σ a diagonal matrix with nonnegative diagonal entries

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{min(q,r)} \geq 0.$$

An alternative way of formulating the SVD is:

$$P = \sum_{k=1}^{min(q,r)} \sigma_k U_k V_k^T$$

where U_k, V_k are the k-th column of U and V respectively.

Proposition B.2 Let P be a $q \times r$ matrix with rank n, having a SVD $P = U\Sigma V^T$, and let $k < n$. Denote

$$P_k := U\Sigma_k V^T, \qquad \Sigma_k = \begin{bmatrix} I_k & 0 \end{bmatrix} \Sigma \begin{bmatrix} I_k \\ 0 \end{bmatrix}. \tag{B.3}$$

Then P_k minimizes both

$$\|P - \widetilde{P}\|_2 \quad \text{and} \quad \|P - \widetilde{P}\|_F \tag{B.4}$$

over all matrices $\widetilde{P}$ of rank k.
Additionally

- $\|P - P_k\|_2 = \sigma_{k+1}$, and
- $\|P - P_k\|_F = \left(\sum_{i=k+1}^{min(q,r)} \sigma_i^2 \right)^{\frac{1}{2}}$.

B.3 Projection operations

Let A be an $p \times r$ matrix, and B an $q \times r$ matrix. Then the orthogonal projection of the rows of A onto the row space of B is given by

$$AB^T(BB^T)^{-1}B = AVV^T \tag{B.5}$$

where V is taken from the singular value decomposition $B = U\Sigma V^T$.
Note that this property is related to the least-squares problem

$$\min_X \|A - XB\|_F \tag{B.6}$$

where $X \in \mathbb{R}^{p \times q}$ and $\|\cdot\|_F$ the Frobenius-norm defined by $\|Y\|_F = trace(Y^T Y)$. The solution $\hat{X}$ to (B.6) is given by

$$\hat{X}^T = (BB^T)^{-1}BA^T \tag{B.7}$$

and the projection is thus given by $\hat{X}B = AB^T(BB^T)^{-1}B$.

Let A be an $p \times q$ matrix, and B an $p \times r$ matrix. Then the orthogonal projection of the columns of A onto the column space of B is given by

$$B(B^T B)^{-1} B^T A = UU^T A \tag{B.8}$$

where U is taken from the singular value decomposition $B = U\Sigma V^T$.
This property is related to the least-squares problem

$$\min_X \|A - BX\|_F \tag{B.9}$$

where $X \in \mathbb{R}^{r \times q}$. The solution $\hat{X}$ to (B.9) is given by

$$\hat{X} = (B^T B)^{-1} B^T A \tag{B.10}$$

and the projection is thus given by $B\hat{X} = B(B^T B)^{-1} B^T A$.

For more details on linear algebra see e.g. Noble (1969) and Golub and Van Loan (1983).

B.4 Partial fraction expansion

Let G be a rational transfer function

$$G(z) = \frac{A(z)}{B(z)} = \frac{a_0 + a_1 z^{-1} + a_2 z^{-2} + \cdots + a_m z^{-m}}{1 + b_1 z^{-1} + b_2 z^{-2} + \cdots + b_n z^{-n}}, \tag{B.11}$$

with $m < n$, while

$$B(z) = \prod_{i=1}^{n} (1 - p_i z^{-1}).$$

If all p_i are distinct then the partial fraction expansion of G is given by

$$G(z) = \sum_{i=1}^{n} \frac{\alpha_i}{1 - p_i z^{-1}} \tag{B.12}$$

with

$$\alpha_i = [(1 - p_i z^{-1}) G(z)]_{z=p_i}.$$

If $G(z)$ has distinct poles $p_1, \ldots, p_r$ with multiplicities $\nu_1, \ldots \nu_r$, then the partial fraction expansion of G is given by

$$G(z) = \sum_{i=1}^{r} \sum_{j=1}^{\nu_i} \frac{\alpha_{ij}}{(1 - p_i z^{-1})^j} \tag{B.13}$$

with

$$\alpha_{ij} = \frac{1}{(\nu_i - j)!(-p_i)^{\nu_i - j}} \left[\frac{d^{\nu_i - j}}{d(z^{-1})^{\nu_i - j}} [(1 - p_i z^{-1})^{\nu_i} G(z)] \right]_{z=p_i}.$$

For more details see Oppenheim and Willsky (1997).

B.5 Transfer functions and series expansions

Consider the series expansion

$$\sum_{k=0}^{\infty} a^k z^{-k}$$

with (complex-valued) a, z and $|a| < 1$. This series is convergent for $|az^{-1}| < 1$, leading to

$$\sum_{k=0}^{\infty} (az^{-1})^k = \frac{1}{1 - az^{-1}} \quad \text{for } |z| > |a|. \tag{B.14}$$

Since $|a| < 1$, the rational transfer function has a pole $z = a$ inside the unit circle and the series expansion (B.14) is convergent for z on the unit circle.

In the dual situation, consider the series expansion

$$\sum_{k=1}^{\infty} a^k z^k$$

with $|a| < 1$. This series is convergent for $|az| < 1$, leading to

$$\sum_{k=1}^{\infty} (az)^k = \frac{az}{1 - az} = \frac{-1}{1 - a^{-1}z^{-1}} \quad \text{for } |z| < \frac{1}{|a|}. \tag{B.15}$$

Since $|a| < 1$, the rational transfer function has a pole $z = 1/a$ outside the unit circle and the series expansion (B.15) is convergent for z on the unit circle.

The generic (first-order) component of a causal transfer function $G_+(z)$ that converges on the unit circle is thus given by

$$\frac{b}{1 - az^{-1}} \quad \text{with } |a| < 1$$

where constant terms, related to the term z^0, are considered to be part of the causal part of the system.
For analyzing the anticausal part we write

$$\sum_{k=1}^{\infty} (az)^k = \frac{-1}{1 - a^{-1}z^{-1}}.$$

Therefore the generic (first-order) component of an anticausal transfer function $G_-(z)$ is given by

$$\frac{b}{1 - a^{-1}z^{-1}} \quad \text{with } |a| < 1.$$

When given a general transfer function $G(z)$ there is a unique expansion of G in powers of z^{-1} that is convergent on the unit circle. This expansion decomposes the transfer function into a causal and a noncausal part as follows:
Let $G(z)$ be written as

$$G(z) = \frac{b(z)}{f(z)}$$

with $b(z)$ and $f(z)$ polynomials in z^{-1}, and $f(z) = a(z) \cdot c(z)$ with

$$\begin{aligned} a(z) &= (1 - a_1 z^{-1}) \cdot (1 - a_2 z^{-1}) \cdots (1 - a_n z^{-1}) \\ c(z) &= (1 - c_1^{-1} z^{-1}) \cdot (1 - c_2^{-1} z^{-1}) \cdots (1 - c_m^{-1} z^{-1}) \end{aligned}$$

with $|a_i|, |c_j| \leq 1$ for all i, j. Pole locations a_i are inside the unit circle, while pole locations $1/c_j$ are outside the unit circle.
Decomposition of $G(z)$ now follows by writing

$$G(z) = \frac{g_+(z)}{a(z)} + \frac{g_-(z)}{c(z)} \tag{B.16}$$

with g_+ and g_- polynomials in z^{-1} such that the two quotient expressions $\frac{g_+(z)}{a(z)}$ and $\frac{g_-(z)}{c(z)}$ can both be described as (B.11).
The term $\frac{g_+(z)}{a(z)}$ reflects a causal series expansion $\sum_{k=0}^{\infty} g(k) z^{-k}$ that is convergent on the unit circle; the term $\frac{g_-(z)}{c(z)}$ reflects a strictly anticausal series expansion $\sum_{k=-1}^{-\infty} g(k) z^{-k}$ that is convergent on the unit circle.
Each of the two terms can be written as a summation of first order terms, by applying partial fraction expansion as explained in Appendix B.4. When all roots of $a(z)$ and $c(z)$ are distinct, it follows from the partial fraction expansion that

$$G(z) = \sum_{i=1}^{n} \frac{\alpha_i}{1 - a_i z^{-1}} + \sum_{j=1}^{m} \frac{\beta_j}{1 - c_j^{-1} z^{-1}}$$

The causal part of the system is then given by

$$G_+(z) = \sum_{i=1}^{n} \frac{\alpha_i}{1 - a_i z^{-1}} = \sum_{i=1}^{n} \frac{\alpha_i z}{z - a_i}$$

having poles $z_i = a_i$ within the unit circle, while the anti-causal part is given by

$$G_-(z) = \sum_{j=1}^{m} \frac{\beta_j}{1 - c_j^{-1} z^{-1}} = \sum_{j=1}^{m} \frac{\gamma_j z}{1 - c_j z}$$

having poles $z_j = 1/c_j$ outside the unit circle.

B.6 Trigonometric relations

$$\begin{aligned} 2 \cdot sin(\alpha) \cdot sin(\beta) &= cos(\alpha - \beta) - cos(\alpha + \beta) \\ 2 \cdot cos(\alpha) \cdot cos(\beta) &= cos(\alpha - \beta) + cos(\alpha + \beta) \\ 2 \cdot sin(\alpha) \cdot cos(\beta) &= sin(\alpha - \beta) + sin(\alpha + \beta) \\ cos(\alpha + \beta) &= cos(\alpha) \cdot cos(\beta) - sin(\alpha) \cdot sin(\beta) \\ cos(\alpha - \beta) &= cos(\alpha) \cdot cos(\beta) + sin(\alpha) \cdot sin(\beta) \\ sin(\alpha + \beta) &= sin(\alpha) \cdot cos(\beta) + cos(\alpha) \cdot sin(\beta) \\ sin(\alpha - \beta) &= sin(\alpha) \cdot cos(\beta) - cos(\alpha) \cdot sin(\beta) \end{aligned}$$

Appendix C

Conditional distribution for multivariate Gaussian random variabels

C.1 Determinants and inversion of block matrices

If A and D are square matrices and A is nonsingular, then

$$det \begin{bmatrix} A & B \\ C & D \end{bmatrix} = det(A) \cdot det(D - CA^{-1}B). \tag{B.1}$$

The matrix $D - CA^{-1}B$ is referred to as the Schur complement of matrix A in the block matrix $\begin{bmatrix} A & B \\ C & D \end{bmatrix}$.
If alternatively D is nonsingular then

$$det \begin{bmatrix} A & B \\ C & D \end{bmatrix} = det(D) \cdot det(A - BD^{-1}C). \tag{B.2}$$

Here $A - BD^{-1}C$ is the Schur complement of D within the block matrix.
As a special case it follows that

$$det \begin{bmatrix} A & B \\ 0 & D \end{bmatrix} = det(A) \cdot det(D). \tag{B.3}$$

If A is nonsingular, then

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \left[\begin{array}{c|c} A^{-1} + A^{-1}BECA^{-1} & -A^{-1}BE \\ \hline -ECA^{-1} & E \end{array}\right] \quad \text{with } E = [D - CA^{-1}B]^{-1} \tag{B.4}$$

This result follows from Gaussian elimination, applied to the block matrix. This can be understood by reducing the block matrix to the identity matrix through elementary row operations. Starting with

$$\left[\begin{array}{cc|cc} A & B & I & 0 \\ C & D & 0 & I \end{array}\right]. \tag{B.5}$$

Premultiplying the first block row with A^{-1}, and subtracting C times the new first block row from the second, delivers

$$\left[\begin{array}{cc|cc} I & A^{-1}B & A^{-1} & 0 \\ 0 & D-CA^{-1}B & -CA^{-1} & I \end{array}\right]. \tag{B.6}$$

Multiplying the second block row with $E = (D - CA^{-1}B)^{-1}$ gives

$$\left[\begin{array}{cc|cc} I & A^{-1}B & A^{-1} & 0 \\ 0 & I & -ECA^{-1} & E \end{array}\right]. \tag{B.7}$$

Subtracting $A^{-1}B$ times the second block row from the first one, then provides

$$\left[\begin{array}{cc|cc} I & 0 & A^{-1}+A^{-1}BECA^{-1} & -A^{-1}BE \\ 0 & I & -ECA^{-1} & E \end{array}\right]. \tag{B.8}$$

The above result can equivalently be applied to the situation when not A but D is considered to be invertible. The related result in that situation is:

$$\left[\begin{array}{cc} A & B \\ C & D \end{array}\right]^{-1} = \left[\begin{array}{c|c} F & -FBD^{-1} \\ \hline -D^{-1}CF & D^{-1}+D^{-1}CFBD^{-1} \end{array}\right] \quad \text{with } F = [A - BD^{-1}C]^{-1} \tag{B.9}$$

Equating the two expressions for the inversion of the block matrix leads -among others- to the *matrix inversion lemma* :

$$(A - BD^{-1}C)^{-1} = A^{-1} + A^{-1}B(D - CA^{-1}B)^{-1}CA^{-1}. \tag{B.10}$$

C.2 Multivariate conditional Gaussian distribution

Proposition B.1 Let random variable $\mathbf{z}$ have an $m+n$ multivariate Gaussian distribution, while $\mathbf{z}$ is partitioned according to

$$\mathbf{z} = \left[\begin{array}{c} \mathbf{x} \\ \mathbf{y} \end{array}\right] \quad \mu_{\mathbf{z}} = \left[\begin{array}{c} \mu_{\mathbf{x}} \\ \mu_{\mathbf{y}} \end{array}\right] \quad \Sigma_{\mathbf{z}} = \left[\begin{array}{cc} \Sigma_{\mathbf{x}} & \Sigma_{\mathbf{xy}} \\ \Sigma^*_{\mathbf{xy}} & \Sigma_{\mathbf{y}} \end{array}\right] \tag{B.11}$$

with $dim(\mathbf{x}) = m$ and $dim(\mathbf{y}) = n$, then the conditional distribution $f_{\mathbf{x}|\mathbf{y}}(x|y)$ is given by

$$f_{\mathbf{x}|\mathbf{y}}(x|y) = \frac{1}{(2\pi)^{m/2}\sqrt{det(\Sigma_{\mathbf{x}|\mathbf{y}})}} \cdot e^{-\frac{1}{2}(x-\mu(y))^T\Sigma^{-1}_{\mathbf{x}|\mathbf{y}}(x-\mu(y))} \tag{B.12}$$

with

$$\mu(y) = \mu_{\mathbf{x}} + \Sigma_{\mathbf{xy}}\Sigma^{-1}_{\mathbf{y}}(y - \mu_{\mathbf{y}}) \tag{B.13}$$

$$\Sigma_{\mathbf{x}|\mathbf{y}} = \Sigma_{\mathbf{x}} - \Sigma_{\mathbf{xy}}\Sigma^{-1}_{\mathbf{y}}\Sigma^*_{\mathbf{xy}}. \tag{B.14}$$

The result of this proposition can be verified by using Bayes' relation:

$$f_{\mathbf{x}|\mathbf{y}}(x|y) = \frac{f_{\mathbf{x},\mathbf{y}}(x,y)}{f_{\mathbf{y}}(y)}$$

with

$$f_{\mathbf{x},\mathbf{y}}(x,y) = \frac{1}{(2\pi)^{(m+n)/2}\sqrt{det(\Sigma_{\mathbf{z}})}} \cdot e^{-\frac{1}{2}[(x-\mu_{\mathbf{x}})^T\ (y-\mu_{\mathbf{y}})^T]\Sigma_{\mathbf{z}}^{-1}[(x-\mu_{\mathbf{x}})^T\ (y-\mu_{\mathbf{y}})^T]^T} \quad \text{(B.15)}$$

$$f_{\mathbf{y}}(y) = \frac{1}{(2\pi)^{n/2}\sqrt{det(\Sigma_{\mathbf{y}})}} \cdot e^{-\frac{1}{2}(y-\mu_{\mathbf{y}})^T\Sigma_{\mathbf{y}}^{-1}(y-\mu_{\mathbf{y}})} \quad \text{(B.16)}$$

The fact that $det(\Sigma_{\mathbf{z}}) = det(\Sigma_{\mathbf{y}}) \cdot det(\Sigma_{\mathbf{x}|\mathbf{y}})$ follows simply from the determinant rule (B.2), thus proving the correctness of the coefficient that multiplies the exponential in (B.12). In order to verify the exponential term, note that the exponential term of the quotient of (B.15) and (B.16) can be written as

$$-\frac{1}{2}[(x-\mu_{\mathbf{x}})^T\ (y-\mu_{\mathbf{y}})^T]\left[\Sigma_{\mathbf{z}}^{-1} - \begin{pmatrix} 0 & 0 \\ 0 & \Sigma_{\mathbf{y}}^{-1} \end{pmatrix}\right]\begin{pmatrix} x-\mu_{\mathbf{x}} \\ y-\mu_{\mathbf{y}} \end{pmatrix}. \quad \text{(B.17)}$$

Applying the matrix inversion (B.9) to $\Sigma_{\mathbf{z}}^{-1}$, using the decomposition (B.11), the above expression reduces to

$$-\frac{1}{2}[(x-\mu_{\mathbf{x}})^T\ (y-\mu_{\mathbf{y}})^T]\left[\left(\begin{array}{c|c} F & -F\Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}-1 \\ \hline -\Sigma_{\mathbf{y}}^{-1}\Sigma_{\mathbf{xy}}^{*}F & \Sigma_{\mathbf{y}}^{-1}\Sigma_{\mathbf{xy}}^{*}F\Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}^{-1} \end{array}\right)\right]\begin{pmatrix} x-\mu_{\mathbf{x}} \\ y-\mu_{\mathbf{y}} \end{pmatrix}, \quad \text{(B.18)}$$

with $F = [\Sigma_{\mathbf{x}} - \Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}^{-1}\Sigma_{\mathbf{xy}}^{*}]^{-1}$.
The block matrix in (B.18) now has become singular and it can be decomposed as

$$\begin{pmatrix} I \\ -\Sigma_{\mathbf{y}}^{-1}\Sigma_{\mathbf{xy}}^{*} \end{pmatrix} \cdot F \cdot [I \quad -\Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}^{-1}]. \quad \text{(B.19)}$$

As a result the full exponential term becomes

$$-\frac{1}{2}[(x-\mu_{\mathbf{x}})^T - (y-\mu_{\mathbf{y}})^T\Sigma_{\mathbf{y}}^{-1}\Sigma_{\mathbf{xy}}^{*}] \cdot F \cdot [x-\mu_{\mathbf{x}} - \Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}^{-1}(y-\mu_{\mathbf{y}})] \quad \text{(B.20)}$$

showing that $m(y) = \mu_{\mathbf{x}} + \Sigma_{\mathbf{xy}}\Sigma_{\mathbf{y}}^{-1}(y-\mu_{\mathbf{y}})$ and $\Sigma_{\mathbf{x}|\mathbf{y}} = F^{-1}$.

Bibliography

Bartlett, M.S. (1946). On the theoretical specification and sampling properties of autocorrelated time series. *J. Royal Statistical Society, B 8*, no. 27.

Brillinger, D.R. (1981). *Time Series - Data Analysis and Theory.* McGraw-Hill, New York, expanded edition.

Cramér, H. (1946). *Mathematical Methods of Statistics.* Princeton University Press, Princeton, New Jersey.

Dekking, F.M., C. Kraaikamp, H.P. Lopuhaä and L.E. Meester (2005). *A Modern Introduction to Probability and Statistics* , Springer texts in statistics, Springer-Verlag, London.

Fante, R.L. (1988). *Signal Analysis and Estimation - An Introduction.* John Wiley & Sons, Inc., New York.

Fisher, R.A. (1922). On the mathematical foundations of theoretical statistics. *Phil. Trans. Roy. Soc. London, Vol. 222*, pp. 309-368.

Fletcher, R.R. (1980). *Practical Methods of Optimization.* John Wiley, UK.

Frieden, B.R. (2001). *Probability, Statistical Optics, and Data Testing.* Springer Series in Information Sciences, Vol. 10. Springer Verlag, Berlin, third edition.

Gleser, L.J. (1981). Estimation in a multivariate "errors in variables" regression model: large sample results. *Ann. Statistics, 9*, pp. 24-44.

Golub, G. and C.I. Van Loan (1983). *Matrix Computations.* John Hopkins Univ. Press, Baltimore.

Hayes, M.H. (1996). *Statistical Digital Signal Processing and Modeling.* John Wiley & Sons Inc., New York, 1996.

Jenkins, G.M. and D.G. Watts (1968). *Spectral Analysis and its Applications.* Holden-Day, Oakland, CA.

Ljung, L. (1999). *System Identification - Theory for the User.* Prentice-Hall, Englewood Cliffs, NJ. Second edition, 1999.

Miller, R.E. (2000). *Optimization - Foundations and Applications.* John Wiley & Sons, New York.

Mood, A.M., F.A. Graybill and D.C. Boes (1974). *Introduction to the Theory of Statistics.* McGraw-Hill Kogakusha, Ltd., Tokyo, Japan.

Noble, B.(1969). *Applied Linear Algebra.* Prentice-Hall, Englewood Cliffs, NJ.

Oppenheim, A.V. and A.S. Willsky (1997). *Signals and Sytems.* Prentice-Hall International, Inc., London, Second Edition.

Papoulis, A. (1991). *Probability, Random Variables, and Stochastic Processes.* McGraw-Hill, third edition.

Phillips, C.L. and J.M. Parr (1995). *Signals, Systems and Transformations.* Prentice Hall, Englewood Cliffs, NJ.

Porat, B. (1994). Digital Processing of Random Signals. Prentice-Hall, Englewood Cliffs,

NJ, 1994.

Priestley, M.B. (1981). *Spectral Analysis and Time Series.* Academic Press, London, UK.

Rao, C.R. (1945). Information and accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc., Vol. 37*, pp. 81-91.

Shanmugan, K.S. and A.M. Breipohl (1988). *Random Signals; Detection, Estimation and Data Analysis.* John Wiley & Sons Inc., New York.

Therrien, C.W. (1992). *Discrete Random Signals and Statistical Signal Processing.* Prentice Hall, Englewood Cliffs, NJ.

van den Bos, A. (1999). Measurement errors. in: J.G. Webster (Ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering, Vol. 12.* J. Wiley & Sons, Inc., New York, pp. 448-459.

van den Bos, A. (2007). *Parameter Estimation for Scientists and Engineers.* John Wiley & Sons, Inc., Hoboken, NJ.

Van Huffel, S. and J. Vandewalle (1989). On the accuracy of total least squares and least squares techniques in the presence of errors on all data. *Automatica, 25*, no. 5, pp. 765-769.

Van Huffel, S. and J. Vandewalle (1991). *The Total Least Squares Problem: Computational Aspects and Analysis.* SIAM Series Frontiers in Applied Mathematics, Vol. 9. SIAM, Philadelphia.

Index