


PRENTICE HALL SIGNAL PROCESSING SERIES

Alan V. Oppenheim, Series Editor

ANDREWS AND HUNT *Digital Image Restoration*
BRIGHAM *The Fast Fourier Transform*
BRIGHAM *The Fast Fourier Transform and Its Applications*
BURDIC *Underwater Acoustic System Analysis, 2/E*
CASTLEMAN *Digital Image Processing*
COWAN AND GRANT *Adaptive Filters*
CROCHIERE AND RABINER *Multirate Digital Signal Processing*
DUDGEON AND MERSEREAU *Multidimensional Digital Signal Processing*
HAMMING *Digital Filters, 3/E*
HAYKIN, ED. *Advances in Spectrum Analysis and Array Processing, Vols. I & II*
HAYKIN, ED. *Array Signal Processing*
JAYANT AND NOLL *Digital Coding of Waveforms*
JOHNSON AND DUDGEON *Array Signal Processing: Concepts and Techniques*
KAY *Fundamentals of Statistical Signal Processing: Estimation Theory*
KAY *Modern Spectral Estimation*
KINO *Acoustic Waves: Devices, Imaging, and Analog Signal Processing*
LEA, ED. *Trends in Speech Recognition*
LIM *Two-Dimensional Signal and Image Processing*
LIM, ED. *Speech Enhancement*
LIM AND OPPENHEIM, EDS. *Advanced Topics in Signal Processing*
MARPLE *Digital Spectral Analysis with Applications*
MCCLELLAN AND RADER *Number Theory in Digital Signal Processing*
MENDEL *Lessons in Digital Estimation Theory*
OPPENHEIM, ED. *Applications of Digital Signal Processing*
OPPENHEIM AND NAWAB, EDS. *Symbolic and Knowledge-Based Signal Processing*
OPPENHEIM, WILLSKY, WITH YOUNG *Signals and Systems*
OPPENHEIM AND SCHAFER *Digital Signal Processing*
OPPENHEIM AND SCHAFER *Discrete-Time Signal Processing*
QUACKENBUSH ET AL. *Objective Measures of Speech Quality*
RABINER AND GOLD *Theory and Applications of Digital Signal Processing*
RABINER AND SCHAFER *Digital Processing of Speech Signals*
ROBINSON AND TREITEL *Geophysical Signal Analysis*
STEARNS AND DAVID *Signal Processing Algorithms*
STEARNS AND HUSH *Digital Signal Analysis, 2/E*
TRIBOLET *Seismic Applications of Homomorphic Signal Processing*
VAIDYANATHAN *Multirate Systems and Filter Banks*
WIDROW AND STEARNS *Adaptive Signal Processing*

Fundamentals of Statistical Signal Processing: Estimation Theory

Steven M. Kay
University of Rhode Island



Prentice Hall PTR
Upper Saddle River, NJ 07458

Contents

Preface	xi
1 Introduction	1
1.1 Estimation in Signal Processing	1
1.2 The Mathematical Estimation Problem	7
1.3 Assessing Estimator Performance	9
1.4 Some Notes to the Reader	12
2 Minimum Variance Unbiased Estimation	15
2.1 Introduction	15
2.2 Summary	15
2.3 Unbiased Estimators	16
2.4 Minimum Variance Criterion	19
2.5 Existence of the Minimum Variance Unbiased Estimator	20
2.6 Finding the Minimum Variance Unbiased Estimator	21
2.7 Extension to a Vector Parameter	22
3 Cramer-Rao Lower Bound	27
3.1 Introduction	27
3.2 Summary	27
3.3 Estimator Accuracy Considerations	28
3.4 Cramer-Rao Lower Bound	30
3.5 General CRLB for Signals in White Gaussian Noise	35
3.6 Transformation of Parameters	37
3.7 Extension to a Vector Parameter	39
3.8 Vector Parameter CRLB for Transformations	45
3.9 CRLB for the General Gaussian Case	47
3.10 Asymptotic CRLB for WSS Gaussian Random Processes	50
3.11 Signal Processing Examples	53
3A Derivation of Scalar Parameter CRLB	67
3B Derivation of Vector Parameter CRLB	70
3C Derivation of General Gaussian CRLB	73
3D Derivation of Asymptotic CRLB	77

4	Linear Models	83
4.1	Introduction	83
4.2	Summary	83
4.3	Definition and Properties	83
4.4	Linear Model Examples	86
4.5	Extension to the Linear Model	94
5	General Minimum Variance Unbiased Estimation	101
5.1	Introduction	101
5.2	Summary	101
5.3	Sufficient Statistics	102
5.4	Finding Sufficient Statistics	104
5.5	Using Sufficiency to Find the MVU Estimator	107
5.6	Extension to a Vector Parameter	116
5A	Proof of Neyman-Fisher Factorization Theorem (Scalar Parameter) . . .	127
5B	Proof of Rao-Blackwell-Lehmann-Scheffe Theorem (Scalar Parameter) .	130
6	Best Linear Unbiased Estimators	133
6.1	Introduction	133
6.2	Summary	133
6.3	Definition of the BLUE	134
6.4	Finding the BLUE	136
6.5	Extension to a Vector Parameter	139
6.6	Signal Processing Example	141
6A	Derivation of Scalar BLUE	151
6B	Derivation of Vector BLUE	153
7	Maximum Likelihood Estimation	157
7.1	Introduction	157
7.2	Summary	157
7.3	An Example	158
7.4	Finding the MLE	162
7.5	Properties of the MLE	164
7.6	MLE for Transformed Parameters	173
7.7	Numerical Determination of the MLE	177
7.8	Extension to a Vector Parameter	182
7.9	Asymptotic MLE	190
7.10	Signal Processing Examples	191
7A	Monte Carlo Methods	205
7B	Asymptotic PDF of MLE for a Scalar Parameter	211
7C	Derivation of Conditional Log-Likelihood for EM Algorithm Example .	214
8	Least Squares	219
8.1	Introduction	219
8.2	Summary	219

8.3	The Least Squares Approach	220
8.4	Linear Least Squares	223
8.5	Geometrical Interpretations	226
8.6	Order-Recursive Least Squares	232
8.7	Sequential Least Squares	242
8.8	Constrained Least Squares	251
8.9	Nonlinear Least Squares	254
8.10	Signal Processing Examples	260
8A	Derivation of Order-Recursive Least Squares	282
8B	Derivation of Recursive Projection Matrix	285
8C	Derivation of Sequential Least Squares	286
9	Method of Moments	289
9.1	Introduction	289
9.2	Summary	289
9.3	Method of Moments	289
9.4	Extension to a Vector Parameter	292
9.5	Statistical Evaluation of Estimators	294
9.6	Signal Processing Example	299
10	The Bayesian Philosophy	309
10.1	Introduction	309
10.2	Summary	309
10.3	Prior Knowledge and Estimation	310
10.4	Choosing a Prior PDF	316
10.5	Properties of the Gaussian PDF	321
10.6	Bayesian Linear Model	325
10.7	Nuisance Parameters	328
10.8	Bayesian Estimation for Deterministic Parameters	330
10A	Derivation of Conditional Gaussian PDF	337
11	General Bayesian Estimators	341
11.1	Introduction	341
11.2	Summary	341
11.3	Risk Functions	342
11.4	Minimum Mean Square Error Estimators	344
11.5	Maximum A Posteriori Estimators	350
11.6	Performance Description	359
11.7	Signal Processing Example	365
11A	Conversion of Continuous-Time System to Discrete-Time System	375
12	Linear Bayesian Estimators	379
12.1	Introduction	379
12.2	Summary	379
12.3	Linear MMSE Estimation	380

12.4 Geometrical Interpretations	384
12.5 The Vector LMMSE Estimator	389
12.6 Sequential LMMSE Estimation	392
12.7 Signal Processing Examples - Wiener Filtering	400
12A Derivation of Sequential LMMSE Estimator	415
13 Kalman Filters	419
13.1 Introduction	419
13.2 Summary	419
13.3 Dynamical Signal Models	420
13.4 Scalar Kalman Filter	431
13.5 Kalman Versus Wiener Filters	442
13.6 Vector Kalman Filter	446
13.7 Extended Kalman Filter	449
13.8 Signal Processing Examples	452
13A Vector Kalman Filter Derivation	471
13B Extended Kalman Filter Derivation	476
14 Summary of Estimators	479
14.1 Introduction	479
14.2 Estimation Approaches	479
14.3 Linear Model	486
14.4 Choosing an Estimator	489
15 Extensions for Complex Data and Parameters	493
15.1 Introduction	493
15.2 Summary	493
15.3 Complex Data and Parameters	494
15.4 Complex Random Variables and PDFs	500
15.5 Complex WSS Random Processes	513
15.6 Derivatives, Gradients, and Optimization	517
15.7 Classical Estimation with Complex Data	524
15.8 Bayesian Estimation	532
15.9 Asymptotic Complex Gaussian PDF	535
15.10 Signal Processing Examples	539
15A Derivation of Properties of Complex Covariance Matrices	555
15B Derivation of Properties of Complex Gaussian PDF	558
15C Derivation of CRLB and MLE Formulas	563
A1 Review of Important Concepts	567
A1.1 Linear and Matrix Algebra	567
A1.2 Probability, Random Processes, and Time Series Models	574
A2 Glossary of Symbols and Abbreviations	583
INDEX	589

Preface

Parameter estimation is a subject that is standard fare in the many books available on statistics. These books range from the highly theoretical expositions written by statisticians to the more practical treatments contributed by the many users of applied statistics. This text is an attempt to strike a balance between these two extremes. The particular audience we have in mind is the community involved in the design and implementation of signal processing algorithms. As such, the primary focus is on obtaining optimal estimation algorithms that may be implemented on a digital computer. The data sets are therefore assumed to be samples of a continuous-time waveform or a sequence of data points. The choice of topics reflects what we believe to be the important approaches to obtaining an optimal estimator and analyzing its performance. As a consequence, some of the deeper theoretical issues have been omitted with references given instead.

It is the author's opinion that the best way to assimilate the material on parameter estimation is by exposure to and working with good examples. Consequently, there are numerous examples that illustrate the theory and others that apply the theory to actual signal processing problems of current interest. Additionally, an abundance of homework problems have been included. They range from simple applications of the theory to extensions of the basic concepts. A solutions manual is available from the publisher. To aid the reader, summary sections have been provided at the beginning of each chapter. Also, an overview of all the principal estimation approaches and the rationale for choosing a particular estimator can be found in Chapter 14. Classical estimation is first discussed in Chapters 2-9, followed by Bayesian estimation in Chapters 10-13. This delineation will, hopefully, help to clarify the basic differences between these two principal approaches. Finally, again in the interest of clarity, we present the estimation principles for scalar parameters first, followed by their vector extensions. This is because the matrix algebra required for the vector estimators can sometimes obscure the main concepts.

This book is an outgrowth of a one-semester graduate level course on estimation theory given at the University of Rhode Island. It includes somewhat more material than can actually be covered in one semester. We typically cover most of Chapters 1-12, leaving the subjects of Kalman filtering and complex data/parameter extensions to the student. The necessary background that has been assumed is an exposure to the basic theory of digital signal processing, probability and random processes, and linear

and matrix algebra. This book can also be used for self-study and so should be useful to the practicing engineer as well as the student.

The author would like to acknowledge the contributions of the many people who over the years have provided stimulating discussions of research problems, opportunities to apply the results of that research, and support for conducting research. Thanks are due to my colleagues L. Jackson, R. Kumaresan, L. Pakula, and D. Tufts of the University of Rhode Island, and L. Scharf of the University of Colorado. Exposure to practical problems, leading to new research directions, has been provided by H. Woodsum of Sonetech, Bedford, New Hampshire, and by D. Mook, S. Lang, C. Myers, and D. Morgan of Lockheed-Sanders, Nashua, New Hampshire. The opportunity to apply estimation theory to sonar and the research support of J. Kelly of the Naval Undersea Warfare Center, Newport, Rhode Island, J. Salisbury of Analysis and Technology, Middletown, Rhode Island (formerly of the Naval Undersea Warfare Center), and D. Sheldon of the Naval Undersea Warfare Center, New London, Connecticut, are also greatly appreciated. Thanks are due to J. Sjogren of the Air Force Office of Scientific Research, whose continued support has allowed the author to investigate the field of statistical estimation. A debt of gratitude is owed to all my current and former graduate students. They have contributed to the final manuscript through many hours of pedagogical and research discussions as well as by their specific comments and questions. In particular, P. Djurić of the State University of New York proofread much of the manuscript, and V. Nagesha of the University of Rhode Island proofread the manuscript and helped with the problem solutions.

—

Steven M. Kay
University of Rhode Island
Kingston, RI 02881

Chapter 1

Introduction

1.1 Estimation in Signal Processing

Modern estimation theory can be found at the heart of many electronic signal processing systems designed to extract information. These systems include

1. Radar
2. Sonar
3. Speech
4. Image analysis
5. Biomedicine
6. Communications
7. Control
8. Seismology,

and all share the common problem of needing to estimate the values of a group of parameters. We briefly describe the first three of these systems. In radar we are interested in determining the position of an aircraft, as for example, in airport surveillance radar [Skolnik 1980]. To determine the range R we transmit an electromagnetic pulse that is reflected by the aircraft, causing an echo to be received by the antenna τ_0 seconds later, as shown in Figure 1.1a. The range is determined by the equation $\tau_0 = 2R/c$, where c is the speed of electromagnetic propagation. Clearly, if the round trip delay τ_0 can be measured, then so can the range. A typical transmit pulse and received waveform are shown in Figure 1.1b. The received echo is decreased in amplitude due to propagation losses and hence may be obscured by environmental noise. Its onset may also be perturbed by time delays introduced by the electronics of the receiver. Determination of the round trip delay can therefore require more than just a means of detecting a jump in the power level at the receiver. It is important to note that a typical modern




Figure 1.1 Radar system

radar system will input the received continuous-time waveform into a digital computer by taking samples via an analog-to-digital converter. Once the waveform has been sampled, the data compose a *time series*. (See also Examples 3.13 and 7.15 for a more detailed description of this problem and optimal estimation procedures.)

Another common application is in sonar, in which we are also interested in the position of a target, such as a submarine [Knight et al. 1981, Burdic 1984]. A typical passive sonar is shown in Figure 1.2a. The target radiates noise due to machinery on board, propellor action, etc. This noise, which is actually the *signal* of interest, propagates through the water and is received by an array of sensors. The sensor outputs





Figure 1.2 Passive sonar system

are then transmitted to a tow ship for input to a digital computer. Because of the positions of the sensors relative to the arrival angle of the target signal, we receive the signals shown in Figure 1.2b. By measuring τ_0 , the delay between sensors, we can determine the bearing β from the expression

$$\beta = \arccos\left(\frac{c\tau_0}{d}\right) \quad (1.1)$$

where c is the speed of sound in water and d is the distance between sensors (see Examples 3.15 and 7.17 for a more detailed description). Again, however, the received




Figure 1.3 Examples of speech sounds

waveforms are not “clean” as shown in Figure 1.2b but are embedded in noise, making the determination of τ_0 more difficult. The value of β obtained from (1.1) is then only an estimate.

Another application is in speech processing systems [Rabiner and Schafer 1978]. A particularly important problem is speech recognition, which is the recognition of speech by a machine (digital computer). The simplest example of this is in recognizing individual speech sounds or phonemes. Phonemes are the vowels, consonants, etc., or the fundamental sounds of speech. As an example, the vowels /a/ and /e/ are shown in Figure 1.3. Note that they are periodic waveforms whose period is called the pitch. To recognize whether a sound is an /a/ or an /e/ the following simple strategy might be employed. Have the person whose speech is to be recognized say each vowel three times and store the waveforms. To recognize the spoken vowel, compare it to the stored vowels and choose the one that is closest to the spoken vowel or the one that




Figure 1.4 LPC spectral modeling

minimizes some distance measure. Difficulties arise if the pitch of the speaker's voice changes from the time he or she records the sounds (the training session) to the time when the speech recognizer is used. This is a natural variability due to the nature of human speech. In practice, attributes, other than the waveforms themselves, are used to measure distance. Attributes are chosen that are less susceptible to variation. For example, the spectral envelope will not change with pitch since the Fourier transform of a periodic signal is a sampled version of the Fourier transform of one period of the signal. The period affects only the spacing between frequency samples, not the values. To extract the spectral envelope we employ a model of speech called *linear predictive coding* (LPC). The parameters of the model determine the spectral envelope. For the speech sounds in Figure 1.3 the power spectrum (magnitude-squared Fourier transform divided by the number of time samples) or periodogram and the estimated LPC spectral envelope are shown in Figure 1.4. (See Examples 3.16 and 7.18 for a description of how

the parameters of the model are estimated and used to find the spectral envelope.) It is interesting that in this example a human interpreter can easily discern the spoken vowel. The real problem then is to design a machine that is able to do the same. In the radar/sonar problem a human interpreter would be unable to determine the target position from the received waveforms, so that the machine acts as an indispensable tool.

In all these systems we are faced with the problem of extracting values of parameters based on continuous-time waveforms. Due to the use of digital computers to sample and store the continuous-time waveform, we have the equivalent problem of extracting parameter values from a discrete-time waveform or a data set. Mathematically, we have the N -point data set $\{x[0], x[1], \dots, x[N-1]\}$ which depends on an unknown parameter θ . We wish to determine θ based on the data or to define an *estimator*

$$\hat{\theta} = g(x[0], x[1], \dots, x[N-1]) \quad (1.2)$$

where g is some function. This is the problem of *parameter estimation*, which is the subject of this book. Although electrical engineers at one time designed systems based on analog signals and analog circuits, the current and future trend is based on discrete-time signals or sequences and digital circuitry. With this transition the estimation problem has evolved into one of estimating a parameter based on a *time series*, which is just a discrete-time process. Furthermore, because the amount of data is necessarily finite, we are faced with the determination of g as in (1.2). Therefore, our problem has now evolved into one which has a long and glorious history, dating back to Gauss who in 1795 used least squares data analysis to predict planetary movements [Gauss 1963 (English translation)]. All the theory and techniques of statistical estimation are at our disposal [Cox and Hinkley 1974, Kendall and Stuart 1976–1979, Rao 1973, Zacks 1981].

Before concluding our discussion of application areas we complete the previous list.

4. Image analysis - estimate the position and orientation of an object from a camera image, necessary when using a robot to pick up an object [Jain 1989]
5. Biomedicine - estimate the heart rate of a fetus [Widrow and Stearns 1985]
6. Communications - estimate the carrier frequency of a signal so that the signal can be demodulated to baseband [Proakis 1983]
7. Control - estimate the position of a powerboat so that corrective navigational action can be taken, as in a LORAN system [Dabbous 1988]
8. Seismology - estimate the underground distance of an oil deposit based on sound reflections due to the different densities of oil and rock layers [Justice 1985].

Finally, the multitude of applications stemming from analysis of data from physical experiments, economics, etc., should also be mentioned [Box and Jenkins 1970, Holm and Hovem 1979, Schuster 1898, Taylor 1986].




Figure 1.5 Dependence of PDF on unknown parameter

1.2 The Mathematical Estimation Problem

In determining good estimators the first step is to mathematically model the data. Because the data are inherently random, we describe it by its probability density function (PDF) or $p(x[0], x[1], \dots, x[N-1]; \theta)$. The PDF is *parameterized* by the unknown parameter θ , i.e., we have a class of PDFs where each one is different due to a different value of θ . We will use a semicolon to denote this dependence. As an example, if $N = 1$ and θ denotes the mean, then the PDF of the data might be

$$p(x[0]; \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} (x[0] - \theta)^2 \right]$$

which is shown in Figure 1.5 for various values of θ . It should be intuitively clear that because the value of θ affects the probability of $x[0]$, we should be able to *infer the value of θ from the observed value of $x[0]$* . For example, if the value of $x[0]$ is negative, it is doubtful that $\theta = \theta_2$. The value $\theta = \theta_1$ might be more reasonable. This specification of the PDF is critical in determining a good estimator. In an actual problem we are not given a PDF but must choose one that is not only consistent with the problem constraints and any prior knowledge, but one that is also mathematically tractable. To illustrate the approach consider the hypothetical Dow-Jones industrial average shown in Figure 1.6. It might be conjectured that this data, although appearing to fluctuate wildly, actually is "on the average" increasing. To determine if this is true we could assume that the data actually consist of a straight line embedded in random noise or

$$x[n] = A + Bn + w[n] \quad n = 0, 1, \dots, N-1.$$

A reasonable model for the noise is that $w[n]$ is white Gaussian noise (WGN) or each sample of $w[n]$ has the PDF $\mathcal{N}(0, \sigma^2)$ (denotes a Gaussian distribution with a mean of 0 and a variance of σ^2) and is uncorrelated with all the other samples. Then, the unknown parameters are A and B , which arranged as a vector become the vector parameter $\theta = [A \ B]^T$. Letting $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[N-1]]^T$, the PDF is

$$p(\mathbf{x}; \theta) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn)^2 \right]. \quad (1.3)$$

The choice of a straight line for the signal component is consistent with the knowledge that the Dow-Jones average is hovering around 3000 (A models this) and the conjecture




Figure 1.6 Hypothetical Dow-Jones average

that it is increasing ($B > 0$ models this). The assumption of WGN is justified by the need to formulate a mathematically tractable model so that closed form estimators can be found. Also, it is reasonable unless there is strong evidence to the contrary, such as highly correlated noise. Of course, the performance of any estimator obtained will be critically dependent on the PDF assumptions. We can only hope the estimator obtained is robust, in that slight changes in the PDF do not severely affect the performance of the estimator. More conservative approaches utilize robust statistical procedures [Huber 1981].

Estimation based on PDFs such as (1.3) is termed *classical estimation* in that the parameters of interest are assumed to be deterministic but unknown. In the Dow-Jones average example we know a priori that the mean is somewhere around 3000. It seems inconsistent with reality, then, to choose an estimator of A that can result in values as low as 2000 or as high as 4000. We might be more willing to constrain the estimator to produce values of A in the range $[2800, 3200]$. To incorporate this prior knowledge we can assume that A is no longer deterministic but a random variable and assign it a PDF, possibly uniform over the $[2800, 3200]$ interval. Then, any subsequent estimator will yield values in this range. Such an approach is termed *Bayesian estimation.* The parameter we are attempting to estimate is then viewed as a *realization* of the random variable θ . As such, the data are described by the *joint PDF*

$$p(\mathbf{x}, \theta) = p(\mathbf{x}|\theta)p(\theta)$$

where $p(\theta)$ is the prior PDF, summarizing our knowledge about θ before any data are observed, and $p(\mathbf{x}|\theta)$ is a conditional PDF, summarizing our knowledge provided by the data $\mathbf{x}$ conditioned on knowing θ . The reader should compare the notational differences between $p(\mathbf{x}; \theta)$ (a family of PDFs) and $p(\mathbf{x}|\theta)$ (a conditional PDF), as well as the implied interpretations (see also Problem 1.3).




Figure 1.7 Realization of DC level in noise

Once the PDF has been specified, the problem becomes one of determining an optimal estimator or function of the data, as in (1.2). Note that an estimator may depend on other parameters, but only if they are *known*. An estimator may be thought of as a rule that assigns a value to θ for each realization of $\mathbf{x}$. The estimate of θ is the value of θ obtained for a given realization of $\mathbf{x}$. This distinction is analogous to a random variable (which is a function defined on the sample space) and the value it takes on. Although some authors distinguish between the two by using capital and lowercase letters, we will not do so. The meaning will, hopefully, be clear from the context.

1.3 Assessing Estimator Performance

Consider the data set shown in Figure 1.7. From a cursory inspection it appears that $x[n]$ consists of a DC level A in noise. (The use of the term DC is in reference to direct current, which is equivalent to the constant function.) We could model the data as

$$x[n] = A + w[n]$$

where $w[n]$ denotes some zero mean noise process. Based on the data set $\{x[0], x[1], \dots, x[N-1]\}$, we would like to estimate A . Intuitively, since A is the average level of $x[n]$ ($w[n]$ is zero mean), it would be reasonable to estimate A as

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

or by the sample mean of the data. Several questions come to mind:

1. How close will $\hat{A}$ be to A ?
2. Are there better estimators than the sample mean?

For the data set in Figure 1.7 it turns out that $\hat{A} = 0.9$, which is close to the true value of $A = 1$. Another estimator might be

$$\check{A} = x[0].$$

Intuitively, we would not expect this estimator to perform as well since it does not make use of all the data. There is no *averaging* to reduce the noise effects. However, for the data set in Figure 1.7, $\check{A} = 0.95$, which is closer to the true value of A than the sample mean estimate. Can we conclude that $\check{A}$ is a better estimator than $\hat{A}$? The answer is of course no. Because an estimator is a function of the data, which are random variables, it too is a random variable, subject to many possible outcomes. The fact that $\check{A}$ is closer to the true value only means that for the *given realization of data*, as shown in Figure 1.7, the estimate $\check{A} = 0.95$ (or realization of $\check{A}$) is closer to the true value than the estimate $\hat{A} = 0.9$ (or realization of $\hat{A}$). To assess performance we must do so statistically. One possibility would be to repeat the experiment that generated the data and apply each estimator to every data set. Then, we could ask which estimator produces a better estimate in the majority of the cases. Suppose we repeat the experiment by fixing $A = 1$ and adding different noise realizations of $w[n]$ to generate an ensemble of realizations of $x[n]$. Then, we determine the values of the two estimators for each data set and finally plot the histograms. (A histogram describes the number of times the estimator produces a given range of values and is an approximation to the PDF.) For 100 realizations the histograms are shown in Figure 1.8. It should now be evident that $\hat{A}$ is a better estimator than $\check{A}$ because the values obtained are more concentrated about the true value of $A = 1$. Hence, $\hat{A}$ will *usually* produce a value closer to the true one than $\check{A}$. The skeptic, however, might argue that if we repeat the experiment 1000 times instead, then the histogram of $\hat{A}$ will be more concentrated. To dispel this notion, we cannot repeat the experiment 1000 times, for surely the skeptic would then reassert his or her conjecture for 10,000 experiments. To *prove* that $\hat{A}$ is better we could establish that the variance is less. The modeling assumptions that we must employ are that the $w[n]$'s, in addition to being zero mean, are uncorrelated and have equal variance σ^2 . Then, we first show that the mean of each estimator is the true value or

$$\begin{aligned} E(\hat{A}) &= E\left(\frac{1}{N} \sum_{n=0}^{N-1} x[n]\right) \\ &= \frac{1}{N} \sum_{n=0}^{N-1} E(x[n]) \\ &= A \\ E(\check{A}) &= E(x[0]) \\ &= A \end{aligned}$$

so that on the average the estimators produce the true value. Second, the variances are

$$\text{var}(\hat{A}) = \text{var}\left(\frac{1}{N} \sum_{n=0}^{N-1} x[n]\right)$$




Figure 1.8 Histograms for sample mean and first sample estimator

$$\begin{aligned}
 &= \frac{1}{N^2} \sum_{n=0}^{N-1} \text{var}(x[n]) \\
 &= \frac{1}{N^2} N \sigma^2 \\
 &= \frac{\sigma^2}{N}
 \end{aligned}$$

since the $w[n]$'s are uncorrelated and thus

$$\begin{aligned}
 \text{var}(\tilde{A}) &= \text{var}(x[0]) \\
 &= \sigma^2 \\
 &> \text{var}(\hat{A}).
 \end{aligned}$$

Furthermore, if we could assume that $w[n]$ is Gaussian, we could also conclude that the probability of a given magnitude error is less for $\hat{A}$ than for $\check{A}$ (see Problem 2.7).

Several important points are illustrated by the previous example, which should always be kept in mind.

1. An estimator is a random variable. As such, its performance can only be completely described statistically or by its PDF.
2. The use of computer simulations for assessing estimation performance, although quite valuable for gaining insight and motivating conjectures, is never conclusive. At best, the true performance may be obtained to the desired degree of accuracy. At worst, for an insufficient number of experiments and/or errors in the simulation techniques employed, erroneous results may be obtained (see Appendix 7A for a further discussion of Monte Carlo computer techniques).

Another theme that we will repeatedly encounter is the tradeoff between performance and computational complexity. As in the previous example, even though $\hat{A}$ has better performance, it also requires more computation. We will see that optimal estimators can sometimes be difficult to implement, requiring a multidimensional optimization or integration. In these situations, alternative estimators that are suboptimal, but which can be implemented on a digital computer, may be preferred. For any particular application, the user must determine whether the loss in performance is offset by the reduced computational complexity of a suboptimal estimator.

1.4 Some Notes to the Reader

Our philosophy in presenting a theory of estimation is to provide the user with the main ideas necessary for determining optimal estimators. We have included results that we deem to be most useful in practice, omitting some important theoretical issues. The latter can be found in many books on statistical estimation theory which have been written from a more theoretical viewpoint [Cox and Hinkley 1974, Kendall and Stuart 1976–1979, Rao 1973, Zacks 1981]. As mentioned previously, our goal is to obtain an optimal estimator, and we resort to a suboptimal one if the former cannot be found or is not implementable. The sequence of chapters in this book follows this approach, so that optimal estimators are discussed first, followed by approximately optimal estimators, and finally suboptimal estimators. In Chapter 14 a “road map” for finding a good estimator is presented along with a summary of the various estimators and their properties. The reader may wish to read this chapter first to obtain an overview.

We have tried to maximize insight by including many examples and minimizing long mathematical expositions, although much of the tedious algebra and proofs have been included as appendices. The DC level in noise described earlier will serve as a standard example in introducing almost all the estimation approaches. It is hoped that in doing so the reader will be able to develop his or her own intuition by building upon previously assimilated concepts. Also, where possible, the scalar estimator is

presented first followed by the vector estimator. This approach reduces the tendency of vector/matrix algebra to obscure the main ideas. Finally, classical estimation is described first, followed by Bayesian estimation, again in the interest of not obscuring the main issues. The estimators obtained using the two approaches, although similar in appearance, are fundamentally different.

The mathematical notation for all common symbols is summarized in Appendix 2. The distinction between a continuous-time waveform and a discrete-time waveform or sequence is made through the symbolism $x(t)$ for continuous-time and $x[n]$ for discrete-time. Plots of $x[n]$, however, appear continuous in time, the points having been connected by straight lines for easier viewing. All vectors and matrices are boldface with all vectors being *column* vectors. All other symbolism is defined within the context of the discussion.

References

- Box, G.E.P., G.M. Jenkins, *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco, 1970.
- Burdic, W.S., *Underwater Acoustic System Analysis*, Prentice-Hall, Englewood Cliffs, N.J., 1984.
- Cox, D.R., D.V. Hinkley, *Theoretical Statistics*, Chapman and Hall, New York, 1974.
- Dabbous, T.E., N.U. Ahmed, J.C. McMillan, D.F. Liang, "Filtering of Discontinuous Processes Arising in Marine Integrated Navigation," *IEEE Trans. Aerosp. Electron. Syst.*, Vol. 24, pp. 85-100, 1988.
- Gauss, K.G., *Theory of Motion of Heavenly Bodies*, Dover, New York, 1963.
- Holm, S., J.M. Hovem, "Estimation of Scalar Ocean Wave Spectra by the Maximum Entropy Method," *IEEE J. Ocean Eng.*, Vol. 4, pp. 76-83, 1979.
- Huber, P.J., *Robust Statistics*, J. Wiley, New York, 1981.
- Jain, A.K., *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1989.
- Justice, J.H., "Array Processing in Exploration Seismology," in *Array Signal Processing*, S. Haykin, ed., Prentice-Hall, Englewood Cliffs, N.J., 1985.
- Kendall, Sir M., A. Stuart, *The Advanced Theory of Statistics*, Vols. 1-3, Macmillan, New York, 1976-1979.
- Knight, W.S., R.G. Pridham, S.M. Kay, "Digital Signal Processing for Sonar," *Proc. IEEE*, Vol. 69, pp. 1451-1506, Nov. 1981.
- Proakis, J.G., *Digital Communications*, McGraw-Hill, New York, 1983.
- Rabiner, L.R., R.W. Schafer, *Digital Processing of Speech Signals*, Prentice-Hall, Englewood Cliffs, N.J., 1978.
- Rao, C.R., *Linear Statistical Inference and Its Applications*, J. Wiley, New York, 1973.
- Schuster, A., "On the Investigation of Hidden Periodicities with Application to a Supposed 26 Day Period of Meteorological Phenomena," *Terrestrial Magnetism*, Vol. 3, pp. 13-41, March 1898.
- Skolnik, M.I., *Introduction to Radar Systems*, McGraw-Hill, New York, 1980.
- Taylor, S., *Modeling Financial Time Series*, J. Wiley, New York, 1986.
- Widrow, B., Stearns, S.D., *Adaptive Signal Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1985.
- Zacks, S., *Parametric Statistical Inference*, Pergamon, New York, 1981.

Problems

1. In a radar system an estimator of round trip delay τ_0 has the PDF $\hat{\tau}_0 \sim \mathcal{N}(\tau_0, \sigma_{\hat{\tau}_0}^2)$, where τ_0 is the true value. If the range is to be estimated, propose an estimator $\hat{R}$ and find its PDF. Next determine the standard deviation $\sigma_{\hat{\tau}_0}$ so that 99% of the time the range estimate will be within 100 m of the true value. Use $c = 3 \times 10^8$ m/s for the speed of electromagnetic propagation.
2. An unknown parameter θ influences the outcome of an experiment which is modeled by the random variable x . The PDF of x is

$$p(x; \theta) = \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2}(x - \theta)^2 \right].$$

A series of experiments is performed, and x is found to always be in the interval [97, 103]. As a result, the investigator concludes that θ must have been 100. Is this assertion correct?

3. Let $x = \theta + w$, where w is a random variable with PDF $p_w(w)$. If θ is a deterministic parameter, find the PDF of x in terms of p_w and denote it by $p(x; \theta)$. Next assume that θ is a random variable independent of w and find the conditional PDF $p(x|\theta)$. Finally, do not assume that θ and w are independent and determine $p(x|\theta)$. What can you say about $p(x; \theta)$ versus $p(x|\theta)$?
4. It is desired to estimate the value of a DC level A in WGN or

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is zero mean and uncorrelated, and each sample has variance $\sigma^2 = 1$. Consider the two estimators

$$\begin{aligned} \hat{A} &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] \\ \hat{A} &= \frac{1}{N+2} \left(2x[0] + \sum_{n=1}^{N-2} x[n] + 2x[N-1] \right). \end{aligned}$$

Which one is better? Does it depend on the value of A ?

5. For the same data set as in Problem 1.4 the following estimator is proposed:

$$\hat{A} = \begin{cases} x[0] & \frac{A^2}{\sigma^2} = A^2 > 1000 \\ \frac{1}{N} \sum_{n=0}^{N-1} x[n] & \frac{A^2}{\sigma^2} = A^2 \leq 1000. \end{cases}$$

The rationale for this estimator is that for a high enough signal-to-noise ratio (SNR) or A^2/σ^2 , we do not need to reduce the effect of noise by averaging and hence can avoid the added computation. Comment on this approach.

Chapter 2

Minimum Variance Unbiased Estimation

2.1 Introduction

In this chapter we will begin our search for good estimators of unknown deterministic parameters. We will restrict our attention to estimators which on the average yield the true parameter value. Then, within this class of estimators the goal will be to find the one that exhibits the least variability. Hopefully, the estimator thus obtained will produce values close to the true value most of the time. The notion of a minimum variance unbiased estimator is examined within this chapter, but the means to find it will require some more theory. Succeeding chapters will provide that theory as well as apply it to many of the typical problems encountered in signal processing.

2.2 Summary

An unbiased estimator is defined by (2.1), with the important proviso that this holds for all possible values of the unknown parameter. Within this class of estimators the one with the minimum variance is sought. The unbiased constraint is shown by example to be desirable from a practical viewpoint since the more natural error criterion, the minimum mean square error, defined in (2.5), generally leads to unrealizable estimators. Minimum variance unbiased estimators do not, in general, exist. When they do, several methods can be used to find them. The methods rely on the Cramer-Rao lower bound and the concept of a sufficient statistic. If a minimum variance unbiased estimator does not exist or if both of the previous two approaches fail, a further constraint on the estimator, to being linear in the data, leads to an easily implemented, but suboptimal, estimator.

2.3 Unbiased Estimators

For an estimator to be unbiased we mean that on the average the estimator will yield the true value of the unknown parameter. Since the parameter value may in general be anywhere in the interval $a < \theta < b$, unbiasedness asserts that no matter what the true value of θ , our estimator will yield it on the average. Mathematically, an estimator is unbiased if

$$E(\hat{\theta}) = \theta \quad a < \theta < b \quad (2.1)$$

where (a, b) denotes the range of possible values of θ .

Example 2.1 - Unbiased Estimator for DC Level in White Gaussian Noise

Consider the observations

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where A is the parameter to be estimated and $w[n]$ is WGN. The parameter A can take on any value in the interval $-\infty < A < \infty$. Then, a reasonable estimator for the average value of $x[n]$ is

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n] \quad (2.2)$$

or the sample mean. Due to the linearity properties of the expectation operator

$$\begin{aligned} E(\hat{A}) &= E\left[\frac{1}{N} \sum_{n=0}^{N-1} x[n]\right] \\ &= \frac{1}{N} \sum_{n=0}^{N-1} E(x[n]) \\ &= \frac{1}{N} \sum_{n=0}^{N-1} A \\ &= A \end{aligned}$$

for all A . The sample mean estimator is therefore unbiased. $\diamond$

In this example A can take on any value, although in general the values of an unknown parameter may be restricted by physical considerations. Estimating the resistance R of an unknown resistor, for example, would necessitate an interval $0 < R < \infty$.

Unbiased estimators tend to have symmetric PDFs centered about the true value of θ , although this is not necessary (see Problem 2.5). For Example 2.1 the PDF is shown in Figure 2.1 and is easily shown to be $\mathcal{N}(A, \sigma^2/N)$ (see Problem 2.3).

The restriction that $E(\hat{\theta}) = \theta$ for all θ is an important one. Letting $\hat{\theta} = g(\mathbf{x})$, where $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[N-1]]^T$, it asserts that

$$E(\hat{\theta}) = \int g(\mathbf{x}) p(\mathbf{x}; \theta) d\mathbf{x} = \theta \quad \text{for all } \theta. \quad (2.3)$$




Figure 2.1 Probability density function for sample mean estimator

It is possible, however, that (2.3) may hold for some values of θ and not others, as the next example illustrates.

Example 2.2 - Biased Estimator for DC Level in White Noise

Consider again Example 2.1 but with the *modified* sample mean estimator

$$\tilde{A} = \frac{1}{2N} \sum_{n=0}^{N-1} x[n].$$

Then,

$$\begin{aligned} E(\tilde{A}) &= \frac{1}{2} A \\ &= A \text{ if } A = 0 \\ &\neq A \text{ if } A \neq 0. \end{aligned}$$

It is seen that (2.3) holds for the modified estimator only for $A = 0$. Clearly, $\tilde{A}$ is a biased estimator. $\diamond$

That an estimator is unbiased does not necessarily mean that it is a good estimator. It only guarantees that on the average it will attain the true value. On the other hand, biased estimators are ones that are characterized by a systematic error, which presumably should not be present. A persistent bias will always result in a poor estimator. As an example, the unbiased property has an important implication when several estimators are combined (see Problem 2.4). It sometimes occurs that multiple estimates of the same parameter are available, i.e., $\{\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_n\}$. A reasonable procedure is to combine these estimates into, hopefully, a better one by averaging them to form

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n \hat{\theta}_i. \quad (2.4)$$

Assuming the estimators are unbiased, with the same variance, and uncorrelated with each other,

$$E(\hat{\theta}) = \theta$$




Figure 2.2 Effect of combining estimators

and

$$\begin{aligned}\text{var}(\hat{\theta}) &= \frac{1}{n^2} \sum_{i=1}^n \text{var}(\hat{\theta}_i) \\ &= \frac{\text{var}(\hat{\theta}_1)}{n}\end{aligned}$$

so that as more estimates are averaged, the variance will decrease. Ultimately, as $n \rightarrow \infty$, $\hat{\theta} \rightarrow \theta$. However, if the estimators are biased or $E(\hat{\theta}_i) = \theta + b(\theta)$, then

$$\begin{aligned}E(\hat{\theta}) &= \frac{1}{n} \sum_{i=1}^n E(\hat{\theta}_i) \\ &= \theta + b(\theta)\end{aligned}$$

and no matter how many estimators are averaged, $\hat{\theta}$ will not converge to the true value. This is depicted in Figure 2.2. Note that, in general,

$$b(\theta) = E(\hat{\theta}) - \theta$$

is defined as the *bias* of the estimator.

2.4 Minimum Variance Criterion

In searching for optimal estimators we need to adopt some optimality criterion. A natural one is the mean square error (MSE), defined as

$$\text{mse}(\hat{\theta}) = E \left[(\hat{\theta} - \theta)^2 \right]. \quad (2.5)$$

This measures the average mean squared deviation of the estimator from the true value. Unfortunately, adoption of this natural criterion leads to unrealizable estimators, ones that cannot be written solely as a function of the data. To understand the problem which arises we first rewrite the MSE as

$$\begin{aligned} \text{mse}(\hat{\theta}) &= E \left\{ \left[(\hat{\theta} - E(\hat{\theta})) + (E(\hat{\theta}) - \theta) \right]^2 \right\} \\ &= \text{var}(\hat{\theta}) + [E(\hat{\theta}) - \theta]^2 \\ &= \text{var}(\hat{\theta}) + b^2(\theta) \end{aligned} \quad (2.6)$$

which shows that the MSE is composed of errors due to the variance of the estimator as well as the bias. As an example, for the problem in Example 2.1 consider the modified estimator

$$\check{A} = a \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

for some constant a . We will attempt to find the a which results in the minimum MSE. Since $E(\check{A}) = aA$ and $\text{var}(\check{A}) = a^2\sigma^2/N$, we have, from (2.6),

$$\text{mse}(\check{A}) = \frac{a^2\sigma^2}{N} + (a-1)^2A^2.$$

Differentiating the MSE with respect to a yields

$$\frac{d \text{mse}(\check{A})}{da} = \frac{2a\sigma^2}{N} + 2(a-1)A^2$$

which upon setting to zero and solving yields the optimum value

$$a_{\text{opt}} = \frac{A^2}{A^2 + \sigma^2/N}.$$

It is seen that, unfortunately, the optimal value of a depends upon the unknown parameter A . The estimator is therefore not realizable. In retrospect the estimator depends upon A since the bias term in (2.6) is a function of A . It would seem that any criterion which depends on the bias will lead to an unrealizable estimator. Although this is generally true, on occasion realizable minimum MSE estimators can be found [Bibby and Toutenburg 1977, Rao 1973, Stoica and Moses 1990].




Figure 2.3 Possible dependence of estimator variance with θ

From a practical viewpoint the minimum MSE estimator needs to be abandoned. An alternative approach is to constrain the bias to be zero and find the estimator which minimizes the variance. Such an estimator is termed the **minimum variance unbiased (MVU) estimator**. Note that from (2.6) that the MSE of an unbiased estimator is just the variance.

Minimizing the variance of an unbiased estimator also has the effect of concentrating the PDF of the estimation error, $\hat{\theta} - \theta$, about zero (see Problem 2.7). The estimation error will therefore be less likely to be large.

2.5 Existence of the Minimum Variance Unbiased Estimator

The question arises as to whether a MVU estimator exists, i.e., an unbiased estimator with minimum variance for all θ . Two possible situations are described in Figure 2.3. If there are three unbiased estimators that exist and whose variances are shown in Figure 2.3a, then clearly $\hat{\theta}_3$ is the MVU estimator. If the situation in Figure 2.3b exists, however, then there is no MVU estimator since for $\theta < \theta_0$, $\hat{\theta}_2$ is better, while for $\theta > \theta_0$, $\hat{\theta}_3$ is better. In the former case $\hat{\theta}_3$ is sometimes referred to as the **uniformly minimum variance unbiased estimator** to emphasize that the variance is smallest for all θ . In general, the MVU estimator does not always exist, as the following example illustrates.

Example 2.3 - Counterexample to Existence of MVU Estimator

If the form of the PDF changes with θ , then it would be expected that the best estimator would also change with θ . Assume that we have two independent observations $x[0]$ and $x[1]$ with PDF

$$\begin{aligned} x[0] &\sim \mathcal{N}(\theta, 1) \\ x[1] &\sim \begin{cases} \mathcal{N}(\theta, 1) & \text{if } \theta \geq 0 \\ \mathcal{N}(\theta, 2) & \text{if } \theta < 0. \end{cases} \end{aligned}$$




Figure 2.4 Illustration of nonexistence of minimum variance unbiased estimator

The two estimators

$$\begin{aligned}\hat{\theta}_1 &= \frac{1}{2}(x[0] + x[1]) \\ \hat{\theta}_2 &= \frac{2}{3}x[0] + \frac{1}{3}x[1]\end{aligned}$$

can easily be shown to be unbiased. To compute the variances we have that

$$\begin{aligned}\text{var}(\hat{\theta}_1) &= \frac{1}{4}(\text{var}(x[0]) + \text{var}(x[1])) \\ \text{var}(\hat{\theta}_2) &= \frac{4}{9}\text{var}(x[0]) + \frac{1}{9}\text{var}(x[1])\end{aligned}$$

so that

$$\text{var}(\hat{\theta}_1) = \begin{cases} \frac{18}{36} & \text{if } \theta \geq 0 \\ \frac{27}{36} & \text{if } \theta < 0 \end{cases}$$

and

$$\text{var}(\hat{\theta}_2) = \begin{cases} \frac{20}{36} & \text{if } \theta \geq 0 \\ \frac{24}{36} & \text{if } \theta < 0. \end{cases}$$

The variances are shown in Figure 2.4. Clearly, between these two estimators no MVU estimator exists. It is shown in Problem 3.6 that for $\theta \geq 0$ the minimum possible variance of an unbiased estimator is $18/36$, while that for $\theta < 0$ is $24/36$. Hence, no single estimator can have a variance uniformly less than or equal to the minima shown in Figure 2.4. $\diamond$

To conclude our discussion of existence we should note that it is also possible that there may not exist even a single unbiased estimator (see Problem 2.11). In this case any search for a MVU estimator is fruitless.

2.6 Finding the Minimum Variance Unbiased Estimator

Even if a MVU estimator exists, we may not be able to find it. There is no known “turn-the-crank” procedure which will always produce the estimator. In the next few chapters we shall discuss several possible approaches. They are:




Figure 2.5 Cramer-Rao lower bound on variance of unbiased estimator

1. Determine the Cramer-Rao lower bound (CRLB) and check to see if some estimator satisfies it (Chapters 3 and 4).
2. Apply the Rao-Blackwell-Lehmann-Scheffe (RBLs) theorem (Chapter 5).
3. Further restrict the class of estimators to be not only unbiased but also *linear*. Then, find the minimum variance estimator within this restricted class (Chapter 6).

Approaches 1 and 2 may produce the MVU estimator, while 3 will yield it only if the MVU estimator is linear in the data.

The CRLB allows us to determine that for *any unbiased estimator* the variance must be greater than or equal to a given value, as shown in Figure 2.5. If an estimator exists whose variance equals the CRLB for each value of θ , then it must be the MVU estimator. In this case, the theory of the CRLB immediately yields the estimator. It may happen that no estimator exists whose variance equals the bound. Yet, a MVU estimator may still exist, as for instance in the case of θ_1 in Figure 2.5. Then, we must resort to the Rao-Blackwell-Lehmann-Scheffe theorem. This procedure first finds a *sufficient statistic*, one which uses all the data efficiently, and then finds a function of the sufficient statistic which is an unbiased estimator of θ . With a slight restriction of the PDF of the data this procedure will then be guaranteed to produce the MVU estimator. The third approach requires the estimator to be linear, a sometimes severe restriction, and chooses the best *linear* estimator. Of course, only for particular data sets can this approach produce the MVU estimator.

2.7 Extension to a Vector Parameter

If $\theta = [\theta_1 \theta_2 \dots \theta_p]^T$ is a vector of unknown parameters, then we say that an estimator $\hat{\theta} = [\hat{\theta}_1 \hat{\theta}_2 \dots \hat{\theta}_p]^T$ is unbiased if

$$E(\hat{\theta}_i) = \theta_i \quad a_i < \theta_i < b_i \quad (2.7)$$

for $i = 1, 2, \dots, p$. By defining

$$E(\hat{\theta}) = \begin{bmatrix} E(\hat{\theta}_1) \\ E(\hat{\theta}_2) \\ \vdots \\ E(\hat{\theta}_p) \end{bmatrix}$$

we can equivalently define an unbiased estimator to have the property

$$E(\hat{\theta}) = \theta$$

for every θ contained within the space defined in (2.7). A MVU estimator has the additional property that $\text{var}(\hat{\theta}_i)$ for $i = 1, 2, \dots, p$ is minimum among all unbiased estimators.

References

- Bibby, J., H. Toutenburg, *Prediction and Improved Estimation in Linear Models*, J. Wiley, New York, 1977.
- Rao, C.R., *Linear Statistical Inference and Its Applications*, J. Wiley, New York, 1973.
- Stoica, P., R. Moses, "On Biased Estimators and the Unbiased Cramer-Rao Lower Bound," *Signal Process.*, Vol. 21, pp. 349–350, 1990.

Problems

- 2.1** The data $\{x[0], x[1], \dots, x[N-1]\}$ are observed where the $x[n]$'s are independent and identically distributed (IID) as $\mathcal{N}(0, \sigma^2)$. We wish to estimate the variance σ^2 as

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} x^2[n].$$

Is this an unbiased estimator? Find the variance of $\hat{\sigma}^2$ and examine what happens as $N \rightarrow \infty$.

- 2.2** Consider the data $\{x[0], x[1], \dots, x[N-1]\}$, where each sample is distributed as $\mathcal{U}[0, \theta]$ and the samples are IID. Can you find an unbiased estimator for θ ? The range of θ is $0 < \theta < \infty$.

- 2.3** Prove that the PDF of $\hat{A}$ given in Example 2.1 is $\mathcal{N}(A, \sigma^2/N)$.

- 2.4** The heart rate h of a patient is automatically recorded by a computer every 100 ms. In 1 s the measurements $\{\hat{h}_1, \hat{h}_2, \dots, \hat{h}_{10}\}$ are averaged to obtain $\hat{h}$. If $E(\hat{h}_i) = \alpha h$ for some constant α and $\text{var}(\hat{h}_i) = 1$ for each i , determine whether averaging improves the estimator if $\alpha = 1$ and $\alpha = 1/2$. Assume each measurement is uncorrelated.

- 2.5** Two samples $\{x[0], x[1]\}$ are independently observed from a $\mathcal{N}(0, \sigma^2)$ distribution. The estimator

$$\hat{\sigma}^2 = \frac{1}{2}(x^2[0] + x^2[1])$$

is unbiased. Find the PDF of $\hat{\sigma}^2$ to determine if it is symmetric about σ^2 .

2.6 For the problem described in Example 2.1 the more general estimator

$$\hat{A} = \sum_{n=0}^{N-1} a_n x[n]$$

is proposed. Find the a_n 's so that the estimator is unbiased and the variance is minimized. Hint: Use Lagrangian multipliers with unbiasedness as the constraint equation.

2.7 Two unbiased estimators are proposed whose variances satisfy $\text{var}(\hat{\theta}) < \text{var}(\tilde{\theta})$. If both estimators are Gaussian, prove that

$$\Pr\{|\hat{\theta} - \theta| > \epsilon\} < \Pr\{|\tilde{\theta} - \theta| > \epsilon\}$$

for any $\epsilon > 0$. This says that the estimator with less variance is to be preferred since its PDF is more concentrated about the true value.

2.8 For the problem described in Example 2.1 show that as $N \rightarrow \infty$, $\hat{A} \rightarrow A$ by using the results of Problem 2.3. To do so prove that

$$\lim_{N \rightarrow \infty} \Pr\{|\hat{A} - A| > \epsilon\} = 0$$

for any $\epsilon > 0$. In this case the estimator $\hat{A}$ is said to be *consistent*. Investigate what happens if the alternative estimator $\tilde{A} = \frac{1}{2N} \sum_{n=0}^{N-1} x[n]$ is used instead.

2.9 This problem illustrates what happens to an unbiased estimator when it undergoes a nonlinear transformation. In Example 2.1, if we choose to estimate the unknown parameter $\theta = A^2$ by

$$\hat{\theta} = \left(\frac{1}{N} \sum_{n=0}^{N-1} x[n] \right)^2,$$

can we say that the estimator is unbiased? What happens as $N \rightarrow \infty$?

2.10 In Example 2.1 assume now that in addition to A , the value of σ^2 is also unknown. We wish to estimate the vector parameter

$$\hat{\theta} = \begin{bmatrix} A \\ \sigma^2 \end{bmatrix}.$$

Is the estimator

$$\hat{\theta} = \begin{bmatrix} \hat{A} \\ \hat{\sigma}^2 \end{bmatrix} = \begin{bmatrix} \frac{1}{N} \sum_{n=0}^{N-1} x[n] \\ \frac{1}{N-1} \sum_{n=0}^{N-1} (x[n] - \hat{A})^2 \end{bmatrix}$$

unbiased?

- 2.11** Given a single observation $x[0]$ from the distribution $\mathcal{U}[0, 1/\theta]$, it is desired to estimate θ . It is assumed that $\theta > 0$. Show that for an estimator $\hat{\theta} = g(x[0])$ to be unbiased we must have

$$\int_0^{\frac{1}{\theta}} g(u) du = 1.$$

Next prove that a function g cannot be found to satisfy this condition for all $\theta > 0$.

Chapter 3

Cramer-Rao Lower Bound

3.1 Introduction

Being able to place a lower bound on the variance of any unbiased estimator proves to be extremely useful in practice. At best, it allows us to assert that an estimator is the MVU estimator. This will be the case if the estimator attains the bound for all values of the unknown parameter. At worst, it provides a benchmark against which we can compare the performance of any unbiased estimator. Furthermore, it alerts us to the physical impossibility of finding an unbiased estimator whose variance is less than the bound. The latter is often useful in signal processing feasibility studies. Although many such variance bounds exist [McAulay and Hofstetter 1971, Kendall and Stuart 1979, Seidman 1970, Ziv and Zakai 1969], the Cramer-Rao lower bound (CRLB) is by far the easiest to determine. Also, the theory allows us to immediately determine if an estimator exists that attains the bound. If no such estimator exists, then all is not lost since estimators can be found that attain the bound in an approximate sense, as described in Chapter 7. For these reasons we restrict our discussion to the CRLB.

3.2 Summary

The CRLB for a scalar parameter is given by (3.6). If the condition (3.7) is satisfied, then the bound will be attained and the estimator that attains it is readily found. An alternative means of determining the CRLB is given by (3.12). For a signal with an unknown parameter in WGN, (3.14) provides a convenient means to evaluate the bound. When a function of a parameter is to be estimated, the CRLB is given by (3.16). Even though an efficient estimator may exist for θ , in general there will not be one for a function of θ (unless the function is linear). For a vector parameter the CRLB is determined using (3.20) and (3.21). As in the scalar parameter case, if condition (3.25) holds, then the bound is attained and the estimator that attains the bound is easily found. For a function of a vector parameter (3.30) provides the bound. A general formula for the Fisher information matrix (used to determine the vector CRLB) for a multivariate Gaussian PDF is given by (3.31). Finally, if the data set comes from a




Figure 3.1 PDF dependence on unknown parameter

WSS Gaussian random process, then an approximate CRLB, that depends on the PSD, is given by (3.34). It is valid asymptotically or as the data record length becomes large.

3.3 Estimator Accuracy Considerations

Before stating the CRLB theorem, it is worthwhile to expose the hidden factors that determine how well we can estimate a parameter. Since all our information is embodied in the observed data and the underlying PDF for that data, it is not surprising that the estimation accuracy depends directly on the PDF. For instance, we should not expect to be able to estimate a parameter with any degree of accuracy if the PDF depends only weakly upon that parameter, or in the extreme case, if the PDF does not depend on it at all. In general, the more the PDF is influenced by the unknown parameter, the better we should be able to estimate it.

Example 3.1 - PDF Dependence on Unknown Parameter

If a single sample is observed as

$$x[0] = A + w[0]$$

where $w[0] \sim \mathcal{N}(0, \sigma^2)$, and it is desired to estimate A , then we expect a better estimate if σ^2 is small. Indeed, a good unbiased estimator is $\hat{A} = x[0]$. The variance is, of course, just σ^2 , so that the estimator accuracy improves as σ^2 decreases. An alternative way of viewing this is shown in Figure 3.1, where the PDFs for two different variances are shown. They are

$$p_i(x[0]; A) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left[-\frac{1}{2\sigma_i^2} (x[0] - A)^2 \right] \quad (3.1)$$

for $i = 1, 2$. The PDF has been plotted versus the *unknown parameter* A for a *given value of* $x[0]$. If $\sigma_1^2 < \sigma_2^2$, then we should be able to estimate A more accurately based on $p_1(x[0]; A)$. We may interpret this result by referring to Figure 3.1. If $x[0] = 3$ and $\sigma_1 = 1/3$, then as shown in Figure 3.1a, the values of $A > 4$ are highly unlikely. To see

this we determine the probability of observing $x[0]$ in the interval $[x[0] - \delta/2, x[0] + \delta/2] = [3 - \delta/2, 3 + \delta/2]$ when A takes on a given value or

$$\Pr \left\{ 3 - \frac{\delta}{2} \leq x[0] \leq 3 + \frac{\delta}{2} \right\} = \int_{3-\frac{\delta}{2}}^{3+\frac{\delta}{2}} p_i(u; A) du$$

which for δ small is $p_i(x[0] = 3; A)\delta$. But $p_1(x[0] = 3; A = 4)\delta = 0.01\delta$, while $p_1(x[0] = 3; A = 3)\delta = 1.20\delta$. The probability of observing $x[0]$ in a small interval centered about $x[0] = 3$ when $A = 4$ is small with respect to that when $A = 3$. Hence, the values of $A > 4$ can be eliminated from consideration. It might be argued that values of A in the interval $3 \pm 3\sigma_1 = [2, 4]$ are viable candidates. For the PDF in Figure 3.1b there is a much weaker dependence on A . Here our viable candidates are in the much wider interval $3 \pm 3\sigma_2 = [0, 6]$. $\diamond$

When the PDF is viewed as a function of the unknown parameter (with $\mathbf{x}$ fixed), it is termed the likelihood function. Two examples of likelihood functions were shown in Figure 3.1. Intuitively, the “sharpness” of the likelihood function determines how accurately we can estimate the unknown parameter. To quantify this notion observe that the sharpness is effectively measured by the negative of the second derivative of the logarithm of the likelihood function at its peak. This is the curvature of the log-likelihood function. In Example 3.1, if we consider the natural logarithm of the PDF

$$\ln p(x[0]; A) = -\ln \sqrt{2\pi\sigma^2} - \frac{1}{2\sigma^2} (x[0] - A)^2$$

then the first derivative is

$$\frac{\partial \ln p(x[0]; A)}{\partial A} = \frac{1}{\sigma^2} (x[0] - A) \quad (3.2)$$

and the negative of the second derivative becomes

$$-\frac{\partial^2 \ln p(x[0]; A)}{\partial A^2} = \frac{1}{\sigma^2}. \quad (3.3)$$

The curvature increases as σ^2 decreases. Since we already know that the estimator $\hat{A} = x[0]$ has variance σ^2 , then for this example

$$\text{var}(\hat{A}) = \frac{1}{-\frac{\partial^2 \ln p(x[0]; A)}{\partial A^2}} \quad (3.4)$$

and the variance decreases as the curvature increases. Although in this example the second derivative does not depend on $x[0]$, in general it will. Thus, a more appropriate measure of curvature is

$$-E \left[\frac{\partial^2 \ln p(x[0]; A)}{\partial A^2} \right] \quad (3.5)$$

which measures the *average* curvature of the log-likelihood function. The expectation is taken with respect to $p(x[0]; A)$, resulting in a function of A only. The expectation acknowledges the fact that the likelihood function, which depends on $x[0]$, is itself a random variable. The larger the quantity in (3.5), the smaller the variance of the estimator.

3.4 Cramer-Rao Lower Bound

We are now ready to state the CRLB theorem.

Theorem 3.1 (Cramer-Rao Lower Bound - Scalar Parameter) *It is assumed that the PDF $p(\mathbf{x}; \theta)$ satisfies the “regularity” condition*

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] = 0 \quad \text{for all } \theta$$

where the expectation is taken with respect to $p(\mathbf{x}; \theta)$. Then, the variance of any unbiased estimator $\hat{\theta}$ must satisfy

$$\text{var}(\hat{\theta}) \geq \frac{1}{-E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]} \quad (3.6)$$

where the derivative is evaluated at the true value of θ and the expectation is taken with respect to $p(\mathbf{x}; \theta)$. Furthermore, an unbiased estimator may be found that attains the bound for all θ if and only if

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = I(\theta)(g(\mathbf{x}) - \theta) \quad (3.7)$$

for some functions g and I . That estimator, which is the MVU estimator, is $\hat{\theta} = g(\mathbf{x})$, and the minimum variance is $1/I(\theta)$.

The expectation in (3.6) is explicitly given by

$$E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right] = \int \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} p(\mathbf{x}; \theta) d\mathbf{x}$$

since the second derivative is a random variable dependent on $\mathbf{x}$. Also, the bound will depend on θ in general, so that it is displayed as in Figure 2.5 (dashed curve). An example of a PDF that does not satisfy the regularity condition is given in Problem 3.1. For a proof of the theorem see Appendix 3A.

Some examples are now given to illustrate the evaluation of the CRLB.

Example 3.2 - CRLB for Example 3.1

For Example 3.1 we see that from (3.3) and (3.6)

$$\text{var}(\hat{A}) \geq \sigma^2 \quad \text{for all } A.$$

Thus, no unbiased estimator can exist whose variance is lower than σ^2 for even a single value of A . But in fact we know that if $\hat{A} = x[0]$, then $\text{var}(\hat{A}) = \sigma^2$ for all A . Since $x[0]$ is unbiased and attains the CRLB, it must therefore be the MVU estimator. Had we been unable to guess that $x[0]$ would be a good estimator, we could have used (3.7). From (3.2) and (3.7) we make the identification

$$\begin{aligned}\theta &= A \\ I(\theta) &= \frac{1}{\sigma^2} \\ g(x[0]) &= x[0]\end{aligned}$$

so that (3.7) is satisfied. Hence, $\hat{A} = g(x[0]) = x[0]$ is the MVU estimator. Also, note that $\text{var}(\hat{A}) = \sigma^2 = 1/I(\theta)$, so that according to (3.6) we must have

$$I(\theta) = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right].$$

We will return to this after the next example. See also Problem 3.2 for a generalization to the non-Gaussian case. $\diamond$

Example 3.3 - DC Level in White Gaussian Noise

Generalizing Example 3.1, consider the multiple observations

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 . To determine the CRLB for A

$$\begin{aligned}p(\mathbf{x}; A) &= \prod_{n=0}^{N-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} (x[n] - A)^2 \right] \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right].\end{aligned}$$

Taking the first derivative

$$\begin{aligned}\frac{\partial \ln p(\mathbf{x}; A)}{\partial A} &= \frac{\partial}{\partial A} \left[-\ln[(2\pi\sigma^2)^{\frac{N}{2}}] - \frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right] \\ &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A) \\ &= \frac{N}{\sigma^2} (\bar{x} - A)\end{aligned} \tag{3.8}$$

Example 3.4 - Phase Estimation

Assume that we wish to estimate the phase ϕ of a sinusoid embedded in WGN or

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n] \quad n = 0, 1, \dots, N-1.$$

The amplitude A and frequency f_0 are assumed known (see Example 3.14 for the case when they are unknown). The PDF is

$$p(\mathbf{x}; \phi) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} [x[n] - A \cos(2\pi f_0 n + \phi)]^2 \right\}.$$

Differentiating the log-likelihood function produces

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \phi)}{\partial \phi} &= -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} [x[n] - A \cos(2\pi f_0 n + \phi)] A \sin(2\pi f_0 n + \phi) \\ &= -\frac{A}{\sigma^2} \sum_{n=0}^{N-1} [x[n] \sin(2\pi f_0 n + \phi) - \frac{A}{2} \sin(4\pi f_0 n + 2\phi)] \end{aligned}$$

and

$$\frac{\partial^2 \ln p(\mathbf{x}; \phi)}{\partial \phi^2} = -\frac{A}{\sigma^2} \sum_{n=0}^{N-1} [x[n] \cos(2\pi f_0 n + \phi) - A \cos(4\pi f_0 n + 2\phi)].$$

Upon taking the negative expected value we have

$$\begin{aligned} -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \phi)}{\partial \phi^2} \right] &= \frac{A}{\sigma^2} \sum_{n=0}^{N-1} [A \cos^2(2\pi f_0 n + \phi) - A \cos(4\pi f_0 n + 2\phi)] \\ &= \frac{A^2}{\sigma^2} \sum_{n=0}^{N-1} \left[\frac{1}{2} + \frac{1}{2} \cos(4\pi f_0 n + 2\phi) - \cos(4\pi f_0 n + 2\phi) \right] \\ &\approx \frac{NA^2}{2\sigma^2} \end{aligned}$$

since

$$\frac{1}{N} \sum_{n=0}^{N-1} \cos(4\pi f_0 n + 2\phi) \approx 0$$

for f_0 not near 0 or $1/2$ (see Problem 3.7). Therefore,

$$\text{var}(\hat{\phi}) \geq \frac{2\sigma^2}{NA^2}.$$

In this example the condition for the bound to hold is not satisfied. Hence, a phase estimator does not exist which is unbiased and attains the CRLB. It is still possible, however, that an MVU estimator may exist. At this point we do not know how to




Figure 3.2 Efficiency vs. minimum variance

determine whether an MVU estimator exists, and if it does, how to find it. The theory of sufficient statistics presented in Chapter 5 will allow us to answer these questions. $\diamond$

An estimator which is unbiased and attains the CRLB, as the sample mean estimator in Example 3.3 does, is said to be efficient in that it efficiently uses the data. An MVU estimator may or may not be efficient. For instance, in Figure 3.2 the variances of all possible estimators (for purposes of illustration there are three unbiased estimators) are displayed. In Figure 3.2a, $\hat{\theta}_1$ is efficient in that it attains the CRLB. Therefore, it is also the MVU estimator. On the other hand, in Figure 3.2b, $\hat{\theta}_1$ does not attain the CRLB, and hence it is not efficient. However, since its variance is uniformly less than that of all other unbiased estimators, it is the MVU estimator.

The CRLB given by (3.6) may also be expressed in a slightly different form. Although (3.6) is usually more convenient for evaluation, the alternative form is sometimes useful for theoretical work. It follows from the identity (see Appendix 3A)

$$E \left[\left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 \right] = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right] \quad (3.11)$$

so that

$$\text{var}(\hat{\theta}) \geq \frac{1}{E \left[\left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 \right]} \quad (3.12)$$

(see Problem 3.8).

The denominator in (3.6) is referred to as the *Fisher information* $I(\theta)$ for the data $\mathbf{x}$ or

$$I(\theta) = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]. \quad (3.13)$$

As we saw previously, when the CRLB is attained, the variance is the reciprocal of the Fisher information. Intuitively, the more information, the lower the bound. It has the essential properties of an information measure in that it is

1. nonnegative due to (3.11)
2. additive for independent observations.

The latter property leads to the result that the CRLB for N IID observations is $1/N$ times that for one observation. To verify this, note that for independent observations

$$\ln p(\mathbf{x}; \theta) = \sum_{n=0}^{N-1} \ln p(x[n]; \theta).$$

This results in

$$-E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right] = - \sum_{n=0}^{N-1} E \left[\frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right]$$

and finally for identically distributed observations

$$I(\theta) = Ni(\theta)$$

where

$$i(\theta) = -E \left[\frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right]$$

is the Fisher information for one sample. For nonindependent samples we might expect that the information will be less than $Ni(\theta)$, as Problem 3.9 illustrates. For completely dependent samples, as for example, $x[0] = x[1] = \dots = x[N-1]$, we will have $I(\theta) = i(\theta)$ (see also Problem 3.9). Therefore, additional observations carry no information, and the CRLB will not decrease with increasing data record length.

3.5 General CRLB for Signals in White Gaussian Noise

Since it is common to assume white Gaussian noise, it is worthwhile to derive the CRLB for this case. Later, we will extend this to nonwhite Gaussian noise and a vector parameter as given by (3.31). Assume that a deterministic signal with an unknown parameter θ is observed in WGN as

$$x[n] = s[n; \theta] + w[n] \quad n = 0, 1, \dots, N-1.$$

The dependence of the signal on θ is explicitly noted. The likelihood function is

$$p(\mathbf{x}; \theta) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - s[n; \theta])^2 \right\}.$$

Differentiating once produces

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - s[n; \theta]) \frac{\partial s[n; \theta]}{\partial \theta}$$

and a second differentiation results in

$$\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left\{ (x[n] - s[n; \theta]) \frac{\partial^2 s[n; \theta]}{\partial \theta^2} - \left(\frac{\partial s[n; \theta]}{\partial \theta} \right)^2 \right\}.$$

Taking the expected value yields

$$E \left(\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right) = -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left(\frac{\partial s[n; \theta]}{\partial \theta} \right)^2$$

so that finally

$$\text{var}(\hat{\theta}) \geq \frac{\sigma^2}{\sum_{n=0}^{N-1} \left(\frac{\partial s[n; \theta]}{\partial \theta} \right)^2}. \quad (3.14)$$

The form of the bound demonstrates the importance of the signal dependence on θ . Signals that change rapidly as the unknown parameter changes result in accurate estimators. A simple application of (3.14) to Example 3.3, in which $s[n; \theta] = \theta$, produces a CRLB of σ^2/N . The reader should also verify the results of Example 3.4. As a final example we examine the problem of frequency estimation.

Example 3.5 - Sinusoidal Frequency Estimation

We assume that the signal is sinusoidal and is represented as

$$s[n; f_0] = A \cos(2\pi f_0 n + \phi) \quad 0 < f_0 < \frac{1}{2}$$

where the amplitude and phase are known (see Example 3.14 for the case when they are unknown). From (3.14) the CRLB becomes

$$\text{var}(\hat{f}_0) \geq \frac{\sigma^2}{A^2 \sum_{n=0}^{N-1} [2\pi n \sin(2\pi f_0 n + \phi)]^2}. \quad (3.15)$$

The CRLB is plotted in Figure 3.3 versus frequency for an SNR of $A^2/\sigma^2 = 1$, a data record length of $N = 10$, and a phase of $\phi = 0$. It is interesting to note that there appear to be preferred frequencies (see also Example 3.14) for an approximation to (3.15). Also, as $f_0 \rightarrow 0$, the CRLB goes to infinity. This is because for f_0 close to zero a slight change in frequency will not alter the signal significantly. $\diamond$




Figure 3.3 Cramer-Rao lower bound for sinusoidal frequency estimation

3.6 Transformation of Parameters

It frequently occurs in practice that the parameter we wish to estimate is a function of some more fundamental parameter. For instance, in Example 3.3 we may not be interested in the sign of A but instead may wish to estimate A^2 or the power of the signal. Knowing the CRLB for A , we can easily obtain it for A^2 or in general for any function of A . As shown in Appendix 3A, if it is desired to estimate $\alpha = g(\theta)$, then the CRLB is

$$\text{var}(\hat{\alpha}) \geq \frac{\left(\frac{\partial g}{\partial \theta}\right)^2}{-E\left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2}\right]} \quad (3.16)$$

For the present example this becomes $\alpha = g(A) = A^2$ and

$$\boxed{\text{var}(\hat{A}^2) \geq \frac{(2A)^2}{N/\sigma^2} = \frac{4A^2\sigma^2}{N}} \quad (3.17)$$

Note that in using (3.16) the CRLB is expressed in terms of θ .

We saw in Example 3.3 that the sample mean estimator was efficient for A . It might be supposed that $\bar{x}^2$ is efficient for A^2 . To quickly dispel this notion we first show that $\bar{x}^2$ is not even an unbiased estimator. Since $\bar{x} \sim \mathcal{N}(A, \sigma^2/N)$

$$\begin{aligned} E(\bar{x}^2) &= E^2(\bar{x}) + \text{var}(\bar{x}) = A^2 + \frac{\sigma^2}{N} \\ &\neq A^2 \end{aligned} \quad (3.18)$$

Hence, we immediately conclude that the efficiency of an estimator is destroyed by a nonlinear transformation. That it is maintained for linear (actually affine) transformations is easily verified. Assume that an efficient estimator for θ exists and is given

by $\hat{\theta}$. It is desired to estimate $g(\theta) = a\theta + b$. As our estimator of $g(\theta)$, we choose $\widehat{g(\theta)} = g(\hat{\theta}) = a\hat{\theta} + b$. Then,

$$\begin{aligned} E(a\hat{\theta} + b) &= aE(\hat{\theta}) + b = a\theta + b \\ &= g(\theta), \end{aligned}$$

so that $\widehat{g(\theta)}$ is unbiased. The CRLB for $g(\theta)$, is from (3.16),

$$\begin{aligned} \text{var}(\widehat{g(\theta)}) &\geq \frac{\left(\frac{\partial g}{\partial \theta}\right)^2}{I(\theta)} \\ &= \left(\frac{\partial g(\theta)}{\partial \theta}\right)^2 \text{var}(\hat{\theta}) \\ &= a^2 \text{var}(\hat{\theta}). \end{aligned}$$

But $\text{var}(\widehat{g(\theta)}) = \text{var}(a\hat{\theta} + b) = a^2 \text{var}(\hat{\theta})$, so that the CRLB is achieved.

Although efficiency is preserved only over linear transformations, it is *approximately maintained over nonlinear transformations if the data record is large enough*. This has great practical significance in that we are frequently interested in estimating functions of parameters. To see why this property holds, we return to the previous example of estimating A^2 by $\bar{x}^2$. Although $\bar{x}^2$ is biased, we note from (3.18) that $\bar{x}^2$ is *asymptotically unbiased* or unbiased as $N \rightarrow \infty$. Furthermore, since $\bar{x} \sim \mathcal{N}(A, \sigma^2/N)$, we can evaluate the variance

$$\text{var}(\bar{x}^2) = E(\bar{x}^4) - E^2(\bar{x}^2)$$

by using the result that if $\xi \sim \mathcal{N}(\mu, \sigma^2)$, then

$$\begin{aligned} E(\xi^2) &= \mu^2 + \sigma^2, \\ E(\xi^4) &= \mu^4 + 6\mu^2\sigma^2 + 3\sigma^4, \end{aligned}$$

and therefore

$$\begin{aligned} \text{var}(\xi^2) &= E(\xi^4) - E^2(\xi^2) \\ &= 4\mu^2\sigma^2 + 2\sigma^4. \end{aligned}$$

For our problem we have then

$$\text{var}(\bar{x}^2) = \frac{4A^2\sigma^2}{N} + \frac{2\sigma^4}{N^2} \quad (3.19)$$

Hence, as $N \rightarrow \infty$, the variance approaches $4A^2\sigma^2/N$, the last term in (3.19) converging to zero faster than the first. But this is just the CRLB as given by (3.17). Our assertion that $\bar{x}^2$ is an *asymptotically efficient estimator of A^2* is verified. Intuitively, this situation occurs due to the *statistical linearity* of the transformation, as illustrated in Figure 3.4. As N increases, the PDF of $\bar{x}$ becomes more concentrated about the mean A . Therefore,




Figure 3.4 Statistical linearity of nonlinear transformations

the values of $\bar{x}$ that are observed lie in a small interval about $\bar{x} = A$ (the ± 3 standard deviation interval is displayed). Over this small interval the nonlinear transformation is approximately linear. Therefore, the transformation may be replaced by a linear one since a value of $\bar{x}$ in the *nonlinear region* rarely occurs. In fact, if we linearize g about A , we have the approximation

$$g(\bar{x}) \approx g(A) + \frac{dg(A)}{dA}(\bar{x} - A).$$

It follows that, to within this approximation,

$$E[g(\bar{x})] = g(A) = A^2$$

or the estimator is unbiased (asymptotically). Also,

$$\begin{aligned} \text{var}[g(\bar{x})] &= \left[\frac{dg(A)}{dA} \right]^2 \text{var}(\bar{x}) \\ &= \frac{(2A)^2 \sigma^2}{N} \\ &= \frac{4A^2 \sigma^2}{N} \end{aligned}$$

so that the estimator achieves the CRLB (asymptotically). Therefore, it is asymptotically efficient. This result also yields insight into the form of the CRLB given by (3.16).

3.7 Extension to a Vector Parameter

We now extend the results of the previous sections to the case where we wish to estimate a vector parameter $\theta = [\theta_1 \theta_2 \dots \theta_p]^T$. We will assume that the estimator $\hat{\theta}$ is unbiased

as defined in Section 2.7. The vector parameter CRLB will allow us to place a bound on the variance of each element. As derived in Appendix 3B, the CRLB is found as the $[i, i]$ element of the inverse of a matrix or

$$\text{var}(\hat{\theta}_i) \geq [\mathbf{I}^{-1}(\boldsymbol{\theta})]_{ii} \quad (3.20)$$

where $\mathbf{I}(\boldsymbol{\theta})$ is the $p \times p$ Fisher information matrix. The latter is defined by

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right] \quad (3.21)$$

for $i = 1, 2, \dots, p; j = 1, 2, \dots, p$. In evaluating (3.21) the true value of $\boldsymbol{\theta}$ is used. Note that in the scalar case ($p = 1$), $\mathbf{I}(\boldsymbol{\theta}) = I(\theta)$ and we have the scalar CRLB. Some examples follow.

Example 3.6 - DC Level in White Gaussian Noise (Revisited)

We now extend Example 3.3 to the case where in addition to A the noise variance σ^2 is also unknown. The parameter vector is $\boldsymbol{\theta} = [A \ \sigma^2]^T$, and hence $p = 2$. The 2×2 Fisher information matrix is

$$\mathbf{I}(\boldsymbol{\theta}) = \begin{bmatrix} -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A^2} \right] & -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A \partial \sigma^2} \right] \\ -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \sigma^2 \partial A} \right] & -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \sigma^2^2} \right] \end{bmatrix}.$$

It is clear from (3.21) that the matrix is symmetric since the order of partial differentiation may be interchanged and can also be shown to be positive definite (see Problem 3.10). The log-likelihood function is, from Example 3.3,

$$\ln p(\mathbf{x}; \boldsymbol{\theta}) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \ln \sigma^2 - \frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2.$$

The derivatives are easily found as

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A) \\ \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \sigma^2} &= -\frac{N}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{n=0}^{N-1} (x[n] - A)^2 \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A^2} &= -\frac{N}{\sigma^2} \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A \partial \sigma^2} &= -\frac{1}{\sigma^4} \sum_{n=0}^{N-1} (x[n] - A) \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \sigma^2^2} &= \frac{N}{2\sigma^4} - \frac{1}{\sigma^6} \sum_{n=0}^{N-1} (x[n] - A)^2. \end{aligned}$$

Upon taking the negative expectations, the Fisher information matrix becomes

$$\mathbf{I}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{N}{\sigma^2} & 0 \\ 0 & \frac{N}{2\sigma^4} \end{bmatrix}.$$

Although not true in general, for this example the Fisher information matrix is diagonal and hence easily inverted to yield

$$\begin{aligned} \text{var}(\hat{A}) &\geq \frac{\sigma^2}{N} \\ \text{var}(\hat{\sigma}^2) &\geq \frac{2\sigma^4}{N}. \end{aligned}$$

Note that the CRLB for $\hat{A}$ is the same as for the case when σ^2 is known due to the diagonal nature of the matrix. Again this is not true in general, as the next example illustrates. $\diamond$

Example 3.7 - Line Fitting

Consider the problem of line fitting or given the observations

$$x[n] = A + Bn + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN, determine the CRLB for the slope B and the intercept A . The parameter vector in this case is $\boldsymbol{\theta} = [A B]^T$. We need to first compute the 2×2 Fisher information matrix,

$$\mathbf{I}(\boldsymbol{\theta}) = \begin{bmatrix} -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A^2} \right] & -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A \partial B} \right] \\ -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial B \partial A} \right] & -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial B^2} \right] \end{bmatrix}.$$

The likelihood function is

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn)^2 \right\}$$

from which the derivatives follow as

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn) \\ \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial B} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn)n \end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A^2} &= -\frac{N}{\sigma^2} \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial A \partial B} &= -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} n \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial B^2} &= -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} n^2.\end{aligned}$$

Since the second-order derivatives do not depend on $\mathbf{x}$, we have immediately that

$$\begin{aligned}\mathbf{I}(\boldsymbol{\theta}) &= \frac{1}{\sigma^2} \begin{bmatrix} N & \sum_{n=0}^{N-1} n \\ \sum_{n=0}^{N-1} n & \sum_{n=0}^{N-1} n^2 \end{bmatrix} \\ &= \frac{1}{\sigma^2} \begin{bmatrix} N & \frac{N(N-1)}{2} \\ \frac{N(N-1)}{2} & \frac{N(N-1)(2N-1)}{6} \end{bmatrix}\end{aligned}$$

where we have used the identities

$$\begin{aligned}\sum_{n=0}^{N-1} n &= \frac{N(N-1)}{2} \\ \sum_{n=0}^{N-1} n^2 &= \frac{N(N-1)(2N-1)}{6}.\end{aligned}\tag{3.22}$$

Inverting the matrix yields

$$\mathbf{I}^{-1}(\boldsymbol{\theta}) = \sigma^2 \begin{bmatrix} \frac{2(2N-1)}{N(N+1)} & -\frac{6}{N(N+1)} \\ -\frac{6}{N(N+1)} & \frac{12}{N(N^2-1)} \end{bmatrix}.$$

It follows from (3.20) that the CRLB is

$$\begin{aligned}\text{var}(\hat{A}) &\geq \frac{2(2N-1)\sigma^2}{N(N+1)} \\ \text{var}(\hat{B}) &\geq \frac{12\sigma^2}{N(N^2-1)}.\end{aligned}$$




Figure 3.5 Sensitivity of observations to parameter changes—no noise

Some interesting observations follow from examination of the CRLB. Note first that the CRLB for A has increased over that obtained when B is known, for in the latter case we have

$$\text{var}(\hat{A}) \geq -\frac{1}{E \left[\frac{\partial^2 \ln p(\mathbf{x}; A)}{\partial A^2} \right]} = \frac{\sigma^2}{N}$$

and for $N \geq 2$, $2(2N-1)/(N+1) > 1$. This is a quite general result that asserts that the CRLB always increases as we estimate more parameters (see Problems 3.11 and 3.12). A second point is that

$$\frac{\text{CRLB}(\hat{A})}{\text{CRLB}(\hat{B})} = \frac{(2N-1)(N-1)}{6} > 1$$

for $N \geq 3$. Hence, B is easier to estimate, its CRLB decreasing as $1/N^3$ as opposed to the $1/N$ dependence for the CRLB of A . These differing dependences indicate that $x[n]$ is more sensitive to changes in B than to changes in A . A simple calculation reveals

$$\begin{aligned} \Delta x[n] &\approx \frac{\partial x[n]}{\partial A} \Delta A = \Delta A \\ \Delta x[n] &\approx \frac{\partial x[n]}{\partial B} \Delta B = n \Delta B. \end{aligned}$$

Changes in B are magnified by n , as illustrated in Figure 3.5. This effect is reminiscent of (3.14), and indeed a similar type of relationship is obtained in the vector parameter case (see (3.33)). See Problem 3.13 for a generalization of this example. $\diamond$

As an alternative means of computing the CRLB we can use the identity

$$E \left[\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_j} \right] = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right] \quad (3.23)$$

as shown in Appendix 3B. The form given on the right-hand side is usually easier to evaluate, however.

We now formally state the CRLB theorem for a vector parameter. Included in the theorem are conditions for equality. The bound is stated in terms of the covariance matrix of $\hat{\theta}$, denoted by $\mathbf{C}_{\hat{\theta}}$, from which (3.20) follows.

Theorem 3.2 (Cramer-Rao Lower Bound - Vector Parameter) It is assumed that the PDF $p(\mathbf{x}; \theta)$ satisfies the "regularity" conditions

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] = 0 \quad \text{for all } \theta,$$

where the expectation is taken with respect to $p(\mathbf{x}; \theta)$. Then, the covariance matrix of any unbiased estimator $\hat{\theta}$ satisfies

$$\mathbf{C}_{\hat{\theta}} - \mathbf{I}^{-1}(\theta) \geq 0 \quad (3.24)$$

where ≥ 0 is interpreted as meaning that the matrix is positive semidefinite. The Fisher information matrix $\mathbf{I}(\theta)$ is given as

$$[\mathbf{I}(\theta)]_{ij} = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta_i \partial \theta_j} \right]$$

where the derivatives are evaluated at the true value of θ and the expectation is taken with respect to $p(\mathbf{x}; \theta)$. Furthermore, an unbiased estimator may be found that attains the bound in that $\mathbf{C}_{\hat{\theta}} = \mathbf{I}^{-1}(\theta)$ if and only if

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \mathbf{I}(\theta)(\mathbf{g}(\mathbf{x}) - \theta) \quad (3.25)$$

for some p -dimensional function $\mathbf{g}$ and some $p \times p$ matrix $\mathbf{I}$. That estimator, which is the MVU estimator, is $\hat{\theta} = \mathbf{g}(\mathbf{x})$, and its covariance matrix is $\mathbf{I}^{-1}(\theta)$.

The proof is given in Appendix 3B. That (3.20) follows from (3.24) is shown by noting that for a positive semidefinite matrix the diagonal elements are nonnegative. Hence,

$$[\mathbf{C}_{\hat{\theta}} - \mathbf{I}^{-1}(\theta)]_{ii} \geq 0$$

and therefore

$$\text{var}(\hat{\theta}_i) = [\mathbf{C}_{\hat{\theta}}]_{ii} \geq [\mathbf{I}^{-1}(\theta)]_{ii}. \quad (3.26)$$

Additionally, when equality holds or $\mathbf{C}_{\hat{\theta}} = \mathbf{I}^{-1}(\theta)$, then (3.26) holds with equality also. The conditions for the CRLB to be attained are of particular interest since then $\hat{\theta} = \mathbf{g}(\mathbf{x})$ is efficient and hence is the MVU estimator. An example of equality occurs in Example 3.7. There we found that

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \begin{bmatrix} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial A} \\ \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial B} \end{bmatrix} \quad (3.27)$$

$$= \begin{bmatrix} \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn) \\ \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A - Bn)n \end{bmatrix}. \quad (3.28)$$

Although not obvious, this may be rewritten as

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \begin{bmatrix} \frac{N}{\sigma^2} & \frac{N(N-1)}{2\sigma^2} \\ \frac{N(N-1)}{2\sigma^2} & \frac{N(N-1)(2N-1)}{6\sigma^2} \end{bmatrix} \begin{bmatrix} \hat{A} - A \\ \hat{B} - B \end{bmatrix} \quad (3.29)$$

where

$$\begin{aligned} \hat{A} &= \frac{2(2N-1)}{N(N+1)} \sum_{n=0}^{N-1} x[n] - \frac{6}{N(N+1)} \sum_{n=0}^{N-1} nx[n] \\ \hat{B} &= -\frac{6}{N(N+1)} \sum_{n=0}^{N-1} x[n] + \frac{12}{N(N^2-1)} \sum_{n=0}^{N-1} nx[n]. \end{aligned}$$

Hence, the conditions for equality are satisfied and $[\hat{A} \hat{B}]^T$ is an efficient and therefore MVU estimator. Furthermore, the matrix in (3.29) is the inverse of the covariance matrix.

If the equality conditions hold, the reader may ask whether we can be assured that $\hat{\boldsymbol{\theta}}$ is unbiased. Because the regularity conditions

$$E \left[\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right] = 0$$

are always assumed to hold, we can apply them to (3.25). This then yields $E[\mathbf{g}(\mathbf{x})] = E(\hat{\boldsymbol{\theta}}) = \boldsymbol{\theta}$.

In finding MVU estimators for a vector parameter the CRLB theorem provides a powerful tool. In particular, it allows us to find the MVU estimator for an important class of data models. This class is the linear model and is described in detail in Chapter 4. The line fitting example just discussed is a special case. Suffice it to say that if we can model our data in the linear model form, then the MVU estimator and its performance are easily found.

3.8 Vector Parameter CRLB for Transformations

The discussion in Section 3.6 extends readily to the vector case. Assume that it is desired to estimate $\boldsymbol{\alpha} = \mathbf{g}(\boldsymbol{\theta})$ for $\mathbf{g}$, an r -dimensional function. Then, as shown in Appendix 3B

$$\mathbf{C}_{\hat{\boldsymbol{\alpha}}} - \frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{I}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{g}(\boldsymbol{\theta})^T}{\partial \boldsymbol{\theta}} \geq \mathbf{0} \quad (3.30)$$

where, as before, ≥ 0 is to be interpreted as positive semidefinite. In (3.30), $\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$ is the $r \times p$ **Jacobian matrix** defined as

$$\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \begin{bmatrix} \frac{\partial g_1(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial g_1(\boldsymbol{\theta})}{\partial \theta_2} & \cdots & \frac{\partial g_1(\boldsymbol{\theta})}{\partial \theta_p} \\ \frac{\partial g_2(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial g_2(\boldsymbol{\theta})}{\partial \theta_2} & \cdots & \frac{\partial g_2(\boldsymbol{\theta})}{\partial \theta_p} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial g_r(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial g_r(\boldsymbol{\theta})}{\partial \theta_2} & \cdots & \frac{\partial g_r(\boldsymbol{\theta})}{\partial \theta_p} \end{bmatrix}.$$

Example 3.8 - CRLB for Signal-to-Noise Ratio

Consider a DC level in WGN with A and σ^2 unknown. We wish to estimate

$$\alpha = \frac{A^2}{\sigma^2}$$

which can be considered to be the SNR for a single sample. Here $\boldsymbol{\theta} = [A \ \sigma^2]^T$ and $g(\boldsymbol{\theta}) = \theta_1^2 / \theta_2 = A^2 / \sigma^2$. Then, as shown in Example 3.6,

$$\mathbf{I}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{N}{\sigma^2} & 0 \\ 0 & \frac{N}{2\sigma^4} \end{bmatrix}.$$

The Jacobian is

$$\begin{aligned} \frac{\partial g(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} &= \begin{bmatrix} \frac{\partial g(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial g(\boldsymbol{\theta})}{\partial \theta_2} \end{bmatrix} = \begin{bmatrix} \frac{\partial g(\boldsymbol{\theta})}{\partial A} & \frac{\partial g(\boldsymbol{\theta})}{\partial \sigma^2} \end{bmatrix} \\ &= \begin{bmatrix} \frac{2A}{\sigma^2} & -\frac{A^2}{\sigma^4} \end{bmatrix} \end{aligned}$$

so that

$$\begin{aligned} \frac{\partial g(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{I}^{-1}(\boldsymbol{\theta}) \frac{\partial g(\boldsymbol{\theta})^T}{\partial \boldsymbol{\theta}} &= \begin{bmatrix} \frac{2A}{\sigma^2} & -\frac{A^2}{\sigma^4} \end{bmatrix} \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{2\sigma^4}{N} \end{bmatrix} \begin{bmatrix} \frac{2A}{\sigma^2} \\ -\frac{A^2}{\sigma^4} \end{bmatrix} \\ &= \frac{4A^2}{N\sigma^2} + \frac{2A^4}{N\sigma^4} \\ &= \frac{4\alpha + 2\alpha^2}{N}. \end{aligned}$$

Finally, since α is a scalar

$$\text{var}(\hat{\alpha}) \geq \frac{4\alpha + 2\alpha^2}{N}.$$

◇

As discussed in Section 3.6, efficiency is maintained over linear transformations

$$\alpha = g(\theta) = A\theta + b$$

where A is an $r \times p$ matrix and b is an $r \times 1$ vector. If $\hat{\alpha} = A\hat{\theta} + b$, and $\hat{\theta}$ is efficient or $C_{\hat{\theta}} = I^{-1}(\theta)$, then

$$E(\hat{\alpha}) = A\theta + b = \alpha$$

so that $\hat{\alpha}$ is unbiased and

$$\begin{aligned} C_{\hat{\alpha}} &\Rightarrow AC_{\hat{\theta}}A^T = AI^{-1}(\theta)A^T \\ &= \frac{\partial g(\theta)}{\partial \theta} I^{-1}(\theta) \frac{\partial g(\theta)^T}{\partial \theta} \end{aligned}$$

the latter being the CRLB. For nonlinear transformations efficiency is maintained only as $N \rightarrow \infty$. (This assumes that the PDF of $\hat{\theta}$ becomes concentrated about the true value of θ as $N \rightarrow \infty$ or that $\hat{\theta}$ is consistent.) Again this is due to the statistical linearity of $g(\theta)$ about the true value of θ .

3.9 CRLB for the General Gaussian Case

It is quite convenient at times to have a general expression for the CRLB. In the case of Gaussian observations we can derive the CRLB that generalizes (3.14). Assume that

$$\mathbf{x} \sim \mathcal{N}(\mu(\theta), C(\theta))$$

so that both the mean and covariance may depend on θ . Then, as shown in Appendix 3C, the Fisher information matrix is given by

$$\begin{aligned} [I(\theta)]_{ij} &= \left[\frac{\partial \mu(\theta)}{\partial \theta_i} \right]^T C^{-1}(\theta) \left[\frac{\partial \mu(\theta)}{\partial \theta_j} \right] \\ &+ \frac{1}{2} \text{tr} \left[C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_i} C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_j} \right] \end{aligned} \quad (3.31)$$

where

$$\frac{\partial \mu(\theta)}{\partial \theta_i} = \begin{bmatrix} \frac{\partial [\mu(\theta)]_1}{\partial \theta_i} \\ \frac{\partial [\mu(\theta)]_2}{\partial \theta_i} \\ \vdots \\ \frac{\partial [\mu(\theta)]_N}{\partial \theta_i} \end{bmatrix}$$

and

$$\frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_i} = \begin{bmatrix} \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{11}}{\partial \theta_i} & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{12}}{\partial \theta_i} & \dots & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{1N}}{\partial \theta_i} \\ \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{21}}{\partial \theta_i} & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{22}}{\partial \theta_i} & \dots & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{2N}}{\partial \theta_i} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{N1}}{\partial \theta_i} & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{N2}}{\partial \theta_i} & \dots & \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{NN}}{\partial \theta_i} \end{bmatrix}$$

For the scalar parameter case in which

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\theta), \mathbf{C})$$

this reduces to

$$I(\theta) = \left[\frac{\partial \boldsymbol{\mu}(\theta)}{\partial \theta} \right]^T \mathbf{C}^{-1} \left[\frac{\partial \boldsymbol{\mu}(\theta)}{\partial \theta} \right] + \frac{1}{2} \text{tr} \left[\left(\mathbf{C}^{-1} - \frac{\partial \mathbf{C}(\theta)}{\partial \theta} \right)^2 \right] \quad (3.32)$$

which generalizes (3.14). We now illustrate the computation with some examples.

Example 3.9 - Parameters of a Signal in White Gaussian Noise

Assume that we wish to estimate a scalar signal parameter θ for the data set

$$x[n] = s[n; \theta] + w[n] \quad n = 0, \dots, N-1$$

where $w[n]$ is WGN. The covariance matrix is $\mathbf{C} = \sigma^2 \mathbf{I}$ and does not depend on θ . The second term in (3.32) is therefore zero. The first term yields

$$\begin{aligned} I(\theta) &= \frac{1}{\sigma^2} \left[\frac{\partial \boldsymbol{\mu}(\theta)}{\partial \theta} \right]^T \left[\frac{\partial \boldsymbol{\mu}(\theta)}{\partial \theta} \right] \\ &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left(\frac{\partial s[n; \theta]}{\partial \theta} \right)^2 \\ &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left(\frac{\partial s[n; \theta]}{\partial \theta} \right)^2 \end{aligned}$$

which agrees with (3.14). ◇

Generalizing to a vector signal parameter estimator in the presence of WGN, we have from (3.31)

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = \left[\frac{\partial \boldsymbol{\mu}(\boldsymbol{\theta})}{\partial \theta_i} \right]^T \mathbf{C}^{-1} \left[\frac{\partial \boldsymbol{\mu}(\boldsymbol{\theta})}{\partial \theta_j} \right]$$

which yields

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \frac{\partial s[n; \boldsymbol{\theta}]}{\partial \theta_i} \frac{\partial s[n; \boldsymbol{\theta}]}{\partial \theta_j} \quad (3.33)$$

as the elements of the Fisher information matrix.

Example 3.10 - Parameter of Noise

Assume that we observe

$$x[n] = w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with unknown variance $\theta = \sigma^2$. Then, according to (3.32), since $\mathbf{C}(\sigma^2) = \sigma^2 \mathbf{I}$, we have

$$\begin{aligned} I(\sigma^2) &= \frac{1}{2} \text{tr} \left[\left(\mathbf{C}^{-1}(\sigma^2) \frac{\partial \mathbf{C}(\sigma^2)}{\partial \sigma^2} \right)^2 \right] \\ &= \frac{1}{2} \text{tr} \left[\left(\left(\frac{1}{\sigma^2} \right) (\mathbf{I}) \right)^2 \right] \\ &= \frac{1}{2} \text{tr} \left[\frac{1}{\sigma^4} \mathbf{I} \right] \\ &= \frac{N}{2\sigma^4} \end{aligned}$$

which agrees with the results in Example 3.6. A slightly more complicated example follows. $\diamond$

Example 3.11 - Random DC Level in WGN

Consider the data

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN and A , the DC level, is a Gaussian *random variable* with zero mean and variance σ_A^2 . Also, A is independent of $w[n]$. The power of the signal or variance σ_A^2 is the unknown parameter. Then, $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[N-1]]^T$ is Gaussian with zero mean and an $N \times N$ covariance matrix whose $[i, j]$ element is

$$\begin{aligned} [\mathbf{C}(\sigma_A^2)]_{ij} &= E[x[i-1]x[j-1]] \\ &= E[(A + w[i-1])(A + w[j-1])] \\ &= \sigma_A^2 + \sigma^2 \delta_{ij}. \end{aligned}$$

Therefore,

$$\mathbf{C}(\sigma_A^2) = \sigma_A^2 \mathbf{1}\mathbf{1}^T + \sigma^2 \mathbf{I}$$

where $\mathbf{1} = [1 \ 1 \dots 1]^T$. Using Woodbury's identity (see Appendix 1), we have

$$\mathbf{C}^{-1}(\sigma_A^2) = \frac{1}{\sigma^2} \left(\mathbf{I} - \frac{\sigma_A^2}{\sigma^2 + N\sigma_A^2} \mathbf{1}\mathbf{1}^T \right).$$

Also, since

$$\frac{\partial \mathbf{C}(\sigma_A^2)}{\partial \sigma_A^2} = \mathbf{1}\mathbf{1}^T$$

we have that

$$\mathbf{C}^{-1}(\sigma_A^2) \frac{\partial \mathbf{C}(\sigma_A^2)}{\partial \sigma_A^2} = \frac{1}{\sigma^2 + N\sigma_A^2} \mathbf{1}\mathbf{1}^T.$$

Substituting this in (3.32) produces

$$\begin{aligned} I(\sigma_A^2) &= \frac{1}{2} \text{tr} \left[\left(\frac{1}{\sigma^2 + N\sigma_A^2} \right)^2 \mathbf{1}\mathbf{1}^T \mathbf{1}\mathbf{1}^T \right] \\ &= \frac{N}{2} \left(\frac{1}{\sigma^2 + N\sigma_A^2} \right)^2 \text{tr}(\mathbf{1}\mathbf{1}^T) \\ &= \frac{1}{2} \left(\frac{N}{\sigma^2 + N\sigma_A^2} \right)^2 \end{aligned}$$

so that the CRLB is

$$\text{var}(\sigma_A^2) \geq 2 \left(\sigma_A^2 + \frac{\sigma^2}{N} \right)^2.$$

Note that even as $N \rightarrow \infty$, the CRLB does not decrease below $2\sigma_A^4$. This is because each additional data sample yields the same value of A (see Problem 3.14). $\diamond$

3.10 Asymptotic CRLB for WSS Gaussian Random Processes

At times it is difficult to analytically compute the CRLB using (3.31) due to the need to invert the covariance matrix. Of course, we can always resort to a computer evaluation. An alternative form which can be applied to Gaussian processes that are WSS is very useful. It is easily computed and provides much insight due to its simplified form. The principal drawback is that strictly speaking it is valid only as $N \rightarrow \infty$ or asymptotically. In practice it provides a good approximation to the true CRLB if the data record length N is much greater than the correlation time of the process. The correlation time is defined as the maximum lag k of the ACF $r_{xx}[k] = E[x[n]x[n+k]]$ for which the ACF is essentially nonzero. Hence, for processes with broad PSDs the approximation will be good for moderate length data records, while for narrowband processes longer length data records are required.




Figure 3.6 Signal PSD for center frequency estimation

As shown in Appendix 3D, the elements of the Fisher information are approximately (as $N \rightarrow \infty$)

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{\partial \ln P_{xx}(f; \boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \ln P_{xx}(f; \boldsymbol{\theta})}{\partial \theta_j} df \quad (3.34)$$

where $P_{xx}(f; \boldsymbol{\theta})$ is the PSD of the process with the explicit dependence on $\boldsymbol{\theta}$ shown. It is assumed that the mean of $x[n]$ is zero. This form, which is somewhat reminiscent of (3.33), allows us to examine the accuracy with which PSD, or equivalently, covariance parameters may be estimated.

Example 3.12 - Center Frequency of Process

A typical problem is to estimate the center frequency f_c of a PSD which otherwise is known. Given

$$P_{xx}(f; f_c) = Q(f - f_c) + Q(-f - f_c) + \sigma^2$$

we wish to determine the CRLB for f_c assuming that $Q(f)$ and σ^2 are known. We may view the process as consisting of a random signal embedded in WGN. The center frequency of the signal PSD is to be estimated. The real function $Q(f)$ and the signal PSD $P_{ss}(f; f_c)$ are shown in Figure 3.6. Note that the possible center frequencies are constrained to be in the interval $[f_1, 1/2 - f_2]$. For these center frequencies the signal PSD for $f \geq 0$ will be contained in the $[0, 1/2]$ interval. Then, since $\boldsymbol{\theta} = f_c$ is a scalar,

we have from (3.34)

$$\text{var}(\hat{f}_c) \geq \frac{1}{\frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{\partial \ln P_{xx}(f; f_c)}{\partial f_c} \right)^2 df}.$$

But

$$\begin{aligned} \frac{\partial \ln P_{xx}(f; f_c)}{\partial f_c} &= \frac{\partial \ln [Q(f - f_c) + Q(-f - f_c) + \sigma^2]}{\partial f_c} \\ &= \frac{\frac{\partial Q(f - f_c)}{\partial f_c} + \frac{\partial Q(-f - f_c)}{\partial f_c}}{Q(f - f_c) + Q(-f - f_c) + \sigma^2}. \end{aligned}$$

This is an odd function of f , so that

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{\partial \ln P_{xx}(f; f_c)}{\partial f_c} \right)^2 df = 2 \int_0^{\frac{1}{2}} \left(\frac{\partial \ln P_{xx}(f; f_c)}{\partial f_c} \right)^2 df.$$

Also, for $f \geq 0$ we have that $Q(-f - f_c) = 0$, and thus its derivative is zero due to the assumption illustrated in Figure 3.6. It follows that

$$\begin{aligned} \text{var}(\hat{f}_c) &\geq \frac{1}{N \int_0^{\frac{1}{2}} \left(\frac{\frac{\partial Q(f - f_c)}{\partial f_c}}{Q(f - f_c) + \sigma^2} \right)^2 df} \\ &= \frac{1}{N \int_0^{\frac{1}{2}} \left(\frac{\frac{\partial Q(f - f_c)}{\partial (f - f_c)} (-1)}{Q(f - f_c) + \sigma^2} \right)^2 df} \\ &= \frac{1}{N \int_{-f_c}^{\frac{1}{2} - f_c} \left(\frac{\frac{\partial Q(f')}{\partial f'}}{Q(f') + \sigma^2} \right)^2 df'} \end{aligned}$$

where we have let $f' = f - f_c$. But $1/2 - f_c \geq 1/2 - f_{c_{\max}} = f_2$ and $-f_c \leq -f_{c_{\min}} = -f_1$, so that we may change the limits of integration to the interval $[-1/2, 1/2]$. Thus,

$$\begin{aligned} \text{var}(\hat{f}_c) &\geq \frac{1}{N \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{\frac{\partial Q(f)}{\partial f}}{Q(f) + \sigma^2} \right)^2 df} \\ &= \frac{1}{N \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{\partial \ln(Q(f) + \sigma^2)}{\partial f} \right)^2 df}. \end{aligned}$$

As an example, consider

$$Q(f) = \exp \left[-\frac{1}{2} \left(\frac{f}{\sigma_f} \right)^2 \right]$$

where $\sigma_f \ll 1/2$, so that $Q(f)$ is bandlimited as shown in Figure 3.6. Then, if $Q(f) \gg \sigma^2$, we have approximately

$$\text{var}(\hat{f}_c) \geq \frac{1}{N \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{f^2}{\sigma_f^4} df} = \frac{12\sigma_f^4}{N}.$$

Narrower bandwidth (smaller σ_f^2) spectra yield lower bounds for the center frequency since the PSD changes more rapidly as f_c changes. See also Problem 3.16 for another example. $\diamond$

3.11 Signal Processing Examples

We now apply the theory of the CRLB to several signal processing problems of interest. The problems to be considered and some of their areas of application are:

1. Range estimation - sonar, radar, robotics
2. Frequency estimation - sonar, radar, econometrics, spectrometry
3. Bearing estimation - sonar, radar
4. Autoregressive parameter estimation - speech, econometrics.

These examples will be revisited in Chapter 7, in which actual estimators that asymptotically attain the CRLB will be studied.

Example 3.13 - Range Estimation

In radar or active sonar a signal pulse is transmitted. The round trip delay τ_0 from the transmitter to the target and back is related to the range R as $\tau_0 = 2R/c$, where c is the speed of propagation. Estimation of range is therefore equivalent to estimation of the time delay, assuming that c is known. If $s(t)$ is the transmitted signal, a simple model for the received continuous waveform is

$$\mathbf{x}(t) = \mathbf{s}(t - \tau_0) + \mathbf{w}(t) \quad 0 \leq t \leq T.$$

The transmitted signal pulse is assumed to be nonzero over the interval $[0, T_s]$. Additionally, the signal is assumed to be essentially bandlimited to B Hz. If the maximum time delay is $\tau_{0\max}$, then the observation interval is chosen to include the entire signal by letting $T = T_s + \tau_{0\max}$. The noise is modeled as Gaussian with PSD and ACF as




Figure 3.7 Properties of Gaussian observation noise

shown in Figure 3.7. The bandlimited nature of the noise results from filtering the continuous waveform to the signal bandwidth of B Hz. The continuous received waveform is sampled at the Nyquist rate or samples are taken every $\Delta = 1/(2B)$ seconds to form the observed data

$$x(n\Delta) = s(n\Delta - \tau_0) + w(n\Delta) \quad n = 0, 1, \dots, N-1.$$

Letting $x[n]$ and $w[n]$ be the sampled sequences, we have our discrete data model

$$x[n] = s(n\Delta - \tau_0) + w[n]. \quad (3.35)$$

Note that $w[n]$ is WGN since the samples are separated by $k\Delta = k/(2B)$, which corresponds to the zeros of the ACF of $w(t)$ as shown in Figure 3.7. Also, $w[n]$ has variance $\sigma^2 = r_{ww}(0) = N_0 B$. Because the signal is nonzero only over the interval $\tau_0 \leq t \leq \tau_0 + T_s$, (3.35) reduces to

$$x[n] = \begin{cases} w[n] & 0 \leq n \leq n_0 - 1 \\ s(n\Delta - \tau_0) + w[n] & n_0 \leq n \leq n_0 + M - 1 \\ w[n] & n_0 + M \leq n \leq N - 1 \end{cases} \quad (3.36)$$

where M is the length of the sampled signal and $n_0 = \tau_0/\Delta$ is the delay in samples. (For simplicity we assume that Δ is so small that τ_0/Δ can be approximated by an integer.) With this formulation we can apply (3.14) in evaluating the CRLB.

$$\begin{aligned} \text{var}(\hat{\tau}_0) &\geq \frac{\sigma^2}{\sum_{n=0}^{N-1} \left(\frac{\partial s[n; \tau_0]}{\partial \tau_0} \right)^2} \\ &= \frac{\sigma^2}{\sum_{n=n_0}^{n_0+M-1} \left(\frac{\partial s(n\Delta - \tau_0)}{\partial \tau_0} \right)^2} \\ &= \frac{\sigma^2}{\sum_{n=n_0}^{n_0+M-1} \left(\left. \frac{ds(t)}{dt} \right|_{t=n\Delta - \tau_0} \right)^2} \end{aligned}$$

$$= \frac{\sigma^2}{\sum_{n=0}^{M-1} \left(\left. \frac{ds(t)}{dt} \right|_{t=n\Delta} \right)^2}$$

since $\tau_0 = n_0\Delta$. Assuming that Δ is small enough to approximate the sum by an integral, we have

$$\text{var}(\hat{\tau}_0) \geq \frac{\sigma^2}{\frac{1}{\Delta} \int_0^{T_s} \left(\frac{ds(t)}{dt} \right)^2 dt}$$

Finally, noting that $\Delta = 1/(2B)$ and $\sigma^2 = N_0B$, we have

$$\text{var}(\hat{\tau}_0) \geq \frac{\frac{N_0}{2}}{\int_0^{T_s} \left(\frac{ds(t)}{dt} \right)^2 dt} \quad (3.37)$$

An alternative form observes that the energy $\mathcal{E}$ is

$$\mathcal{E} = \int_0^{T_s} s^2(t) dt$$

which results in

$$\text{var}(\hat{\tau}_0) \geq \frac{1}{\frac{\mathcal{E}}{N_0/2} \overline{F^2}} \quad (3.38)$$

where

$$\overline{F^2} = \frac{\int_0^{T_s} \left(\frac{ds(t)}{dt} \right)^2 dt}{\int_0^{T_s} s^2(t) dt}$$

It can be shown that $\mathcal{E}/(N_0/2)$ is a SNR [Van Trees 1968]. Also, $\overline{F^2}$ is a measure of the bandwidth of the signal since, using standard Fourier transform properties,

$$\overline{F^2} = \frac{\int_{-\infty}^{\infty} (2\pi F)^2 |S(F)|^2 dF}{\int_{-\infty}^{\infty} |S(F)|^2 dF} \quad (3.39)$$

where F denotes continuous-time frequency, and $S(F)$ is the Fourier transform of $s(t)$. In this form it becomes clear that $\overline{F^2}$ is the mean square bandwidth of the signal. From (3.38) and (3.39), the larger the mean square bandwidth, the lower the CRLB. For instance, assume that the signal is a Gaussian pulse given by $s(t) = \exp(-\frac{1}{2}\sigma_F^2(t - T_s/2)^2)$ and that $s(t)$ is essentially nonzero over the interval $[0, T_s]$. Then $|\hat{S}(F)| =$

$(\sigma_F/\sqrt{2\pi}) \exp(-2\pi^2 F^2/\sigma_F^2)$ and $\overline{F^2} = \sigma_F^2/2$. As the mean square bandwidth increases, the signal pulse becomes narrower and it becomes easier to estimate the time delay.

Finally, by noting that $R = c\tau_0/2$ and using (3.16), the CRLB for range is

$$\text{var}(\hat{R}) \geq \frac{c^2/4}{\frac{\mathcal{E}}{N_0/2} F^2}. \quad (3.40)$$

◇

Example 3.14 - Sinusoidal Parameter Estimation

In many fields we are confronted with the problem of estimating the parameters of a sinusoidal signal. Economic data which are cyclical in nature may naturally fit such a model, while in sonar and radar physical mechanisms cause the observed signal to be sinusoidal. Hence, we examine the determination of the CRLB for the amplitude A , frequency f_0 , and phase ϕ of a sinusoid embedded in WGN. This example generalizes Examples 3.4 and 3.5. The data are assumed to be

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n] \quad n = 0, 1, \dots, N-1$$

where $A > 0$ and $0 < f_0 < 1/2$ (otherwise the parameters are not identifiable, as is verified by considering $A = 1, \phi = 0$ versus $A = -1, \phi = \pi$ or $f_0 = 0$ with $A = 1/2, \phi = 0$ versus $A = 1/\sqrt{2}, \phi = \pi/4$). Since multiple parameters are unknown, we use (3.33)

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \frac{\partial s[n; \boldsymbol{\theta}]}{\partial \theta_i} \frac{\partial s[n; \boldsymbol{\theta}]}{\partial \theta_j}$$

for $\boldsymbol{\theta} = [A f_0 \phi]^T$. In evaluating the CRLB it is assumed that f_0 is not near 0 or 1/2, which allows us to make certain simplifications based on the approximations [Stoica 1989] (see also Problem 3.7):

$$\begin{aligned} \frac{1}{N^{i+1}} \sum_{n=0}^{N-1} n^i \sin(4\pi f_0 n + 2\phi) &\approx 0 \\ \frac{1}{N^{i+1}} \sum_{n=0}^{N-1} n^i \cos(4\pi f_0 n + 2\phi) &\approx 0 \end{aligned}$$

for $i = 0, 1, 2$. Using these approximations and letting $\alpha = 2\pi f_0 n + \phi$, we have

$$[\mathbf{I}(\boldsymbol{\theta})]_{11} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \cos^2 \alpha = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left(\frac{1}{2} + \frac{1}{2} \cos 2\alpha \right) \approx \frac{N}{2\sigma^2}$$

$$[\mathbf{I}(\boldsymbol{\theta})]_{12} = -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} A 2\pi n \cos \alpha \sin \alpha = -\frac{\pi A}{\sigma^2} \sum_{n=0}^{N-1} n \sin 2\alpha \approx 0$$

$$[\mathbf{I}(\boldsymbol{\theta})]_{13} = -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} A \cos \alpha \sin \alpha = -\frac{A}{2\sigma^2} \sum_{n=0}^{N-1} \sin 2\alpha \approx 0$$

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{22} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} A^2 (2\pi n)^2 \sin^2 \alpha = \frac{(2\pi A)^2}{\sigma^2} \sum_{n=0}^{N-1} n^2 \left(\frac{1}{2} - \frac{1}{2} \cos 2\alpha \right) \\ &\approx \frac{(2\pi A)^2}{2\sigma^2} \sum_{n=0}^{N-1} n^2 \end{aligned}$$

$$[\mathbf{I}(\boldsymbol{\theta})]_{23} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} A^2 2\pi n \sin^2 \alpha \approx \frac{\pi A^2}{\sigma^2} \sum_{n=0}^{N-1} n$$

$$[\mathbf{I}(\boldsymbol{\theta})]_{33} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} A^2 \sin^2 \alpha \approx \frac{NA^2}{2\sigma^2}.$$

The Fisher information matrix becomes

$$\mathbf{I}(\boldsymbol{\theta}) = \frac{1}{\sigma^2} \begin{bmatrix} \frac{N}{2} & 0 & 0 \\ 0 & 2A^2\pi^2 \sum_{n=0}^{N-1} n^2 & \pi A^2 \sum_{n=0}^{N-1} n \\ 0 & \pi A^2 \sum_{n=0}^{N-1} n & \frac{NA^2}{2} \end{bmatrix}.$$

Using (3.22), we have upon inversion

$$\begin{aligned} \text{var}(\hat{A}) &\geq \frac{2\sigma^2}{N} \\ \text{var}(\hat{f}_0) &\geq \frac{12}{(2\pi)^2 \eta N(N^2 - 1)} \\ \text{var}(\hat{\phi}) &\geq \frac{2(2N - 1)}{\eta N(N + 1)} \end{aligned} \quad (3.41)$$

where $\eta = A^2/(2\sigma^2)$ is the SNR. Frequency estimation of a sinusoid is of considerable interest. Note that the CRLB for the frequency decreases as the SNR increases and that the bound decreases as $1/N^3$, making it quite sensitive to data record length. See also Problem 3.17 for a variation of this example.

◇

Example 3.15 - Bearing Estimation

In sonar it is of interest to estimate the bearing to a target as shown in Figure 3.8. To do so the acoustic pressure field is observed by an array of equally spaced sensors in a




Figure 3.8 Geometry of array for bearing estimation

line. Assuming that the target radiates a sinusoidal signal $A \cos(2\pi F_0 t + \phi)$, then the received signal at the n th sensor is $A \cos(2\pi F_0(t - t_n) + \phi)$, where t_n is the propagation time to the n th sensor. If the array is located far from the target, then the circular wavefronts can be considered to be planar at the array. As shown in Figure 3.8, the wavefront at the $(n-1)$ st sensor lags that at the n th sensor by $d \cos \beta / c$ due to the extra propagation distance. Thus, the propagation time to the n th sensor is

$$t_n = t_0 - n \frac{d}{c} \cos \beta \quad n = 0, 1, \dots, M-1$$

where t_0 is the propagation time to the zeroth sensor, and the observed signal at the n th sensor is

$$s_n(t) = A \cos \left[2\pi F_0(t - t_0 + n \frac{d}{c} \cos \beta) + \phi \right].$$

If a single “snapshot” of data is taken or the array element outputs are sampled at a given time t_s , then

$$s_n(t_s) = A \cos \left[2\pi(F_0 \frac{d}{c} \cos \beta)n + \phi' \right] \quad (3.42)$$

where $\phi' = \phi + 2\pi F_0(t_s - t_0)$. In this form it becomes clear that the spatial observations are sinusoidal with frequency $f_s = F_0(d/c) \cos \beta$. To complete the description of the data we assume that the sensor outputs are corrupted by Gaussian noise with zero mean and variance σ^2 which is independent from sensor to sensor. The data are modeled as

$$x[n] = s_n(t_s) + w[n] \quad n = 0, 1, \dots, M-1$$

where $w[n]$ is WGN. Since typically A, ϕ are unknown, as well as β , we have the problem of estimating $\{A, f_s, \phi'\}$ based on (3.42) as in Example 3.14. Once the CRLB for these parameters is determined, we can use the transformation of parameters formula. The transformation is for $\theta = [A f_s \phi']^T$

$$\alpha = g(\theta) = \begin{bmatrix} A \\ \beta \\ \phi' \end{bmatrix} = \begin{bmatrix} A \\ \arccos \left(\frac{c f_s}{F_0 d} \right) \\ \phi' \end{bmatrix}.$$

The Jacobian is

$$\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\frac{c}{F_0 d \sin \beta} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

so that from (3.30)

$$\left[\mathbf{C}_{\hat{\alpha}} - \frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{I}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{g}(\boldsymbol{\theta})^T}{\partial \boldsymbol{\theta}} \right]_{22} \geq 0.$$

Because of the diagonal Jacobian this yields

$$\text{var}(\hat{\beta}) \geq \left[\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right]_{22}^2 [\mathbf{I}^{-1}(\boldsymbol{\theta})]_{22}.$$

But from (3.41) we have

$$[\mathbf{I}^{-1}(\boldsymbol{\theta})]_{22} = \frac{12}{(2\pi)^2 \eta M (M^2 - 1)}$$

and therefore

$$\text{var}(\hat{\beta}) \geq \frac{12}{(2\pi)^2 \eta M (M^2 - 1)} \frac{c^2}{F_0^2 d^2 \sin^2 \beta}$$

or finally

$$\text{var}(\hat{\beta}) \geq \frac{12}{(2\pi)^2 M \eta \frac{M+1}{M-1} \left(\frac{L}{\lambda} \right)^2 \sin^2 \beta} \quad (3.43)$$

where $\lambda = c/F_0$ is the wavelength of the propagating plane wave and $L = (M-1)d$ is the length of the array. Note that it is easiest to estimate bearing if $\beta = 90^\circ$, and impossible if $\beta = 0^\circ$. Also, the bound depends critically on the array length in wavelengths or L/λ , as well as the SNR at the array output or $M\eta$. The use of the CRLB for feasibility studies is examined in Problem 3.19. $\diamond$

Example 3.16 - Autoregressive Parameter Estimation

In speech processing an important model for speech production is the autoregressive (AR) process. As shown in Figure 3.9, the data are modeled as the output of a causal all-pole discrete filter excited at the input by WGN $u[n]$. The excitation noise $u[n]$ is an inherent part of the model, necessary to ensure that $x[n]$ is a WSS random process. The all-pole filter acts to model the vocal tract, while the excitation noise models the forcing of air through a constriction in the throat necessary to produce an unvoiced sound such as an "s." The effect of the filter is to color the white noise so as to model PSDs with several resonances. This model is also referred to as a **linear predictive coding (LPC) model** [Makhoul 1975]. Since the AR model is capable of producing a variety of PSDs, depending on the choice of the AR filter parameters $\{a[1], a[2], \dots, a[p]\}$ and




Figure 3.9 Autoregressive model

$$A(z) = 1 + \sum_{m=1}^p a[m]z^{-m}$$

excitation white noise variance σ_u^2 , it has also been successfully used for high-resolution spectral estimation. Based on observed data $\{x[0], x[1], \dots, x[N-1]\}$, the parameters are estimated (see Example 7.18), and hence the PSD is estimated as [Kay 1988]

$$\hat{P}_{xx}(f) = \frac{\hat{\sigma}_u^2}{\left| 1 + \sum_{m=1}^p \hat{a}[m] \exp(-j2\pi f m) \right|^2}$$

The derivation of the CRLB proves to be a difficult task. The interested reader may consult [Box and Jenkins 1970, Porat and Friedlander 1987] for details. In practice the asymptotic CRLB given by (3.34) is quite accurate, even for short data records of about $N = 100$ points if the poles are not too close to the unit circle in the z plane. Therefore, we now determine the asymptotic CRLB. The PSD implied by the AR model is

$$P_{xx}(f; \theta) = \frac{\sigma_u^2}{|A(f)|^2}$$

where $\theta = [a[1] a[2] \dots a[p] \sigma_u^2]^T$ and $A(f) = 1 + \sum_{m=1}^p a[m] \exp(-j2\pi f m)$. The partial derivatives are

$$\begin{aligned} \frac{\partial \ln P_{xx}(f; \theta)}{\partial a[k]} &= -\frac{\partial \ln |A(f)|^2}{\partial a[k]} \\ &= -\frac{1}{|A(f)|^2} [A(f) \exp(j2\pi f k) + A^*(f) \exp(-j2\pi f k)] \\ \frac{\partial \ln P_{xx}(f; \theta)}{\partial \sigma_u^2} &= \frac{1}{\sigma_u^2} \end{aligned}$$

For $k = 1, 2, \dots, p$; $l = 1, 2, \dots, p$ we have from (3.34)

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{kl} &= \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{|A(f)|^4} [A(f) \exp(j2\pi f k) + A^*(f) \exp(-j2\pi f k)] \\ &\quad \cdot [A(f) \exp(j2\pi f l) + A^*(f) \exp(-j2\pi f l)] df \\ &= \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\frac{1}{A^*(f)^2} \exp[j2\pi f(k+l)] + \frac{1}{|A(f)|^2} \exp[j2\pi f(k-l)] \right. \\ &\quad \left. + \frac{1}{|A(f)|^2} \exp[j2\pi f(l-k)] + \frac{1}{A^2(f)} \exp[-j2\pi f(k+l)] \right] df. \end{aligned}$$

Noting that

$$\begin{aligned} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{A^*(f)^2} \exp[j2\pi f(k+l)] df &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{A^2(f)} \exp[-j2\pi f(k+l)] df \\ \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{|A(f)|^2} \exp[j2\pi f(k-l)] df &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{|A(f)|^2} \exp[j2\pi f(l-k)] df \end{aligned}$$

which follows from the Hermitian property of the integrand (due to $A(-f) = A^*(f)$), we have

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{kl} &= N \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{|A(f)|^2} \exp[j2\pi f(k-l)] df \\ &\quad + N \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{A^*(f)^2} \exp[j2\pi f(k+l)] df. \end{aligned}$$

The second integral is the inverse Fourier transform of $1/A^*(f)^2$ evaluated at $n = k+l > 0$. This term is zero since the sequence is the convolution of two anticausal sequences, that is, if

$$\mathcal{F}^{-1} \left\{ \frac{1}{A(f)} \right\} = \begin{cases} h[n] & n \geq 0 \\ 0 & n < 0 \end{cases}$$

then

$$\begin{aligned} \mathcal{F}^{-1} \left\{ \frac{1}{A^*(f)^2} \right\} &= h[-n] \star h[-n] \\ &= 0 \quad \text{for } n > 0. \end{aligned}$$

Therefore,

$$[\mathbf{I}(\boldsymbol{\theta})]_{kl} = \frac{N}{\sigma_u^2} r_{xx}[k-l]$$

For $k = 1, 2, \dots, p; l = p + 1$

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{kl} &= -\frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{\sigma_u^2 |A(f)|^2} [A(f) \exp(j2\pi f k) + A^*(f) \exp(-j2\pi f k)] df \\ &= -\frac{N}{\sigma_u^2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{A^*(f)} \exp(j2\pi f k) df \\ &= 0 \end{aligned}$$

where again we have used the Hermitian property of the integrand and the anticausality of $\mathcal{F}^{-1}\{1/A^*(f)\}$. Finally, for $k = p + 1; l = p + 1$

$$[\mathbf{I}(\boldsymbol{\theta})]_{kl} = \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{\sigma_u^4} df = \frac{N}{2\sigma_u^4}$$

so that

$$\mathbf{I}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{N}{\sigma_u^2} \mathbf{R}_{xx} & \mathbf{0} \\ \mathbf{0}^T & \frac{N}{2\sigma_u^4} \end{bmatrix} \quad (3.44)$$

where $[\mathbf{R}_{xx}]_{ij} = r_{xx}[i - j]$ is a $p \times p$ Toeplitz autocorrelation matrix and $\mathbf{0}$ is a $p \times 1$ vector of zeros. Upon inverting the Fisher information matrix we have that

$$\begin{aligned} \text{var}(\hat{a}[k]) &\geq \frac{\sigma_u^2}{N} [\mathbf{R}_{xx}^{-1}]_{kk} \quad k = 1, 2, \dots, p \\ \text{var}(\hat{\sigma}_u^2) &\geq \frac{2\sigma_u^4}{N} \end{aligned} \quad (3.45)$$

As an illustration, if $p = 1$,

$$\text{var}(\hat{a}[1]) \geq \frac{\sigma_u^2}{Nr_{xx}[0]}.$$

But

$$r_{xx}[0] = \frac{\sigma_u^2}{1 - a^2[1]}$$

so that

$$\text{var}(\hat{a}[1]) \geq \frac{1}{N} (1 - a^2[1])$$

indicating that it is easier to estimate the filter parameter when $|a[1]|$ is closer to one than to zero. Since the pole of the filter is at $-a[1]$, this means that the filter parameters of processes with PSDs having sharp peaks are more easily estimated (see also Problem 3.20).

References

- Box, G.E.P., G.M. Jenkins, *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco, 1970.
- Brockwell, P.J., R.A. Davis, *Time Series: Theory and Methods*, Springer-Verlag, New York, 1987.
- Kay, S.M., *Modern Spectral Estimation: Theory and Application*, Prentice-Hall, Englewood Cliffs, N.J., 1988.
- Kendall, Sir M., A. Stuart, *The Advanced Theory of Statistics*, Vol. 2, Macmillan, New York, 1979.
- Makhoul, J., "Linear Prediction: A Tutorial Review," *IEEE Proc.*, Vol. 63, pp. 561-580, April 1975.
- McAulay, R.J., E.M. Hofstetter, "Barankin Bounds on Parameter Estimation," *IEEE Trans. Inform. Theory*, Vol. 17, pp. 669-676, Nov. 1971.
- Porat, B., B. Friedlander, "Computation of the Exact Information Matrix of Gaussian Time Series With Stationary Random Components," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 34, pp. 118-130, Feb. 1986.
- Porat, B., B. Friedlander, "The Exact Cramer-Rao Bound for Gaussian Autoregressive Processes," *IEEE Trans. Aerosp. Electron. Syst.*, Vol. 23, pp. 537-541, July 1987.
- Seidman, L.P., "Performance Limitations and Error Calculations for Parameter Estimation," *Proc. IEEE*, Vol. 58, pp. 644-652, May 1970.
- Stoica, P., R.L. Moses, B. Friedlander, T. Soderstrom, "Maximum Likelihood Estimation of the Parameters of Multiple Sinusoids from Noisy Measurements," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 37, pp. 378-392, March 1989.
- Van Trees, H.L., *Detection, Estimation, and Modulation Theory, Part I*, J. Wiley, New York, 1968.
- Ziv, J., M. Zakai, "Some Lower Bounds on Signal Parameter Estimation," *IEEE Trans. Inform. Theory*, Vol. 15, pp. 386-391, May 1969.

Problems

- 3.1 If $x[n]$ for $n = 0, 1, \dots, N-1$ are IID according to $\mathcal{U}[0, \theta]$, show that the regularity condition does not hold or that

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] \neq 0 \quad \text{for all } \theta > 0.$$

Hence, the CRLB cannot be applied to this problem.

- 3.2 In Example 3.1 assume that $w[0]$ has the PDF $p(w[0])$ which can now be arbitrary. Show that the CRLB for A is

$$\text{var}(\hat{A}) \geq \left[\int_{-\infty}^{\infty} \frac{\left(\frac{dp(u)}{du} \right)^2}{p(u)} du \right]^{-1}.$$

Evaluate this for the Laplacian PDF

$$p(w[0]) = \frac{1}{\sqrt{2}\sigma} \exp \left(-\frac{\sqrt{2}|w[0]|}{\sigma} \right)$$

and compare the result to the Gaussian case.

- 3.3** The data $\mathbf{x}[n] = Ar^n + w[n]$ for $n = 0, 1, \dots, N-1$ are observed, where $w[n]$ is WGN with variance σ^2 and $r > 0$ is known. Find the CRLB for A . Show that an efficient estimator exists and find its variance. What happens to the variance as $N \rightarrow \infty$ for various values of r ?
- 3.4** If $\mathbf{x}[n] = r^n + w[n]$ for $n = 0, 1, \dots, N-1$ are observed, where $w[n]$ is WGN with variance σ^2 and r is to be estimated, find the CRLB. Does an efficient estimator exist and if so find its variance?
- 3.5** If $\mathbf{x}[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$ are observed and $\mathbf{w} = [w[0] \ w[1] \ \dots \ w[N-1]]^T \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, find the CRLB for A . Does an efficient estimator exist and if so what is its variance?
- 3.6** For Example 2.3 compute the CRLB. Does it agree with the results given?
- 3.7** Prove that in Example 3.4

$$\frac{1}{N} \sum_{n=0}^{N-1} \cos(4\pi f_0 n + 2\phi) \approx 0.$$

What conditions on f_0 are required for this to hold? Hint: Note that

$$\sum_{n=0}^{N-1} \cos(\alpha n + \beta) = \operatorname{Re} \left(\sum_{n=0}^{N-1} \exp[j(\alpha n + \beta)] \right)$$

and use the geometric progression sum formula.

- 3.8** Repeat the computation of the CRLB for Example 3.3 by using the alternative expression (3.12).
- 3.9** We observe two samples of a DC level in *correlated* Gaussian noise

$$\begin{aligned} x[0] &= A + w[0] \\ x[1] &= A + w[1] \end{aligned}$$

where $\mathbf{w} = [w[0] \ w[1]]^T$ is zero mean with covariance matrix

$$\mathbf{C} = \sigma^2 \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}.$$

The parameter ρ is the correlation coefficient between $w[0]$ and $w[1]$. Compute the CRLB for A and compare it to the case when $w[n]$ is WGN or $\rho = 0$. Also, explain what happens when $\rho \rightarrow \pm 1$. Finally, comment on the additivity property of the Fisher information for nonindependent observations.

3.10 By using (3.23) prove that the Fisher information matrix is positive semidefinite for all θ . In practice, we assume it to be positive definite and hence invertible, although this is not always the case. Consider the data model in Problem 3.3 with the modification that r is unknown. Find the Fisher information matrix for $\theta = [Ar]^T$. Are there any values of θ for which $\mathbf{I}(\theta)$ is not positive definite?

3.11 For a 2×2 Fisher information matrix

$$\mathbf{I}(\theta) = \begin{bmatrix} a & b \\ b & c \end{bmatrix}$$

which is positive definite, show that

$$[\mathbf{I}^{-1}(\theta)]_{11} = \frac{c}{ac - b^2} \geq \frac{1}{a} = \frac{1}{[\mathbf{I}(\theta)]_{11}}.$$

What does this say about estimating a parameter when a second parameter is either known or unknown? When does equality hold and why?

3.12 Prove that

$$[\mathbf{I}^{-1}(\theta)]_{ii} \geq \frac{1}{[\mathbf{I}(\theta)]_{ii}}.$$

This generalizes the result of Problem 3.11. Additionally, it provides another lower bound on the variance, although it is usually not attainable. Under what conditions will the new bound be achieved? Hint: Apply the Cauchy-Schwarz inequality to $\mathbf{e}_i^T \sqrt{\mathbf{I}(\theta)} \sqrt{\mathbf{I}^{-1}(\theta)} \mathbf{e}_i$, where $\mathbf{e}_i$ is the vectors of all zeros except for a 1 as the i th element. The square root of a positive definite matrix $\mathbf{A}$ is defined to be the matrix with the same eigenvectors as $\mathbf{A}$ but whose eigenvalues are the square roots of those of $\mathbf{A}$.

3.13 Consider a generalization of the line fitting problem as described in Example 3.7, termed *polynomial or curve fitting*. The data model is

$$x[n] = \sum_{k=0}^{p-1} A_k n^k + w[n]$$

for $n = 0, 1, \dots, N-1$. As before, $w[n]$ is WGN with variance σ^2 . It is desired to estimate $\{A_0, A_1, \dots, A_{p-1}\}$. Find the Fisher information matrix for this problem.

3.14 For the data model in Example 3.11 consider the estimator $\hat{\sigma}_A^2 = (\hat{A})^2$, where $\hat{A}$ is the sample mean. Assume we observe a given data set in which the realization of the random variable A is the value A_0 . Show that $\hat{A} \rightarrow A_0$ as $N \rightarrow \infty$ by verifying that

$$\begin{aligned} E(\hat{A} | A = A_0) &= A_0 \\ \text{var}(\hat{A} | A = A_0) &= \frac{\sigma^2}{N}. \end{aligned}$$

Hence, $\hat{\sigma}_A^2 \rightarrow A_0^2$ as $N \rightarrow \infty$ for the given realization $A = A_0$. Next find the variance of $\hat{\sigma}_A^2$ as $N \rightarrow \infty$ by determining $\text{var}(A^2)$, where $A \sim \mathcal{N}(0, \sigma_A^2)$, and compare it to the CRLB. Explain why σ_A^2 cannot be estimated without error even for $N \rightarrow \infty$.

- 3.15** Independent bivariate Gaussian samples $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[N-1]\}$ are observed. Each observation is a 2×1 vector which is distributed as $\mathbf{x}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$ and

$$\mathbf{C} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}.$$

Find the CRLB for the correlation coefficient ρ . Hint: Use (3.32).

- 3.16** It is desired to estimate the total power P_0 of a WSS random process, whose PSD is given as

$$P_{xx}(f) = P_0 Q(f)$$

where

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} Q(f) df = 1$$

and $Q(f)$ is known. If N observations are available, find the CRLB for the total power using the exact form (3.32) as well as the asymptotic approximation (3.34) and compare.

- 3.17** If in Example 3.14 the data are observed over the interval $n \leftarrow -M, \dots, 0, \dots, M$, find the Fisher information matrix. What is the CRLB for the sinusoidal parameters? You can use the same approximations as in the example by assuming M is large. Compare the results to that of the example.

- 3.18** Using the results of Example 3.13, determine the best range estimation accuracy of a sonar if

$$s(t) = \begin{cases} 1 - 100|t - 0.01| & 0 \leq t \leq 0.02 \\ 0 & \text{otherwise.} \end{cases}$$

Let $N_0/2 = 10^{-6}$ and $c = 1500$ m/s.

- 3.19** A line array of antennas with $d = \lambda/2$ spacing for a 1 MHz electromagnetic signal is to be designed to yield a 5° standard deviation at $\beta = 90^\circ$. If the SNR η is 0 dB, comment on the feasibility of the requirement. Assume that the line array is to be mounted on the wing of an aircraft. Use $c = 3 \times 10^8$ m/s.
- 3.20** In Example 3.16 we found the CRLB for the filter parameter of an AR(1) process. Assuming σ_u^2 is known, find the asymptotic CRLB for $P_{xx}(f)$ at a given frequency by using (3.16) and (3.44). For $a[1] = -0.9$, $\sigma_u^2 = 1$, and $N = 100$, plot the CRLB versus frequency. Explain your results.

Appendix 3A

Derivation of Scalar Parameter CRLB

In this appendix we derive the CRLB for a scalar parameter $\alpha = g(\theta)$ where the PDF is parameterized by θ . We consider all unbiased estimators $\hat{\alpha}$ or those for which

$$E(\hat{\alpha}) = \alpha = g(\theta)$$

or

$$\int \hat{\alpha} p(\mathbf{x}; \theta) d\mathbf{x} = g(\theta) \quad (3A.1)$$

Before beginning the derivation we first examine the regularity condition

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] = 0, \quad (3A.2)$$

which is assumed to hold. Note that

$$\begin{aligned} \int \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} &= \int \frac{\partial p(\mathbf{x}; \theta)}{\partial \theta} d\mathbf{x} \\ &= \frac{\partial}{\partial \theta} \int p(\mathbf{x}; \theta) d\mathbf{x} \\ &= \frac{\partial 1}{\partial \theta} \\ &= 0. \end{aligned}$$

Hence, we conclude that the regularity condition will be satisfied if the order of differentiation and integration may be interchanged. This is generally true except when the domain of the PDF for which it is nonzero depends on the unknown parameter such as in Problem 3.1.

Now differentiating both sides of (3A.1) with respect to θ and interchanging the partial differentiation and integration produces

$$\int \hat{\alpha} \frac{\partial p(\mathbf{x}; \theta)}{\partial \theta} d\mathbf{x} = \frac{\partial g(\theta)}{\partial \theta}$$

or

$$\int \hat{\alpha} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} = \frac{\partial g(\theta)}{\partial \theta}. \quad (3A.3)$$

We can modify this using the regularity condition to produce

$$\int (\hat{\alpha} - \alpha) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} = \frac{\partial g(\theta)}{\partial \theta} \quad (3A.4)$$

since

$$\int \alpha \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} = \alpha E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] = 0.$$

We now apply the Cauchy-Schwarz inequality

$$\left[\int w(\mathbf{x}) g(\mathbf{x}) h(\mathbf{x}) d\mathbf{x} \right]^2 \leq \int w(\mathbf{x}) g^2(\mathbf{x}) d\mathbf{x} \int w(\mathbf{x}) h^2(\mathbf{x}) d\mathbf{x} \quad (3A.5)$$

which holds with equality if and only if $g(\mathbf{x}) = ch(\mathbf{x})$ for c some constant not dependent on $\mathbf{x}$. The functions g and h are arbitrary scalar functions, while $w(\mathbf{x}) \geq 0$ for all $\mathbf{x}$. Now let

$$\begin{aligned} w(\mathbf{x}) &= p(\mathbf{x}; \theta) \\ g(\mathbf{x}) &= \hat{\alpha} - \alpha \\ h(\mathbf{x}) &= \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \end{aligned}$$

and apply the Cauchy-Schwarz inequality to (3A.4) to produce

$$\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2 \leq \int (\hat{\alpha} - \alpha)^2 p(\mathbf{x}; \theta) d\mathbf{x} \int \left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 p(\mathbf{x}; \theta) d\mathbf{x}$$

or

$$\boxed{\text{var}(\hat{\alpha}) \geq \frac{\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2}{E \left[\left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 \right]}}$$

Now note that

$$E \left[\left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 \right] = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]$$

This follows from (3A.2) as

$$\begin{aligned} E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] &= 0 \\ \int \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} &= 0 \end{aligned}$$

$$\begin{aligned}\frac{\partial}{\partial \theta} \int \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} &= 0 \\ \int \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} p(\mathbf{x}; \theta) + \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \frac{\partial p(\mathbf{x}; \theta)}{\partial \theta} \right] d\mathbf{x} &= 0\end{aligned}$$

or

$$\begin{aligned}-E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right] &= \int \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} p(\mathbf{x}; \theta) d\mathbf{x} \\ &= E \left[\left(\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right)^2 \right].\end{aligned}$$

Thus,

$$\text{var}(\hat{\alpha}) \geq \frac{\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2}{-E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]}$$

which is (3.16). If $\alpha = g(\theta) = \theta$, we have (3.6).Note that the condition for equality is

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \frac{1}{c} (\hat{\alpha} - \alpha)$$

where c can depend on θ but not on $\mathbf{x}$. If $\alpha = g(\theta) = \theta$, we have

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \frac{1}{c(\theta)} (\hat{\theta} - \theta)$$

The possible dependence of c on θ is noted. To determine $c(\theta)$

$$\begin{aligned}\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} &= -\frac{1}{c(\theta)} + \frac{\partial \left(\frac{1}{c(\theta)} \right)}{\partial \theta} (\hat{\theta} - \theta) \\ -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right] &= \frac{1}{c(\theta)}\end{aligned}$$

or finally

$$\begin{aligned}c(\theta) &= \frac{1}{-E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]} \\ &= \frac{1}{I(\theta)}\end{aligned}$$

which agrees with (3.7).

Appendix 3B

Derivation of Vector Parameter CRLB

In this appendix the CRLB for a vector parameter $\alpha = \mathbf{g}(\theta)$ is derived. The PDF is characterized by θ . We consider unbiased estimators such that

$$E(\hat{\alpha}_i) = \alpha_i = [\mathbf{g}(\theta)]_i \quad i = 1, 2, \dots, r.$$

The regularity conditions are

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right] = \mathbf{0},$$

so that

$$\int (\hat{\alpha}_i - \alpha_i) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_i} p(\mathbf{x}; \theta) d\mathbf{x} = \frac{\partial [\mathbf{g}(\theta)]_i}{\partial \theta_i}. \quad (3B.1)$$

Now consider for $j \neq i$

$$\begin{aligned} \int (\hat{\alpha}_i - \alpha_i) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_j} p(\mathbf{x}; \theta) d\mathbf{x} &= \int (\hat{\alpha}_i - \alpha_i) \frac{\partial p(\mathbf{x}; \theta)}{\partial \theta_j} d\mathbf{x} \\ &= \frac{\partial}{\partial \theta_j} \int \hat{\alpha}_i p(\mathbf{x}; \theta) d\mathbf{x} \\ &\quad - \alpha_i E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_j} \right] \\ &= \frac{\partial \alpha_i}{\partial \theta_j} \\ &= \frac{\partial [\mathbf{g}(\theta)]_i}{\partial \theta_j}. \end{aligned} \quad (3B.2)$$

Combining (3B.1) and (3B.2) into matrix form, we have

$$\int (\hat{\alpha} - \alpha) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta}^T p(\mathbf{x}; \theta) d\mathbf{x} = \frac{\partial \mathbf{g}(\theta)}{\partial \theta}^T$$

Now premultiply by $\mathbf{a}^T$ and postmultiply by $\mathbf{b}$, where $\mathbf{a}$ and $\mathbf{b}$ are arbitrary $r \times 1$ and $p \times 1$ vectors, respectively, to yield

$$\int \mathbf{a}^T (\hat{\alpha} - \alpha) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \mathbf{b} p(\mathbf{x}; \theta) d\mathbf{x} = \mathbf{a}^T \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{b}$$

Now let

$$\begin{aligned} w(\mathbf{x}) &= p(\mathbf{x}; \theta) \\ g(\mathbf{x}) &= \mathbf{a}^T (\hat{\alpha} - \alpha) \\ h(\mathbf{x}) &= \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \mathbf{b} \end{aligned}$$

and apply the Cauchy-Schwarz inequality of (3A.5)

$$\begin{aligned} \left(\mathbf{a}^T \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{b} \right)^2 &\leq \int \mathbf{a}^T (\hat{\alpha} - \alpha) (\hat{\alpha} - \alpha)^T \mathbf{a} p(\mathbf{x}; \theta) d\mathbf{x} \\ &\quad \cdot \int \mathbf{b}^T \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \mathbf{b} p(\mathbf{x}; \theta) d\mathbf{x} \\ &= \mathbf{a}^T \mathbf{C}_{\hat{\alpha}} \mathbf{a} \mathbf{b}^T \mathbf{I}(\theta) \mathbf{b} \end{aligned}$$

since as in the scalar case

$$E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_i} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_j} \right] = -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta_i \partial \theta_j} \right] = [\mathbf{I}(\theta)]_{ij}.$$

Since $\mathbf{b}$ was arbitrary, let

$$\mathbf{b} = \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{a}$$

to yield

$$\left(\mathbf{a}^T \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{a} \right)^2 \leq \mathbf{a}^T \mathbf{C}_{\hat{\alpha}} \mathbf{a} \left(\mathbf{a}^T \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{a} \right).$$

Since $\mathbf{I}(\theta)$ is positive definite, so is $\mathbf{I}^{-1}(\theta)$, and $\frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta}^T$ is at least positive semidefinite. The term inside the parentheses is therefore nonnegative, and we have

$$\mathbf{a}^T \left(\mathbf{C}_{\hat{\alpha}} - \frac{\partial \mathbf{g}(\theta)}{\partial \theta} \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta}^T \right) \mathbf{a} \geq 0.$$

Recall that $\mathbf{a}$ was arbitrary, so that (3.30) follows. If $\alpha = \mathbf{g}(\theta) = \theta$, then $\frac{\partial \mathbf{g}(\theta)}{\partial \theta} = \mathbf{I}$ and (3.24) follows. The conditions for equality are $g(\mathbf{x}) = c h(\mathbf{x})$, where c is a constant not dependent on $\mathbf{x}$. This condition becomes

$$\begin{aligned} \mathbf{a}^T (\hat{\alpha} - \alpha) &= c \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \mathbf{b} \\ &= c \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \mathbf{I}^{-1}(\theta) \frac{\partial \mathbf{g}(\theta)}{\partial \theta}^T \mathbf{a}. \end{aligned}$$

Since $\mathbf{a}$ was arbitrary

$$\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{I}^{-1}(\boldsymbol{\theta}) \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \frac{1}{c} (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}).$$

Consider the case when $\boldsymbol{\alpha} = \mathbf{g}(\boldsymbol{\theta}) = \boldsymbol{\theta}$, so that $\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{I}$. Then,

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \frac{1}{c} \mathbf{I}(\boldsymbol{\theta}) (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}).$$

Noting that c may depend on $\boldsymbol{\theta}$, we have

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i} = \sum_{k=1}^p \frac{[\mathbf{I}(\boldsymbol{\theta})]_{ik}}{c(\boldsymbol{\theta})} (\hat{\theta}_k - \theta_k)$$

and differentiating once more

$$\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} = \sum_{k=1}^p \left(\frac{[\mathbf{I}(\boldsymbol{\theta})]_{ik}}{c(\boldsymbol{\theta})} (-\delta_{kj}) + \frac{\partial \left(\frac{[\mathbf{I}(\boldsymbol{\theta})]_{ik}}{c(\boldsymbol{\theta})} \right)}{\partial \theta_j} (\hat{\theta}_k - \theta_k) \right).$$

Finally, we have

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{ij} &= -E \left[\frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right] \\ &= \frac{[\mathbf{I}(\boldsymbol{\theta})]_{ij}}{c(\boldsymbol{\theta})} \end{aligned}$$

since $E(\hat{\theta}_k) = \theta_k$. Clearly, $c(\boldsymbol{\theta}) = 1$ and the condition for equality follows.

Appendix 3C

Derivation of General Gaussian CRLB

Assume that $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\boldsymbol{\theta}), \mathbf{C}(\boldsymbol{\theta}))$, where $\boldsymbol{\mu}(\boldsymbol{\theta})$ is the $N \times 1$ mean vector and $\mathbf{C}(\boldsymbol{\theta})$ is the $N \times N$ covariance matrix, both of which depend on $\boldsymbol{\theta}$. Then the PDF is

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}[\mathbf{C}(\boldsymbol{\theta})]} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \mathbf{C}^{-1}(\boldsymbol{\theta}) (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) \right].$$

We will make use of the following identities

$$\frac{\partial \ln \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k} = \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right) \quad (3C.1)$$

where $\partial \mathbf{C}(\boldsymbol{\theta}) / \partial \theta_k$ is the $N \times N$ matrix with $[i, j]$ element $\partial [\mathbf{C}(\boldsymbol{\theta})]_{ij} / \partial \theta_k$ and

$$\frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_k} = -\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{C}^{-1}(\boldsymbol{\theta}). \quad (3C.2)$$

To establish (3C.1) we first note

$$\frac{\partial \ln \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k} = \frac{1}{\det[\mathbf{C}(\boldsymbol{\theta})]} \frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k}. \quad (3C.3)$$

Since $\det[\mathbf{C}(\boldsymbol{\theta})]$ depends on all the elements of $\mathbf{C}(\boldsymbol{\theta})$

$$\begin{aligned} \frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k} &= \sum_{i=1}^N \sum_{j=1}^N \frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial [\mathbf{C}(\boldsymbol{\theta})]_{ij}} \frac{\partial [\mathbf{C}(\boldsymbol{\theta})]_{ij}}{\partial \theta_k} \\ &= \text{tr} \left(\frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \mathbf{C}(\boldsymbol{\theta})} \frac{\partial \mathbf{C}^T(\boldsymbol{\theta})}{\partial \theta_k} \right) \end{aligned} \quad (3C.4)$$

where $\partial \det[\mathbf{C}(\boldsymbol{\theta})] / \partial \mathbf{C}(\boldsymbol{\theta})$ is an $N \times N$ matrix with $[i, j]$ element $\partial \det[\mathbf{C}(\boldsymbol{\theta})] / \partial [\mathbf{C}(\boldsymbol{\theta})]_{ij}$ and the identity

$$\text{tr}(\mathbf{A}\mathbf{B}^T) = \sum_{i=1}^N \sum_{j=1}^N [\mathbf{A}]_{ij} [\mathbf{B}]_{ij}$$

has been used. Now by the definition of the determinant

$$\det[\mathbf{C}(\boldsymbol{\theta})] = \sum_{i=1}^N [\mathbf{C}(\boldsymbol{\theta})]_{ij} [\mathbf{M}]_{ij}$$

where $\mathbf{M}$ is the $N \times N$ cofactor matrix and j can take on any value from 1 to N . Thus,

$$\frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial [\mathbf{C}(\boldsymbol{\theta})]_{ij}} = [\mathbf{M}]_{ij}$$

or

$$\frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \mathbf{C}(\boldsymbol{\theta})} = \mathbf{M}.$$

It is well known, however, that

$$\mathbf{C}^{-1}(\boldsymbol{\theta}) = \frac{\mathbf{M}^T}{\det[\mathbf{C}(\boldsymbol{\theta})]}$$

so that

$$\frac{\partial \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \mathbf{C}(\boldsymbol{\theta})} = \mathbf{C}^{-1}(\boldsymbol{\theta}) \det[\mathbf{C}(\boldsymbol{\theta})].$$

Using this in (3C.3) and (3C.4), we have the desired result.

$$\begin{aligned} \frac{\partial \ln \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k} &= \frac{1}{\det[\mathbf{C}(\boldsymbol{\theta})]} \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \det[\mathbf{C}(\boldsymbol{\theta})] \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right) \\ &= \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right). \end{aligned}$$

The second identity (3C.2) is easily established as follows. Consider

$$\mathbf{C}^{-1}(\boldsymbol{\theta}) \mathbf{C}(\boldsymbol{\theta}) = \mathbf{I}.$$

Differentiating each element of the matrices and expressing in matrix form, we have

$$\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} + \frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{C}(\boldsymbol{\theta}) = \mathbf{0}$$

which leads to the desired result.

We are now ready to evaluate the CRLB. Taking the first derivative

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_k} = -\frac{1}{2} \frac{\partial \ln \det[\mathbf{C}(\boldsymbol{\theta})]}{\partial \theta_k} - \frac{1}{2} \frac{\partial}{\partial \theta_k} [(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \mathbf{C}^{-1}(\boldsymbol{\theta}) (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))].$$

The first term has already been evaluated using (3C.1), so that we consider the second term:

$$\frac{\partial}{\partial \theta_k} [(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \mathbf{C}^{-1}(\boldsymbol{\theta}) (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))]$$

$$\begin{aligned}
&= \frac{\partial}{\partial \theta_k} \sum_{i=1}^N \sum_{j=1}^N (x[i] - [\mu(\theta)]_i) [C^{-1}(\theta)]_{ij} (x[j] - [\mu(\theta)]_j) \\
&= \sum_{i=1}^N \sum_{j=1}^N \left\{ (x[i] - [\mu(\theta)]_i) \left[[C^{-1}(\theta)]_{ij} \left(-\frac{\partial [\mu(\theta)]_j}{\partial \theta_k} \right) \right. \right. \\
&\quad \left. \left. + \frac{\partial [C^{-1}(\theta)]_{ij}}{\partial \theta_k} (x[j] - [\mu(\theta)]_j) \right] \right. \\
&\quad \left. + \left(-\frac{\partial [\mu(\theta)]_i}{\partial \theta_k} \right) [C^{-1}(\theta)]_{ij} (x[j] - [\mu(\theta)]_j) \right\} \\
&= -(\mathbf{x} - \mu(\theta))^T C^{-1}(\theta) \frac{\partial \mu(\theta)}{\partial \theta_k} + (\mathbf{x} - \mu(\theta))^T \frac{\partial C^{-1}(\theta)}{\partial \theta_k} (\mathbf{x} - \mu(\theta)) \\
&\quad - \frac{\partial \mu(\theta)^T}{\partial \theta_k} C^{-1}(\theta) (\mathbf{x} - \mu(\theta)) \\
&= -2 \frac{\partial \mu(\theta)^T}{\partial \theta_k} C^{-1}(\theta) (\mathbf{x} - \mu(\theta)) + (\mathbf{x} - \mu(\theta))^T \frac{\partial C^{-1}(\theta)}{\partial \theta_k} (\mathbf{x} - \mu(\theta)).
\end{aligned}$$

Using (3C.1) and the last result, we have

$$\boxed{\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_k} &= -\frac{1}{2} \text{tr} \left(C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_k} \right) + \frac{\partial \mu(\theta)^T}{\partial \theta_k} C^{-1}(\theta) (\mathbf{x} - \mu(\theta)) \\ &\quad - \frac{1}{2} (\mathbf{x} - \mu(\theta))^T \frac{\partial C^{-1}(\theta)}{\partial \theta_k} (\mathbf{x} - \mu(\theta)). \end{aligned}} \quad (3C.5)$$

Let $\mathbf{y} = \mathbf{x} - \mu(\theta)$. Evaluating

$$[\mathbf{I}(\theta)]_{kl} = E \left[\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_k} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_l} \right]$$

which is equivalent to (3.23), yields

$$\begin{aligned}
[\mathbf{I}(\theta)]_{kl} &= \frac{1}{4} \text{tr} \left(C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_k} \right) \text{tr} \left(C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_l} \right) \\
&\quad + \frac{1}{2} \text{tr} \left(C^{-1}(\theta) \frac{\partial C(\theta)}{\partial \theta_k} \right) E \left(\mathbf{y}^T \frac{\partial C^{-1}(\theta)}{\partial \theta_l} \mathbf{y} \right) \\
&\quad + \frac{\partial \mu(\theta)^T}{\partial \theta_k} C^{-1}(\theta) E[\mathbf{y} \mathbf{y}^T] C^{-1}(\theta) \frac{\partial \mu(\theta)}{\partial \theta_l} \\
&\quad + \frac{1}{4} E \left[\mathbf{y}^T \frac{\partial C^{-1}(\theta)}{\partial \theta_k} \mathbf{y} \mathbf{y}^T \frac{\partial C^{-1}(\theta)}{\partial \theta_l} \mathbf{y} \right]
\end{aligned}$$

where we note that all odd order moments are zero. Continuing, we have

$$\begin{aligned}
 & [\mathbf{I}(\boldsymbol{\theta})]_{kl} = \\
 & \frac{1}{4} \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right) \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_l} \right) \\
 & - \frac{1}{2} \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right) \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_l} \right) \\
 & + \frac{\partial \boldsymbol{\mu}(\boldsymbol{\theta})^T}{\partial \theta_k} \mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \boldsymbol{\mu}(\boldsymbol{\theta})}{\partial \theta_l} + \frac{1}{4} E \left[\mathbf{y}^T \frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{y} \mathbf{y}^T \frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_l} \mathbf{y} \right] \quad (3C.6)
 \end{aligned}$$

where $E(\mathbf{y}^T \mathbf{z}) = \text{tr}[E(\mathbf{z} \mathbf{y}^T)]$ for $\mathbf{y}, \mathbf{z}$ $N \times 1$ vectors and (3C.2) have been used. To evaluate the last term we use [Porat and Friedlander 1986]

$$E[\mathbf{y}^T \mathbf{A} \mathbf{y} \mathbf{y}^T \mathbf{B} \mathbf{y}] = \text{tr}(\mathbf{A} \mathbf{C}) \text{tr}(\mathbf{B} \mathbf{C}) + 2 \text{tr}(\mathbf{A} \mathbf{C} \mathbf{B} \mathbf{C})$$

where $\mathbf{C} = E(\mathbf{y} \mathbf{y}^T)$ and $\mathbf{A}$ and $\mathbf{B}$ are symmetric matrices. Thus, this term becomes

$$\begin{aligned}
 & \frac{1}{4} \text{tr} \left(\frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{C}(\boldsymbol{\theta}) \right) \text{tr} \left(\frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_l} \mathbf{C}(\boldsymbol{\theta}) \right) \\
 & + \frac{1}{2} \text{tr} \left(\frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{C}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}^{-1}(\boldsymbol{\theta})}{\partial \theta_l} \mathbf{C}(\boldsymbol{\theta}) \right).
 \end{aligned}$$

Next, using the relationship (3C.2), this term becomes

$$\begin{aligned}
 & \frac{1}{4} \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \right) \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_l} \right) \\
 & + \frac{1}{2} \text{tr} \left(\mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{C}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}(\boldsymbol{\theta})}{\partial \theta_l} \right) \quad (3C.7)
 \end{aligned}$$

and finally, using (3C.7) in (3C.6), produces the desired result.

Appendix 3D

Derivation of Asymptotic CRLB

It can be proven that almost any WSS Gaussian random process $x[n]$ may be represented as the output of a causal linear shift invariant filter driven at the input by white Gaussian noise $u[n]$ [Brockwell and Davis 1987] or

$$x[n] = \sum_{k=0}^{\infty} h[k]u[n-k] \quad (3D.1)$$

where $h[0] = 1$. The only condition is that the PSD must satisfy

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln P_{xx}(f) df > -\infty.$$

With this representation the PSD of $x[n]$ is

$$P_{xx}(f) = |H(f)|^2 \sigma_u^2$$

where σ_u^2 is the variance of $u[n]$ and $H(f) = \sum_{k=0}^{\infty} h[k] \exp(-j2\pi f k)$ is the filter frequency response. If the observations are $\{x[0], x[1], \dots, x[N-1]\}$ and N is large, then the representation is approximated by

$$\begin{aligned} x[n] &= \sum_{k=0}^n h[k]u[n-k] + \sum_{k=n+1}^{\infty} h[k]u[n-k] \\ &\approx \sum_{k=0}^n h[k]u[n-k] \end{aligned} \quad (3D.2)$$

This is equivalent to setting $u[n] = 0$ for $n < 0$. As $n \rightarrow \infty$, the approximate representation becomes better for $x[n]$. It is clear, however, that the beginning samples will be poorly represented unless the impulse response $h[k]$ is small for $k > n$. For large N most of the samples will be accurately represented if N is much greater than the impulse response length. Since

$$r_{xx}[k] = \sigma_u^2 \sum_{n=0}^{\infty} h[n]h[n+k]$$

the correlation time or effective duration of $r_{xx}[k]$ is the same as the impulse response length. Hence, because the CRLB to be derived is based on (3D.2), the asymptotic CRLB will be a good approximation if the data record length is much greater than the correlation time.

To find the PDF of $\mathbf{x}$ we use (3D.2), which is a transformation from $\mathbf{u} = [u[0] u[1] \dots u[N-1]]^T$ to $\mathbf{x} = [x[0] x[1] \dots x[N-1]]^T$ or $\mathbf{x} = \mathbf{H} \cdot \mathbf{u}$

$$\mathbf{x} = \begin{bmatrix} h[0] & 0 & 0 & \dots & 0 \\ h[1] & h[0] & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ h[N-1] & h[N-2] & h[N-3] & \dots & h[0] \end{bmatrix} \mathbf{u}.$$

$\mathbf{H}$

Note that $\mathbf{H}$ has a determinant of $(h[0])^N = 1$ and hence is invertible. Since $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \sigma_u^2 \mathbf{I})$, the PDF of $\mathbf{x}$ is $\mathcal{N}(\mathbf{0}, \sigma_u^2 \mathbf{H} \mathbf{H}^T)$ or

$$p(\mathbf{x}; \theta) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\sigma_u^2 \mathbf{H} \mathbf{H}^T)} \exp \left[-\frac{1}{2} \mathbf{x}^T (\sigma_u^2 \mathbf{H} \mathbf{H}^T)^{-1} \mathbf{x} \right],$$

But

$$\det(\sigma_u^2 \mathbf{H} \mathbf{H}^T) = \sigma_u^{2N} \det^2(\mathbf{H}) = \sigma_u^{2N}.$$

Also,

$$\mathbf{x}^T (\sigma_u^2 \mathbf{H} \mathbf{H}^T)^{-1} \mathbf{x} = \frac{1}{\sigma_u^2} (\mathbf{H}^{-1} \mathbf{x})^T (\mathbf{H}^{-1} \mathbf{x}) = \frac{1}{\sigma_u^2} \mathbf{u}^T \mathbf{u}$$

so that

$$p(\mathbf{x}; \theta) = \frac{1}{(2\pi\sigma_u^2)^{\frac{N}{2}}} \exp \left(-\frac{1}{2\sigma_u^2} \mathbf{u}^T \mathbf{u} \right). \quad (3D.3)$$

From (3D.2) we have approximately

$$X(f) = H(f)U(f)$$

where

$$\begin{aligned} X(f) &= \sum_{n=0}^{N-1} x[n] \exp(-j2\pi fn) \\ U(f) &= \sum_{n=0}^{N-1} u[n] \exp(-j2\pi fn) \end{aligned}$$

are the Fourier transforms of the truncated sequences. By Parseval's theorem

$$\frac{1}{\sigma_u^2} \mathbf{u}^T \mathbf{u} = \frac{1}{\sigma_u^2} \sum_{n=0}^{N-1} u^2[n]$$

$$\begin{aligned}
&= \frac{1}{\sigma_u^2} \int_{-\frac{1}{2}}^{\frac{1}{2}} |U(f)|^2 df \\
&\approx \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{|X(f)|^2}{\sigma_u^2 |H(f)|^2} df \\
&= \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{|X(f)|^2}{P_{xx}(f)} df. \quad (3D.4)
\end{aligned}$$

Also,

$$\begin{aligned}
\ln \sigma_u^2 &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln \sigma_u^2 df \\
&= \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln \left(\frac{P_{xx}(f)}{|H(f)|^2} \right) df \\
&= \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln P_{xx}(f) df - \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |H(f)|^2 df.
\end{aligned}$$

But

$$\begin{aligned}
\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |H(f)|^2 df &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln H(f) + \ln H^*(f) df \\
&= 2 \operatorname{Re} \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln H(f) df \\
&= 2 \operatorname{Re} \oint_C \ln \mathcal{H}(z) \frac{dz}{2\pi j z} \\
&= 2 \operatorname{Re} [\mathcal{Z}^{-1} \{\ln \mathcal{H}(z)\}|_{n=0}]
\end{aligned}$$

where C is the unit circle in the z plane. Since $\mathcal{H}(z)$ corresponds to the system function of a causal filter, it converges outside a circle of radius $r < 1$ (since $\mathcal{H}(z)$ is assumed to exist on the unit circle for the frequency response to exist). Hence, $\ln \mathcal{H}(z)$ also converges outside a circle of radius $r < 1$, so that the corresponding sequence is causal. By the initial value theorem which is valid for a causal sequence

$$\begin{aligned}
\mathcal{Z}^{-1} \{\ln \mathcal{H}(z)\}|_{n=0} &\triangleq \lim_{z \rightarrow \infty} \ln \mathcal{H}(z) \\
&= \ln \lim_{z \rightarrow \infty} \mathcal{H}(z) \\
&= \ln h[0] = 0.
\end{aligned}$$

Therefore,

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |H(f)|^2 df = 0$$

and finally

$$\ln \sigma_u^2 = \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln P_{xx}(f) df. \quad (3D.5)$$

Substituting (3D.4) and (3D.5) into (3D.3) produces for the log PDF

$$\ln p(\mathbf{x}; \boldsymbol{\theta}) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \ln P_{xx}(f) df - \frac{1}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{|X(f)|^2}{P_{xx}(f)} df.$$

Hence, the asymptotic log PDF is

$$\ln p(\mathbf{x}; \boldsymbol{\theta}) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\ln P_{xx}(f) + \frac{\frac{1}{N}|X(f)|^2}{P_{xx}(f)} \right] df. \quad (3D.6)$$

To determine the CRLB

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i} &= -\frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{1}{P_{xx}(f)} - \frac{\frac{1}{N}|X(f)|^2}{P_{xx}^2(f)} \right) \frac{\partial P_{xx}(f)}{\partial \theta_i} df \\ \frac{\partial^2 \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} &= -\frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\frac{1}{P_{xx}(f)} - \frac{\frac{1}{N}|X(f)|^2}{P_{xx}^2(f)} \right) \frac{\partial^2 P_{xx}(f)}{\partial \theta_i \partial \theta_j} \\ &\quad + \left(-\frac{1}{P_{xx}^2(f)} + \frac{\frac{2}{N}|X(f)|^2}{P_{xx}^3(f)} \right) \frac{\partial P_{xx}(f)}{\partial \theta_i} \frac{\partial P_{xx}(f)}{\partial \theta_j} df. \end{aligned} \quad (3D.7)$$

In taking the expected value we encounter the term $E(|X(f)|^2/N)$. For large N this is now shown to be $P_{xx}(f)$. Note that $|X(f)|^2/N$ is termed the periodogram spectral estimator.

$$\begin{aligned} E\left(\frac{1}{N}|X(f)|^2\right) &= E\left(\frac{1}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m]x[n] \exp[-j2\pi f(m-n)]\right) \\ &= \frac{1}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} r_{xx}[m-n] \exp[-j2\pi f(m-n)] \\ &= \sum_{k=-(N-1)}^{N-1} \left(1 - \frac{|k|}{N}\right) r_{xx}[k] \exp(-j2\pi fk) \end{aligned} \quad (3D.8)$$

where we have used the identity

$$\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} g[m-n] = \sum_{k=-(N-1)}^{N-1} (N - |k|)g[k].$$

As $N \rightarrow \infty$,

$$\left(1 - \frac{|k|}{N}\right) r_{xx}[k] \rightarrow r_{xx}[k]$$

assuming that the ACF dies out sufficiently rapidly. Hence,

$$E \left[\frac{1}{N} |X(f)|^2 \right] \approx P_{xx}(f).$$

Upon taking expectations in (3D.7), the first term is zero, and finally,

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\theta})]_{ij} &= \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{P_{xx}^2(f)} \frac{\partial P_{xx}(f)}{\partial \theta_i} \frac{\partial P_{xx}(f)}{\partial \theta_j} df \\ &= \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{\partial \ln P_{xx}(f)}{\partial \theta_i} \frac{\partial \ln P_{xx}(f)}{\partial \theta_j} df \end{aligned}$$

which is (3.34) without the explicit dependence of the PSD on $\boldsymbol{\theta}$ shown.

Chapter 4

Linear Models

4.1 Introduction

The determination of the MVU estimator is *in general* a difficult task. It is fortunate, however, that a large number of signal processing estimation problems can be represented by a data model that allows us to easily determine this estimator. This class of models is the **linear model**. Not only is the MVU estimator immediately evident once the linear model has been identified, but in addition, the statistical performance follows naturally. The key, then, to finding the optimal estimator is in structuring the problem in the linear model form to take advantage of its unique properties.

4.2 Summary

The linear model is defined by (4.8). When this data model can be assumed, the MVU (and also efficient) estimator is given by (4.9), and the covariance matrix by (4.10). A more general model, termed the general linear model, allows the noise to have an arbitrary covariance matrix, as opposed to $\sigma^2 \mathbf{I}$ for the linear model. The MVU (and also efficient) estimator for this model is given by (4.25), and its corresponding covariance matrix by (4.26). A final extension allows for known signal components in the data to yield the MVU (and also efficient) estimator of (4.31). The covariance matrix is the same as for the general linear model.

4.3 Definition and Properties

The linear model has already been encountered in the line fitting problem discussed in Example 3.7. Recall that the problem was to fit a straight line through noise corrupted data. As our model of the data we chose

$$x[n] = A + Bn + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN and the slope B and intercept A were to be estimated. In matrix notation the model is written more compactly as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \quad (4.1)$$

where

$$\begin{aligned} \mathbf{x} &= [x[0] \ x[1] \ \dots \ x[N-1]]^T \\ \mathbf{w} &= [w[0] \ w[1] \ \dots \ w[N-1]]^T \\ \boldsymbol{\theta} &= [A \ B]^T \end{aligned}$$

and

$$\mathbf{H} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ \vdots & \vdots \\ 1 & N-1 \end{bmatrix}$$

The matrix $\mathbf{H}$ is a known matrix of dimension $N \times 2$ and is referred to as the *observation matrix*. The data $\mathbf{x}$ are observed after $\boldsymbol{\theta}$ is operated upon by $\mathbf{H}$. Note also that the noise vector has the statistical characterization $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. The data model in (4.1) is termed the *linear model*. In defining the linear model we assume that the noise vector is Gaussian, although other authors use the term more generally for any noise PDF [Graybill 1976].

As discussed in Chapter 3, it is sometimes possible to determine the MVU estimator if the equality constraints of the CRLB theorem are satisfied. From Theorem 3.2, $\hat{\boldsymbol{\theta}} = \mathbf{g}(\mathbf{x})$ will be the MVU estimator if

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{I}(\boldsymbol{\theta})(\mathbf{g}(\mathbf{x}) - \boldsymbol{\theta}) \quad (4.2)$$

for some function $\mathbf{g}$. Furthermore, the covariance matrix of $\hat{\boldsymbol{\theta}}$ will be $\mathbf{I}^{-1}(\boldsymbol{\theta})$. To determine if this condition is satisfied for the linear model of (4.1), we have

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} &= \frac{\partial}{\partial \boldsymbol{\theta}} \left[-\ln(2\pi\sigma^2)^{\frac{N}{2}} - \frac{1}{2\sigma^2}(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T(\mathbf{x} - \mathbf{H}\boldsymbol{\theta}) \right] \\ &= -\frac{1}{2\sigma^2} \frac{\partial}{\partial \boldsymbol{\theta}} [\mathbf{x}^T \mathbf{x} - 2\mathbf{x}^T \mathbf{H}\boldsymbol{\theta} + \boldsymbol{\theta}^T \mathbf{H}^T \mathbf{H}\boldsymbol{\theta}] \end{aligned}$$

Using the identities

$$\begin{aligned} \frac{\partial \mathbf{b}^T \boldsymbol{\theta}}{\partial \boldsymbol{\theta}} &= \mathbf{b} \\ \frac{\partial \boldsymbol{\theta}^T \mathbf{A} \boldsymbol{\theta}}{\partial \boldsymbol{\theta}} &= 2\mathbf{A}\boldsymbol{\theta} \end{aligned} \quad (4.3)$$

for $\mathbf{A}$ a symmetric matrix, we have

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \frac{1}{\sigma^2} [\mathbf{H}^T \mathbf{x} - \mathbf{H}^T \mathbf{H}\boldsymbol{\theta}]$$

Assuming that $\mathbf{H}^T \mathbf{H}$ is invertible

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \frac{\mathbf{H}^T \mathbf{H}}{\sigma^2} [(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} - \boldsymbol{\theta}], \quad (4.4)$$

which is exactly in the form of (4.2) with

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \quad (4.5)$$

$$\mathbf{I}(\boldsymbol{\theta}) = \frac{\mathbf{H}^T \mathbf{H}}{\sigma^2} \quad (4.6)$$

Hence, the MVU estimator of $\boldsymbol{\theta}$ is given by (4.5), and its covariance matrix is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \mathbf{I}^{-1}(\boldsymbol{\theta}) = \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1} \quad (4.7)$$

Additionally, the MVU estimator for the linear model is efficient in that it attains the CRLB. The reader may now verify the result of (3.29) by substituting $\mathbf{H}$ for the line fitting problem into (4.4). The only detail that requires closer scrutiny is the invertibility of $\mathbf{H}^T \mathbf{H}$. For the line fitting example a direct calculation will verify that the inverse exists (compute the determinant of the matrix given in (3.29)). Alternatively, this follows from the linear independence of the columns of $\mathbf{H}$ (see Problem 4.2). If the columns of $\mathbf{H}$ are not linearly independent, as for example,

$$\mathbf{H} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \\ \vdots & \vdots \\ 1 & 1 \end{bmatrix}$$

and $\mathbf{x} = [2 \ 2 \ \dots \ 2]^T$ so that $\mathbf{x}$ lies in the range space of $\mathbf{H}$, then even in the absence of noise the model parameters will not be identifiable. For then

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta}$$

and for this choice of $\mathbf{H}$ we will have for $x[n]$

$$2 = A + B \quad n = 0, 1, \dots, N-1.$$

As illustrated in Figure 4.1 it is clear that an infinite number of choices can be made for A and B that will result in the same observations or given a noiseless $\mathbf{x}$, $\boldsymbol{\theta}$ is not unique. The situation can hardly hope to improve when the observations are corrupted by noise. Although rarely occurring in practice, this degeneracy sometimes nearly occurs when $\mathbf{H}^T \mathbf{H}$ is ill-conditioned (see Problem 4.3).

The previous discussion, although illustrated by the line fitting example, is completely general, as summarized by the following theorem.

Theorem 4.1 (Minimum Variance Unbiased Estimator for the Linear Model)
If the data observed can be modeled as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \quad (4.8)$$




Figure 4.1 Nonidentifiability of linear model parameters

where $\mathbf{x}$ is an $N \times 1$ vector of observations, $\mathbf{H}$ is a known $N \times p$ observation matrix (with $N > p$) and rank p , $\boldsymbol{\theta}$ is a $p \times 1$ vector of parameters to be estimated, and $\mathbf{w}$ is an $N \times 1$ noise vector with PDF $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, then the MVU estimator is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \quad (4.9)$$

and the covariance matrix of $\hat{\boldsymbol{\theta}}$ is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1}. \quad (4.10)$$

For the linear model the MVU estimator is efficient in that it attains the CRLB.

That $\hat{\boldsymbol{\theta}}$ is unbiased easily follows by substituting (4.8) into (4.9). Also, the statistical performance of $\hat{\boldsymbol{\theta}}$ is completely specified (not just the mean and covariance) because $\hat{\boldsymbol{\theta}}$ is a linear transformation of a Gaussian vector $\mathbf{x}$ and hence

$$\hat{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\theta}, \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1}). \quad (4.11)$$

The Gaussian nature of the MVU estimator for the linear model allows us to determine the exact statistical performance if desired (see Problem 4.4). In the next section we present some examples illustrating the use of the linear model.

4.4 Linear Model Examples

We have already seen how the problem of line fitting is easily handled once we recognize it as a linear model. A simple extension is to the problem of fitting a curve to experimental data.

Example 4.1 - Curve Fitting

In many experimental situations we seek to determine an empirical relationship between a pair of variables. For instance, in Figure 4.2 we present the results of a experiment




Figure 4.2 Experimental data

in which voltage measurements are taken at the time instants $t = t_0, t_1, \dots, t_{N-1}$. By plotting the measurements it appears that the underlying voltage may be a quadratic function of time. That the points do not lie exactly on a curve is attributed to experimental error or noise. Hence, a reasonable model for the data is

$$x(t_n) = \theta_1 + \theta_2 t_n + \theta_3 t_n^2 + w(t_n) \quad n = 0, 1, \dots, N-1.$$

To avail ourselves of the utility of the linear model we assume that $w(t_n)$ are IID Gaussian random variables with zero mean and variance σ^2 or that they are WGN samples. Then, we have the usual linear model form

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where

$$\begin{aligned} \mathbf{x} &= [x(t_0) \ x(t_1) \ \dots \ x(t_{N-1})]^T \\ \boldsymbol{\theta} &= [\theta_1 \ \theta_2 \ \theta_3]^T \end{aligned}$$

and

$$\mathbf{H} = \begin{bmatrix} 1 & t_0 & t_0^2 \\ 1 & t_1 & t_1^2 \\ \vdots & \vdots & \vdots \\ 1 & t_{N-1} & t_{N-1}^2 \end{bmatrix}.$$

In general, if we seek to fit a $(p-1)$ st-order polynomial to experimental data we will have

$$x(t_n) = \theta_1 + \theta_2 t_n + \dots + \theta_p t_n^{p-1} + w(t_n) \quad n = 0, 1, \dots, N-1.$$

The MVU estimator for $\boldsymbol{\theta} = [\theta_1 \ \theta_2 \ \dots \ \theta_p]^T$ follows from (4.9) as

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$$

where

$$\mathbf{x} = [x(t_0) \ x(t_1) \ \dots \ x(t_{N-1})]^T$$

$$\mathbf{H} = \begin{bmatrix} 1 & t_0 & \dots & t_0^{p-1} \\ 1 & t_1 & \dots & t_1^{p-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & t_{N-1} & \dots & t_{N-1}^{p-1} \end{bmatrix} \quad (N \times p).$$

The observation matrix for this example has the special form of a Vandermonde matrix. Note that the resultant curve fit is

$$\hat{s}(t) = \sum_{i=1}^p \hat{\theta}_i t^{i-1}$$

where $s(t)$ denotes the underlying curve or signal. ◇

Example 4.2 - Fourier Analysis

Many signals exhibit cyclical behavior. It is common practice to determine the presence of strong cyclical components by employing a Fourier analysis. Large Fourier coefficients are indicative of strong components. In this example we show that a Fourier analysis is really just an estimation of the linear model parameters. Consider a data model consisting of sinusoids in white Gaussian noise:

$$x[n] = \sum_{k=1}^M a_k \cos\left(\frac{2\pi k n}{N}\right) + \sum_{k=1}^M b_k \sin\left(\frac{2\pi k n}{N}\right) + w[n] \quad n = 0, 1, \dots, N-1 \quad (4.12)$$

where $w[n]$ is WGN. The frequencies are assumed to be harmonically related or multiples of the fundamental $f_1 = 1/N$ as $f_k = k/N$. The amplitudes a_k, b_k of the cosines and sines are to be estimated. To reformulate the problem in terms of the linear model we let

$$\boldsymbol{\theta} = [a_1 \ a_2 \ \dots \ a_M \ b_1 \ b_2 \ \dots \ b_M]^T$$

and

$$\mathbf{H} = \begin{bmatrix} 1 & \dots & 1 & 0 & \dots & 0 \\ \cos\left(\frac{2\pi}{N}\right) & \dots & \cos\left(\frac{2\pi M}{N}\right) & \sin\left(\frac{2\pi}{N}\right) & \dots & \sin\left(\frac{2\pi M}{N}\right) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \cos\left[\frac{2\pi(N-1)}{N}\right] & \dots & \cos\left[\frac{2\pi M(N-1)}{N}\right] & \sin\left[\frac{2\pi(N-1)}{N}\right] & \dots & \sin\left[\frac{2\pi M(N-1)}{N}\right] \end{bmatrix}.$$

Note that $\mathbf{H}$ has dimension $N \times 2M$, where $p = 2M$. Hence, for $\mathbf{H}$ to satisfy $N > p$ we require $M < N/2$. In determining the MVU estimator we can simplify the computations by noting that the columns of $\mathbf{H}$ are orthogonal. Let $\mathbf{H}$ be represented in column form as

$$\mathbf{H} = [\mathbf{h}_1 \mathbf{h}_2 \dots \mathbf{h}_{2M}]$$

where $\mathbf{h}_i$ denotes the i th column of $\mathbf{H}$. Then, it follows that

$$\mathbf{h}_i^T \mathbf{h}_j = 0 \quad \text{for } i \neq j.$$

This property is quite useful in that

$$\begin{aligned} \mathbf{H}^T \mathbf{H} &= \begin{bmatrix} \mathbf{h}_1^T \\ \vdots \\ \mathbf{h}_{2M}^T \end{bmatrix} \begin{bmatrix} \mathbf{h}_1 & \dots & \mathbf{h}_{2M} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{h}_1^T \mathbf{h}_1 & \mathbf{h}_1^T \mathbf{h}_2 & \dots & \mathbf{h}_1^T \mathbf{h}_{2M} \\ \mathbf{h}_2^T \mathbf{h}_1 & \mathbf{h}_2^T \mathbf{h}_2 & \dots & \mathbf{h}_2^T \mathbf{h}_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{h}_{2M}^T \mathbf{h}_1 & \mathbf{h}_{2M}^T \mathbf{h}_2 & \dots & \mathbf{h}_{2M}^T \mathbf{h}_{2M} \end{bmatrix} \end{aligned}$$

becomes a diagonal matrix which is easily inverted. The orthogonality of the columns results from the discrete Fourier transform (DFT) relationships for $i, j = 1, 2, \dots, M < N/2$:

$$\begin{aligned} \sum_{n=0}^{N-1} \cos\left(\frac{2\pi i n}{N}\right) \cos\left(\frac{2\pi j n}{N}\right) &= \frac{N}{2} \delta_{ij} \\ \sum_{n=0}^{N-1} \sin\left(\frac{2\pi i n}{N}\right) \sin\left(\frac{2\pi j n}{N}\right) &= \frac{N}{2} \delta_{ij} \\ \sum_{n=0}^{N-1} \cos\left(\frac{2\pi i n}{N}\right) \sin\left(\frac{2\pi j n}{N}\right) &= 0 \quad \text{for all } i, j. \end{aligned} \quad (4.13)$$

An outline of the orthogonality proof is given in Problem 4.5. Using this property, we have

$$\mathbf{H}^T \mathbf{H} = \begin{bmatrix} \frac{N}{2} & 0 & \dots & 0 \\ 0 & \frac{N}{2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{N}{2} \end{bmatrix} = \frac{N}{2} \mathbf{I}$$

so that the MVU estimator of the amplitudes is

$$\begin{aligned}
 \hat{\boldsymbol{\theta}} &= (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \\
 &= \frac{2}{N} \mathbf{H}^T \mathbf{x} = \frac{2}{N} \begin{bmatrix} \mathbf{h}_1^T \\ \vdots \\ \mathbf{h}_{2M}^T \end{bmatrix} \mathbf{x} \\
 &= \begin{bmatrix} \frac{2}{N} \mathbf{h}_1^T \mathbf{x} \\ \vdots \\ \frac{2}{N} \mathbf{h}_{2M}^T \mathbf{x} \end{bmatrix}
 \end{aligned}$$

or finally,

$$\begin{aligned}
 \hat{a}_k &= \frac{2}{N} \sum_{n=0}^{N-1} x[n] \cos\left(\frac{2\pi kn}{N}\right) \\
 \hat{b}_k &= \frac{2}{N} \sum_{n=0}^{N-1} x[n] \sin\left(\frac{2\pi kn}{N}\right).
 \end{aligned} \quad (4.14)$$

These are recognized as the discrete Fourier transform coefficients. From the properties of the linear model we can immediately conclude that the means are

$$\begin{aligned}
 E(\hat{a}_k) &= a_k \\
 E(\hat{b}_k) &= b_k
 \end{aligned} \quad (4.15)$$

and the covariance matrix is


$$\begin{aligned}
 \mathbf{C}_{\hat{\boldsymbol{\theta}}} &= \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1} \\
 &= \sigma^2 \left(\frac{N}{2} \mathbf{I} \right)^{-1} \\
 &= \frac{2\sigma^2}{N} \mathbf{I}
 \end{aligned} \quad (4.16)$$

Because $\hat{\boldsymbol{\theta}}$ is Gaussian and the covariance matrix is diagonal, the amplitude estimates are independent (see Problem 4.6 for an application to sinusoidal detection).


It is seen from this example that a key ingredient in simplifying the computation of the MVU estimator and its covariance matrix is the orthogonality of the columns of $\mathbf{H}$. Note that this property *does not hold* if the frequencies are arbitrarily chosen. $\diamond$

Example 4.3 - System Identification

It is frequently of interest to be able to identify the model of a system from input/output data. A common model is the tapped delay line (TDL) or finite impulse response (FIR) filter shown in Figure 4.3a. The input $u[n]$ is known and is provided to "probe" the system. Ideally, at the output the sequence $\sum_{k=0}^{P-1} h[k]u[n-k]$ is observed from which



(a) Tapped delay line



(b) Model for noise-corrupted output data

Figure 4.3 System identification model

we would like to estimate the TDL weights $h[k]$, or equivalently, the impulse response of the FIR filter. In practice, however, the output is corrupted by noise, so that the model in Figure 4.3b is more appropriate. Assume that $u[n]$ is provided for $n = 0, 1, \dots, N-1$ and that the output is observed over the same interval. We then have

$$\mathbf{x}[n] = \sum_{k=0}^{p-1} h[k]u[n-k] + w[n] \quad n = 0, 1, \dots, N-1 \quad (4.17)$$

where it is assumed that $u[n] = 0$ for $n < 0$. In matrix form we have

$$\mathbf{x} = \underbrace{\begin{bmatrix} u[0] & 0 & \cdots & 0 \\ u[1] & u[0] & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u[N-1] & u[N-2] & \cdots & u[N-p] \end{bmatrix}}_{\mathbf{H}} \underbrace{\begin{bmatrix} h[0] \\ h[1] \\ \vdots \\ h[p-1] \end{bmatrix}}_{\boldsymbol{\theta}} + \mathbf{w}. \quad (4.18)$$

Assuming $w[n]$ is WGN, (4.18) is in the form of the linear model, and so the MVU estimator of the impulse response is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

The covariance matrix is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1}.$$

A key question in system identification is how to choose the probing signal $u[n]$. It is now shown that the signal should be chosen to be pseudorandom noise (PRN) [MacWilliams and Sloane 1976]. Since the variance of $\hat{\theta}_i$ is

$$\text{var}(\hat{\theta}_i) = \mathbf{e}_i^T \mathbf{C}_{\hat{\boldsymbol{\theta}}} \mathbf{e}_i$$

where $\mathbf{e}_i = [00\dots 010\dots 0]^T$ with the 1 occupying the i th place, and $\mathbf{C}_{\hat{\boldsymbol{\theta}}}^{-1}$ can be factored as $\mathbf{D}^T \mathbf{D}$ with $\mathbf{D}$ an invertible $p \times p$ matrix, we can use the Cauchy-Schwarz inequality as follows. Noting that

$$1 = (\mathbf{e}_i^T \mathbf{D}^T \mathbf{D}^{T^{-1}} \mathbf{e}_i)^2$$

we can let $\boldsymbol{\xi}_1 = \mathbf{D} \mathbf{e}_i$ and $\boldsymbol{\xi}_2 = \mathbf{D}^{T^{-1}} \mathbf{e}_i$ to yield the inequality

$$(\boldsymbol{\xi}_1^T \boldsymbol{\xi}_2)^2 \leq \boldsymbol{\xi}_1^T \boldsymbol{\xi}_1 \boldsymbol{\xi}_2^T \boldsymbol{\xi}_2.$$

Because $\boldsymbol{\xi}_1^T \boldsymbol{\xi}_2 = 1$, we have

$$\begin{aligned} 1 &\leq (\mathbf{e}_i^T \mathbf{D}^T \mathbf{D} \mathbf{e}_i)(\mathbf{e}_i^T \mathbf{D}^{-1} \mathbf{D}^{T^{-1}} \mathbf{e}_i) \\ &= (\mathbf{e}_i^T \mathbf{C}_{\hat{\boldsymbol{\theta}}}^{-1} \mathbf{e}_i)(\mathbf{e}_i^T \mathbf{C}_{\hat{\boldsymbol{\theta}}} \mathbf{e}_i) \end{aligned}$$

or finally

$$\text{var}(\hat{\theta}_i) \geq \frac{1}{\mathbf{e}_i^T \mathbf{C}_{\hat{\boldsymbol{\theta}}}^{-1} \mathbf{e}_i} = \frac{\sigma^2}{[\mathbf{H}^T \mathbf{H}]_{ii}}.$$

Equality holds or the minimum variance is attained if and only if $\boldsymbol{\xi}_1 = c \boldsymbol{\xi}_2$ for c a constant or

$$\mathbf{D} \mathbf{e}_i = c_i \mathbf{D}^{T^{-1}} \mathbf{e}_i$$

or, equivalently, the conditions for all the variances to be minimized are

$$\mathbf{D}^T \mathbf{D} \mathbf{e}_i = c_i \mathbf{e}_i \quad i = 1, 2, \dots, p.$$

Noting that

$$\mathbf{D}^T \mathbf{D} = \mathbf{C}_{\hat{\boldsymbol{\theta}}}^{-1} = \frac{\mathbf{H}^T \mathbf{H}}{\sigma^2}$$

we have

$$\frac{\mathbf{H}^T \mathbf{H}}{\sigma^2} \mathbf{e}_i = c_i \mathbf{e}_i.$$

If we combine these equations in matrix form, then the conditions for achieving the minimum possible variances are

$$\mathbf{H}^T \mathbf{H} = \sigma^2 \begin{bmatrix} c_1 & 0 & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & c_p \end{bmatrix}.$$

It is now clear that to minimize the variance of the MVU estimator, $u[n]$ should be chosen to make $\mathbf{H}^T \mathbf{H}$ diagonal. Since $[\mathbf{H}]_{ij} = u[i - j]$

$$[\mathbf{H}^T \mathbf{H}]_{ij} = \sum_{n=1}^N u[n - i] u[n - j] \quad i = 1, 2, \dots, p; j = 1, 2, \dots, p \quad (4.19)$$

and for N large we have (see Problem 4.7)

$$[\mathbf{H}^T \mathbf{H}]_{ij} \approx \sum_{n=0}^{N-1-|i-j|} u[n] u[n + |i - j|] \quad (4.20)$$

which can be recognized as a correlation function of the deterministic sequence $u[n]$. Also, with this approximation $\mathbf{H}^T \mathbf{H}$ becomes a symmetric Toeplitz autocorrelation matrix

$$\mathbf{H}^T \mathbf{H} = N \begin{bmatrix} r_{uu}[0] & r_{uu}[1] & \dots & r_{uu}[p-1] \\ r_{uu}[1] & r_{uu}[0] & \dots & r_{uu}[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ r_{uu}[p-1] & r_{uu}[p-2] & \dots & r_{uu}[0] \end{bmatrix}$$

where

$$r_{uu}[k] = \frac{1}{N} \sum_{n=0}^{N-1-k} u[n] u[n+k]$$

may be viewed as an autocorrelation function of $u[n]$. For $\mathbf{H}^T \mathbf{H}$ to be diagonal we require

$$r_{uu}[k] = 0 \quad k \neq 0$$

which is approximately realized if we use a PRN sequence as our input signal. Finally, under these conditions $\mathbf{H}^T \mathbf{H} = N r_{uu}[0] \mathbf{I}$, and hence

$$\text{var}(\hat{h}[i]) = \frac{1}{N r_{uu}[0] / \sigma^2} \quad i = 0, 1, \dots, p-1 \quad (4.21)$$

and the TDL weight estimators are independent.

As a final consequence of choosing a PRN sequence, we obtain as our MVU estimator

$$\hat{\theta} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$$

where $\mathbf{H}^T \mathbf{H} = N r_{uu}[0] \mathbf{I}$. Hence, we have

$$\begin{aligned} \hat{h}[i] &= \frac{1}{N r_{uu}[0]} \sum_{n=0}^{N-1} u[n-i] x[n] \\ &= \frac{\frac{1}{N} \sum_{n=0}^{N-1-i} u[n] x[n+i]}{r_{uu}[0]} \end{aligned} \quad (4.22)$$

since $u[n] = 0$ for $n < 0$. The numerator in (4.22) is just the crosscorrelation $r_{ux}[i]$ between the input and output sequences. Hence, if a PRN input is used to identify the system, then the approximate (for large N) MVU estimator is

$$\hat{h}[i] = \frac{r_{ux}[i]}{r_{uu}[0]} \quad i = 0, 1, \dots, p-1 \quad (4.23)$$

where

$$\begin{aligned} r_{ux}[i] &= \frac{1}{N} \sum_{n=0}^{N-1-i} u[n] x[n+i] \\ r_{uu}[0] &= \frac{1}{N} \sum_{n=0}^{N-1} u^2[n]. \end{aligned}$$

See also Problem 4.8 for a spectral interpretation of the system identification problem.

4.5 Extension to the Linear Model

A more general form of the linear model allows for noise that is not white. The general linear model assumes that

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$$

where $\mathbf{C}$ is not necessarily a scaled identity matrix. To determine the MVU estimator, we can repeat the derivation in Section 4.3 (see Problem 4.9). Alternatively, we can use a whitening approach as follows. Since $\mathbf{C}$ is assumed to be positive definite, $\mathbf{C}^{-1}$ is positive definite and so can be factored as

$$\mathbf{C}^{-1} = \mathbf{D}^T \mathbf{D} \quad (4.24)$$

where $\mathbf{D}$ is an $N \times N$ invertible matrix. The matrix $\mathbf{D}$ acts as a whitening transformation when applied to $\mathbf{w}$ since

$$\begin{aligned} E[(\mathbf{D}\mathbf{w})(\mathbf{D}\mathbf{w})^T] &= \mathbf{D}\mathbf{C}\mathbf{D}^T \\ &= \mathbf{D}\mathbf{D}^{-1}\mathbf{D}^{T-1}\mathbf{D}^T = \mathbf{I} \end{aligned}$$

As a consequence, if we transform our generalized model

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

to

$$\begin{aligned}\mathbf{x}' &= \mathbf{D}\mathbf{x} \\ &= \mathbf{D}\mathbf{H}\boldsymbol{\theta} + \mathbf{D}\mathbf{w} \\ &= \mathbf{H}'\boldsymbol{\theta} + \mathbf{w}'\end{aligned}$$

the noise will be whitened since $\mathbf{w}' = \mathbf{D}\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and the usual linear model will result. The MVU estimator of $\boldsymbol{\theta}$ is, from (4.9),

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= (\mathbf{H}'^T \mathbf{H}')^{-1} \mathbf{H}'^T \mathbf{x}' \\ &= (\mathbf{H}^T \mathbf{D}^T \mathbf{D} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{D}^T \mathbf{D} \mathbf{x}\end{aligned}$$

so that

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (4.25)$$

In a similar fashion we find that

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = (\mathbf{H}'^T \mathbf{H}')^{-1}$$

or finally

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \quad (4.26)$$

Of course, if $\mathbf{C} = \sigma^2 \mathbf{I}$, we have our previous results. The use of the general linear model is illustrated by an example.

Example 4.4 - DC Level in Colored Noise

We now extend Example 3.3 to the colored noise case. If $x[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$, where $w[n]$ is *colored* Gaussian noise with $N \times N$ covariance matrix $\mathbf{C}$, it immediately follows from (4.25) that with $\mathbf{H} = \mathbf{1} = [1 \ 1 \ \dots \ 1]^T$, the MVU estimator of the DC level is

$$\begin{aligned}\hat{A} &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \\ &= \frac{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{x}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}\end{aligned}$$

and the variance is, from (4.26),

$$\begin{aligned}\text{var}(\hat{A}) &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \\ &= \frac{1}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}.\end{aligned}$$

If $\mathbf{C} = \sigma^2 \mathbf{I}$, we have as our MVU estimator the sample mean with a variance of σ^2/N . An interesting interpretation of the MVU estimator follows by considering the factorization of $\mathbf{C}^{-1}$ as $\mathbf{D}^T \mathbf{D}$. We noted previously that $\mathbf{D}$ is a whitening matrix. The MVU estimator is expressed as

$$\begin{aligned}\hat{A} &= \frac{\mathbf{1}^T \mathbf{D}^T \mathbf{D} \mathbf{x}}{\mathbf{1}^T \mathbf{D}^T \mathbf{D} \mathbf{1}} \\ &= \frac{(\mathbf{D} \mathbf{1})^T \mathbf{x}'}{\mathbf{1}^T \mathbf{D}^T \mathbf{D} \mathbf{1}} \\ &= \sum_{n=0}^{N-1} d_n x'[n]\end{aligned}\quad (4.27)$$

where $d_n = [\mathbf{D} \mathbf{1}]_n / \mathbf{1}^T \mathbf{D}^T \mathbf{D} \mathbf{1}$. According to (4.27), the data are first prewhitened to form $x'[n]$ and then “averaged” using prewhitened averaging weights d_n . The prewhitening has the effect of decorrelating and equalizing the variances of the noises at each observation time before averaging (see Problems 4.10 and 4.11).

◇

Another extension to the linear model allows for signal components that are known. Assume that $\mathbf{s}$ represents a known signal contained in the data. Then, a linear model that incorporates this signal is

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{s} + \mathbf{w}.$$

To determine the MVU estimator let $\mathbf{x}' = \mathbf{x} - \mathbf{s}$, so that

$$\mathbf{x}' = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

which is now in the form of the linear model. The MVU estimator follows as

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T (\mathbf{x} - \mathbf{s}) \quad (4.28)$$

with covariance

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \sigma^2 (\mathbf{H}^T \mathbf{H})^{-1}. \quad (4.29)$$

Example 4.5 - DC Level and Exponential in White Noise

If $x[n] = A + r^n + w[n]$ for $n = 0, 1, \dots, N-1$, where r is known, A is to be estimated, and $w[n]$ is WGN, the model is

$$\mathbf{x} = A \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} + \mathbf{s} + \mathbf{w}$$

where $\mathbf{s} = [1r \dots r^{N-1}]^T$. The MVU estimator, is from (4.28),

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - r^n)$$

with variance, from (4.29), as

$$\text{var}(\hat{A}) = \frac{\sigma^2}{N}.$$

◇

It should be clear that the two extensions described can be combined to produce the general linear model summarized by the following theorem.

Theorem 4.2 (Minimum Variance Unbiased Estimator for General Linear Model) If the data can be modeled as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{s} + \mathbf{w} \quad (4.30)$$

where $\mathbf{x}$ is an $N \times 1$ vector of observations, $\mathbf{H}$ is a known $N \times p$ observation matrix ($N > p$) of rank p , $\boldsymbol{\theta}$ is a $p \times 1$ vector of parameters to be estimated, $\mathbf{s}$ is an $N \times 1$ vector of known signal samples, and $\mathbf{w}$ is an $N \times 1$ noise vector with PDF $N(\mathbf{0}, \mathbf{C})$, then the MVU estimator is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}) \quad (4.31)$$

and the covariance matrix is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \quad (4.32)$$

For the general linear model the MVU estimator is efficient in that it attains the CRLB.

This theorem is quite powerful in practice since many signal processing problems can be modeled by (4.30).

References

- Graybill, F.A., *Theory and Application of the Linear Model*, Duxbury Press, North Scituate, Mass., 1976.
- MacWilliams, F.J., N.J. Sloane, "Pseudo-Random Sequences and Arrays," *Proc. IEEE*, Vol. 64, pp. 1715-1729, Dec. 1976.

Problems

4.1 We wish to estimate the amplitudes of exponentials in noise. The observed data are

$$x[n] = \sum_{i=1}^p A_i r_i^n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 . Find the MVU estimator of the amplitudes and also their covariance. Evaluate your results for the case when $p = 2$, $r_1 = 1$, $r_2 = -1$, and N is even.

4.2 Prove that the inverse of $\mathbf{H}^T \mathbf{H}$ exists if and only if the columns of $\mathbf{H}$ are linearly independent. Hint: The problem is equivalent to proving that $\mathbf{H}^T \mathbf{H}$ is positive definite and hence invertible if and only if the columns are linearly independent.

4.3 Consider the observation matrix

$$\mathbf{H} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 + \epsilon \end{bmatrix}$$

where ϵ is small. Compute $(\mathbf{H}^T \mathbf{H})^{-1}$ and examine what happens as $\epsilon \rightarrow 0$. If $\mathbf{x} = [2 \ 2 \ 2]^T$, find the MVU estimator and describe what happens as $\epsilon \rightarrow 0$.

4.4 In the linear model it is desired to estimate the signal $\mathbf{s} = \mathbf{H}\boldsymbol{\theta}$. If an MVU estimator of $\boldsymbol{\theta}$ is found, then the signal may be estimated as $\hat{\mathbf{s}} = \mathbf{H}\hat{\boldsymbol{\theta}}$. What is the PDF of $\hat{\mathbf{s}}$? Apply your results to the linear model in Example 4.2.

4.5 Prove that for $k, l = 1, 2, \dots, M < N/2$

$$\sum_{n=0}^{N-1} \cos\left(\frac{2\pi kn}{N}\right) \cos\left(\frac{2\pi ln}{N}\right) = \frac{N}{2} \delta_{kl}$$

by using the trigonometric identity

$$\cos \omega_1 \cos \omega_2 = \frac{1}{2} \cos(\omega_1 + \omega_2) + \frac{1}{2} \cos(\omega_1 - \omega_2)$$

and noting that

$$\sum_{n=0}^{N-1} \cos \alpha n = \operatorname{Re} \left(\sum_{n=0}^{N-1} \exp(j\alpha n) \right).$$

4.6 Assume that in Example 4.2 we have a single sinusoidal component at $f_k = k/N$. The model as given by (4.12) is

$$x[n] = a_k \cos(2\pi f_k n) + b_k \sin(2\pi f_k n) + w[n] \quad n = 0, 1, \dots, N-1.$$

Using the identity $A \cos \omega + B \sin \omega = \sqrt{A^2 + B^2} \cos(\omega - \phi)$, where $\phi = \arctan(B/A)$ we can rewrite the model as

$$x[n] = \sqrt{a_k^2 + b_k^2} \cos(2\pi f_k n - \phi) + w[n].$$

An MVU estimator is used for a_k, b_k , so that the estimated power of the sinusoid is

$$\hat{P} = \frac{\hat{a}_k^2 + \hat{b}_k^2}{2}.$$

A measure of detectability is $E^2(\hat{P})/\text{var}(\hat{P})$. Compare the measure when a sinusoid is present to the case when only noise is present or $a_k = b_k = 0$. Could you use the estimated power to decide if a signal is present?

4.7 Verify (4.20) by letting the limits of the summation in (4.19) be $n = -\infty$ to $n = \infty$ and noting that $u[n] = 0$ for $n < 0$ and $n > N - 1$.

4.8 We wish to estimate the frequency response $H(f)$ of a causal linear system based on the observed input and output processes. The input process $u[n]$ is assumed to be a WSS random process, as is the output $x[n]$. Prove that the cross-power spectral density between the input and output is

$$P_{ux}(f) = H(f)P_{uu}(f)$$

where $P_{ux}(f) = \mathcal{F}\{r_{ux}[k]\}$ and $r_{ux}[k] = E(u[n]x[n+k])$. If the input process is white noise with PSD $P_{uu}(f) = \sigma^2$, propose an estimator for the frequency response. Could you estimate the frequency response if $P_{uu}(f)$ were zero for a band of frequencies? Can you relate your estimator to (4.23) for the case when the linear system is a TDL?

4.9 Derive (4.25) and (4.26) by repeating the derivation in Section 4.3.

4.10 If in Example 4.4 the noise samples are uncorrelated but of unequal variance or

$$\mathbf{C} = \text{diag}(\sigma_0^2, \sigma_1^2, \dots, \sigma_{N-1}^2)$$

find d_n and interpret the results. What would happen to $\hat{A}$ if a single σ_n^2 were equal to zero?

4.11 Assume the data model described in Problem 3.9. Find the MVU estimator of A and its variance. Why is the estimator independent of ρ ? What happens if $\rho \rightarrow \pm 1$?

4.12 In this problem we investigate the estimation of a linear function of θ . Letting this new parameter be $\alpha = \mathbf{A}\theta$, where $\mathbf{A}$ is a known $r \times p$ matrix with $r < p$ and rank r , show that the MVU estimator is

$$\hat{\alpha} = \mathbf{A}\hat{\theta}$$

where $\hat{\theta}$ is the MVU estimator of θ . Also, find the covariance matrix. Hint: Replace $\mathbf{x}$ by

$$\mathbf{x}' = \mathbf{A}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{x}$$

where $\mathbf{x}'$ is $r \times 1$. This results in the *reduced linear model*. It can be shown that $\mathbf{x}'$ contains all the information about θ so that we may base our estimator on this lower dimensionality data set [Graybill 1976].

- 4.13** In practice we sometimes encounter the “linear model” $\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$ but with $\mathbf{H}$ composed of random variables. Suppose we ignore this difference and use our usual estimator

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x},$$

where we assume that the particular realization of $\mathbf{H}$ is known to us. Show that if $\mathbf{H}$ and $\mathbf{w}$ are independent, the mean and covariance of $\hat{\boldsymbol{\theta}}$ are

$$\begin{aligned} E(\hat{\boldsymbol{\theta}}) &= \boldsymbol{\theta} \\ \mathbf{C}_{\hat{\boldsymbol{\theta}}} &= \sigma^2 E_H [(\mathbf{H}^T \mathbf{H})^{-1}] \end{aligned}$$

where E_H denotes the expectation with respect to the PDF of $\mathbf{H}$. What happens if the independence assumption is not made?

- 4.14** Suppose we observe a signal, which is subject to fading, in noise. We model the fading process as resulting in a signal that is either “on” or “off”. As a simple illustration, consider the DC level in WGN or $x[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$. When the signal fades, the data model becomes

$$x[n] = \begin{cases} A + w[n] & n = 0, 1, \dots, M-1 \\ w[n] & n = M, M+1, \dots, N-1 \end{cases}$$

where the probability of a fade is ϵ . Assuming we know when the signal has experienced a fade, use the results of Problem 4.13 to determine an estimator of A and also its variance. Compare your results to the case of no fading.

Chapter 5

General Minimum Variance Unbiased Estimation

5.1 Introduction

We have seen that the evaluation of the CRLB sometimes results in an efficient and hence MVU estimator. In particular, the linear model provides a useful example of this approach. If, however, an efficient estimator does not exist, it is still of interest to be able to find the MVU estimator (assuming of course that it exists). To do so requires the concept of sufficient statistics and the important Rao-Blackwell-Lehmann-Scheffe theorem. Armed with this theory it is possible in many cases to determine the MVU estimator by a simple inspection of the PDF. How this is done is explored in this chapter.

5.2 Summary

In Section 5.3 we define a sufficient statistic as one which summarizes the data in the sense that if we are given the sufficient statistic, the PDF of the data no longer depends on the unknown parameter. The Neyman-Fisher factorization theorem (see Theorem 5.1) enables us to easily find sufficient statistics by examining the PDF. Once the sufficient statistic has been found, the MVU estimator may be determined by finding a function of the sufficient statistic that is an unbiased estimator. The essence of the approach is summarized in Theorem 5.2 and is known as the Rao-Blackwell-Lehmann-Scheffe theorem. In using this theorem we must also ensure that the sufficient statistic is complete or that there exists only one function of it that is an unbiased estimator. The extension to the vector parameter case is given in Theorems 5.3 and 5.4. The approach is basically the same as for a scalar parameter. We must first find as many sufficient statistics as unknown parameters and then determine an unbiased estimator based on the sufficient statistics.

5.3 Sufficient Statistics

In a previous chapter we found that for the problem of estimating a DC level A in WGN (see Example 3.3), the sample mean

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

was the MVU estimator, having minimum variance σ^2/N . If, on the other hand, we had chosen

$$\hat{A} = x[0]$$


as our estimator, it is immediately clear that even though $\hat{A}$ is unbiased, its variance is much larger (being σ^2) than the minimum. Intuitively, the poor performance is a direct result of discarding the data points $\{x[1], x[2], \dots, x[N-1]\}$ which carry information about A . A reasonable question to ask is Which data samples are pertinent to the estimation problem? or Is there a set of data that is sufficient? The following data sets may be claimed to be sufficient in that they may be used to compute $\hat{A}$.

$$\begin{aligned} S_1 &= \{x[0], x[1], \dots, x[N-1]\} \\ S_2 &= \{x[0] + x[1], x[2], x[3], \dots, x[N-1]\} \\ S_3 &= \left\{ \sum_{n=0}^{N-1} x[n] \right\}. \end{aligned}$$

S_1 represents the original data set, which as expected, is always sufficient for the problem. S_2 and S_3 are also sufficient. It is obvious that for this problem there are many sufficient data sets. The data set that contains the least number of elements is called the *minimal one*. If we now think of the elements of these sets as statistics, we say that the N statistics of S_1 are *sufficient* as well as the $(N-1)$ statistics of S_2 and the single statistic of S_3 . This latter statistic, $\sum_{n=0}^{N-1} x[n]$, in addition to being a *sufficient statistic*, is the *minimal sufficient statistic*. For estimation of A , once we know $\sum_{n=0}^{N-1} x[n]$, we no longer need the individual data values since all information has been summarized in the sufficient statistic. To quantify what we mean by this, consider the PDF of the data

$$p(\mathbf{x}; A) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right] \quad (5.1)$$

and assume that $T(\mathbf{x}) = \sum_{n=0}^{N-1} x[n] = T_0$ has been observed. Knowledge of the value of this statistic will change the PDF to the conditional one $p(\mathbf{x} | \sum_{n=0}^{N-1} x[n] = T_0; A)$, which now gives the PDF of the observations after the sufficient statistic has been observed. Since the statistic is sufficient for the estimation of A , this conditional PDF should not



(a) Observations provide information after $T(\mathbf{x})$ observed— $T(\mathbf{x})$ is not sufficient (b) No information from observations after $T(\mathbf{x})$ observed— $T(\mathbf{x})$ is sufficient

Figure 5.1 Sufficient statistic definition

depend on A . If it did, then we could infer some additional information about A from the data in addition to that already provided by the sufficient statistic. As an example, in Figure 5.1a, if $\mathbf{x} = \mathbf{x}_0$ for an arbitrary $\mathbf{x}_0$, then values of A near A_0 would be more likely. This violates our notion that $\sum_{n=0}^{N-1} x[n]$ is a sufficient statistic. On the other hand, in Figure 5.1b, any value of A is as likely as any other, so that after observing $T(\mathbf{x})$ the data may be discarded. Hence, to verify that a statistic is sufficient we need to determine the conditional PDF and confirm that there is no dependence on A .

Example 5.1 - Verification of a Sufficient Statistic

Consider the PDF of (5.1). To prove that $\sum_{n=0}^{N-1} x[n]$ is a sufficient statistic we need to determine $p(\mathbf{x}|T(\mathbf{x}) = T_0; A)$, where $T(\mathbf{x}) = \sum_{n=0}^{N-1} x[n]$. By the definition of the conditional PDF we have

$$p(\mathbf{x}|T(\mathbf{x}) = T_0; A) = \frac{p(\mathbf{x}, T(\mathbf{x}) = T_0; A)}{p(T(\mathbf{x}) = T_0; A)}.$$

But note that $T(\mathbf{x})$ is functionally dependent on $\mathbf{x}$, so that the joint PDF $p(\mathbf{x}, T(\mathbf{x}) = T_0; A)$ takes on nonzero values only when $\mathbf{x}$ satisfies $T(\mathbf{x}) = T_0$. The joint PDF is therefore $p(\mathbf{x}; A)\delta(T(\mathbf{x}) - T_0)$, where δ is the Dirac delta function (see also Appendix 5A for a further discussion). Thus, we have that

$$p(\mathbf{x}|T(\mathbf{x}) = T_0; A) = \frac{p(\mathbf{x}; A)\delta(T(\mathbf{x}) - T_0)}{p(T(\mathbf{x}) = T_0; A)}. \quad (5.2)$$

Clearly, $T(\mathbf{x}) \sim \mathcal{N}(NA, N\sigma^2)$, so that

$$\begin{aligned} & p(\mathbf{x}; A) \delta(T(\mathbf{x}) - T_0) \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right] \delta(T(\mathbf{x}) - T_0) \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \left(\sum_{n=0}^{N-1} x^2[n] - 2AT(\mathbf{x}) + NA^2 \right) \right] \delta(T(\mathbf{x}) - T_0) \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \left(\sum_{n=0}^{N-1} x^2[n] - 2AT_0 + NA^2 \right) \right] \delta(T(\mathbf{x}) - T_0). \end{aligned}$$

From (5.2) we have

$$\begin{aligned} & p(\mathbf{x}|T(\mathbf{x}) = T_0; A) \\ &= \frac{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right] \exp \left[-\frac{1}{2\sigma^2} (-2AT_0 + NA^2) \right]}{\frac{1}{\sqrt{2\pi N\sigma^2}} \exp \left[-\frac{1}{2N\sigma^2} (T_0 - NA)^2 \right]} \delta(T(\mathbf{x}) - T_0) \\ &= \frac{\sqrt{N}}{(2\pi\sigma^2)^{\frac{N-1}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right] \exp \left[\frac{T_0^2}{2N\sigma^2} \right] \delta(T(\mathbf{x}) - T_0) \end{aligned}$$

which as claimed does not depend on A . Therefore, we can conclude that $\sum_{n=0}^{N-1} x[n]$ is a sufficient statistic for the estimation of A . $\diamond$

This example indicates the procedure for verifying that a statistic is sufficient. For many problems the task of evaluating the conditional PDF is formidable, so that an easier approach is needed. Additionally, in Example 5.1 the choice of $\sum_{n=0}^{N-1} x[n]$ for examination as a sufficient statistic was fortuitous. In general an even more difficult problem would be to *identify* potential sufficient statistics. The approach of guessing at a sufficient statistic and then verifying it is, of course, quite unsatisfactory in practice. To alleviate the guesswork we can employ the Neyman-Fisher factorization theorem, which is a simple “turn-the-crank” procedure for finding sufficient statistics.

5.4 Finding Sufficient Statistics

The Neyman-Fisher factorization theorem is now stated, after which we will use it to find sufficient statistics in several examples.

Theorem 5.1 (Neyman-Fisher Factorization) *If we can factor the PDF $p(\mathbf{x}; \theta)$ as*

$$p(\mathbf{x}; \theta) = g(T(\mathbf{x}), \theta) h(\mathbf{x}) \quad (5.3)$$

where g is a function depending on $\mathbf{x}$ only through $T(\mathbf{x})$ and h is a function depending only on $\mathbf{x}$, then $T(\mathbf{x})$ is a sufficient statistic for θ . Conversely, if $T(\mathbf{x})$ is a sufficient statistic for θ , then the PDF can be factored as in (5.3).

A proof of this theorem is contained in Appendix 5A. It should be mentioned that at times it is not obvious if the PDF can be factored in the required form. If this is the case, then a sufficient statistic may not exist. Some examples are now given to illustrate the use of this powerful theorem.

Example 5.2 - DC Level in WGN

We now reexamine the problem discussed in the previous section. There the PDF was given by (5.1), where we note that σ^2 is assumed known. To demonstrate that a factorization exists we observe that the exponent of the PDF may be rewritten as

$$\sum_{n=0}^{N-1} (x[n] - A)^2 = \sum_{n=0}^{N-1} x^2[n] - 2A \sum_{n=0}^{N-1} x[n] + NA^2$$

so that the PDF is factorable as

$$p(\mathbf{x}; A) = \underbrace{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \left(NA^2 - 2A \sum_{n=0}^{N-1} x[n] \right) \right]}_{g(T(\mathbf{x}), A)} \underbrace{\exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right]}_{h(\mathbf{x})}.$$

Clearly then, $T(\mathbf{x}) = \sum_{n=0}^{N-1} x[n]$ is a sufficient statistic for A . Note that $T'(\mathbf{x}) = 2 \sum_{n=0}^{N-1} x[n]$ is also a sufficient statistic for A , and in fact *any one-to-one function of* $\sum_{n=0}^{N-1} x[n]$ is a sufficient statistic (see Problem 5.12). Hence, sufficient statistics are unique only to within one-to-one transformations. $\diamond$

Example 5.3 - Power of WGN

Now consider the PDF of (5.1) with $A = 0$ and σ^2 as the unknown parameter. Then,

$$p(\mathbf{x}; \sigma^2) = \underbrace{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right]}_{g(T(\mathbf{x}), \sigma^2)} \underbrace{1}_{h(\mathbf{x})}.$$

Again it is immediately obvious from the factorization theorem that $T(\mathbf{x}) = \sum_{n=0}^{N-1} x^2[n]$ is a sufficient statistic for σ^2 . See also Problem 5.1. $\diamond$

Example 5.4 - Phase of Sinusoid

Recall the problem in Example 3.4 in which we wish to estimate the phase of a sinusoid embedded in WGN or

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n] \quad n = 0, 1, \dots, N-1.$$

Here, the amplitude A and frequency f_0 of the sinusoid are known, as is the noise variance σ^2 . The PDF is

$$p(\mathbf{x}; \phi) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} [x[n] - A \cos(2\pi f_0 n + \phi)]^2 \right\}.$$

The exponent may be expanded as

$$\begin{aligned} & \sum_{n=0}^{N-1} x^2[n] - 2A \sum_{n=0}^{N-1} x[n] \cos(2\pi f_0 n + \phi) + \sum_{n=0}^{N-1} A^2 \cos^2(2\pi f_0 n + \phi) \\ &= \sum_{n=0}^{N-1} x^2[n] - 2A \left(\sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \right) \cos \phi \\ & \quad + 2A \left(\sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n \right) \sin \phi + \sum_{n=0}^{N-1} A^2 \cos^2(2\pi f_0 n + \phi). \end{aligned}$$

In this example it does not appear that the PDF is factorable as required by the Neyman-Fisher theorem. Hence, no *single* sufficient statistic exists. However, it can be factored as

$$\begin{aligned} p(\mathbf{x}; \phi) = & \underbrace{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \left[\sum_{n=0}^{N-1} A^2 \cos^2(2\pi f_0 n + \phi) - 2AT_1(\mathbf{x}) \cos \phi + 2AT_2(\mathbf{x}) \sin \phi \right] \right\}}_{g(T_1(\mathbf{x}), T_2(\mathbf{x}), \phi)} \\ & \cdot \underbrace{\exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right]}_{h(\mathbf{x})} \end{aligned}$$

where

$$\begin{aligned} T_1(\mathbf{x}) &= \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \\ T_2(\mathbf{x}) &= \sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n. \end{aligned}$$

By a slight generalization of the Neyman-Fisher theorem we can conclude that $T_1(\mathbf{x})$ and $T_2(\mathbf{x})$ are *jointly* sufficient statistics for the estimation of ϕ . However, no *single* sufficient statistic exists. The reason why we wish to restrict our attention to single sufficient statistics will become clear in the next section. $\diamond$

The concept of jointly sufficient statistics is a simple extension of our previous definition. The r statistics $T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_r(\mathbf{x})$ are jointly sufficient statistics if the conditional PDF $p(\mathbf{x}|T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_r(\mathbf{x}); \theta)$ does not depend on θ . The generalization of the Neyman-Fisher theorem asserts that if $p(\mathbf{x}; \theta)$ can be factored as [Kendall and Stuart 1979]

$$p(\mathbf{x}; \theta) = g(T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_r(\mathbf{x}), \theta)h(\mathbf{x}) \quad (5.4)$$

then $\{T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_r(\mathbf{x})\}$ are sufficient statistics for θ . It is clear then that the original data are always sufficient statistics since we can let $r = N$ and

$$T_{n+1}(\mathbf{x}) = x[n] \quad n = 0, 1, \dots, N-1$$

so that

$$\begin{aligned} g &= p(\mathbf{x}; \theta) \\ h &= 1 \end{aligned}$$

and (5.4) holds identically. Of course, they are seldom the *minimal* set of sufficient statistics.

5.5 Using Sufficiency to Find the MVU Estimator

Assuming that we have been able to find a sufficient statistic $T(\mathbf{x})$ for θ , we can make use of the Rao-Blackwell-Lehmann-Scheffe (RBLs) theorem to find the MVU estimator. We will first illustrate the approach with an example and then state the theorem formally.

Example 5.5 - DC Level in WGN

We will continue Example 5.2. Although we already know that $\hat{A} = \bar{x}$ is the MVU estimator (since it is efficient), we will use the RBLs theorem, which can be used even when an efficient estimator does not exist and hence the CRLB method is no longer viable. The procedure for finding the MVU estimator $\hat{A}$ may be implemented in two different ways. They are both based on the sufficient statistic $T(\mathbf{x}) = \sum_{n=0}^{N-1} x[n]$.

1. Find *any* unbiased estimator of A , say $\tilde{A} = x[0]$, and determine $\hat{A} = E(\tilde{A}|T)$. The expectation is taken with respect to $p(\tilde{A}|T)$.
2. Find some function g so that $\hat{A} = g(T)$ is an unbiased estimator of A .

For the first approach we can let the unbiased estimator be $\tilde{A} = x[0]$ and determine $\hat{A} = E(x[0] | \sum_{n=0}^{N-1} x[n])$. To do so we will need some properties of the conditional

Gaussian PDF. For $[x y]^T$ a Gaussian random vector with mean vector $\mu = [E(x) E(y)]^T$ and covariance matrix

$$\mathbf{C} = \begin{bmatrix} \text{var}(x) & \text{cov}(x, y) \\ \text{cov}(y, x) & \text{var}(y) \end{bmatrix},$$

it may be shown that (see Appendix 10A)

$$\begin{aligned} E(x|y) &= \int_{-\infty}^{\infty} xp(x|y) dx \\ &= \int_{-\infty}^{\infty} x \frac{p(x, y)}{p(y)} dx \\ &= E(x) + \frac{\text{cov}(x, y)}{\text{var}(y)}(y - E(y)). \end{aligned} \quad (5.5)$$

Applying this result, we let $x = x[0]$ and $y = \sum_{n=0}^{N-1} x[n]$ and note that

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x[0] \\ \sum_{n=0}^{N-1} x[n] \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 1 & \dots & 1 \end{bmatrix}}_{\mathbf{L}} \begin{bmatrix} x[0] \\ x[1] \\ \vdots \\ x[N-1] \end{bmatrix}.$$

Hence, the PDF of $[x y]^T$ is $\mathcal{N}(\mu, \mathbf{C})$ since this represents a linear transformation of a Gaussian vector, where

$$\begin{aligned} \mu &= \mathbf{L}E(\mathbf{x}) = \mathbf{L}\mathbf{A}\mathbf{1} = \begin{bmatrix} A \\ NA \end{bmatrix} \\ \mathbf{C} &= \sigma^2 \mathbf{L}\mathbf{L}^T = \sigma^2 \begin{bmatrix} 1 & 1 \\ 1 & N \end{bmatrix}. \end{aligned}$$

Hence, we have finally from (5.5) that

$$\begin{aligned} \hat{A} &= E(x|y) = A + \frac{\sigma^2}{N\sigma^2} \left(\sum_{n=0}^{N-1} x[n] - NA \right) \\ &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] \end{aligned}$$

which is the MVU estimator. This approach, requiring evaluation of a conditional expectation, is usually mathematically intractable.

Turning our attention to the second approach, we need to find some function g so that

$$\hat{A} = g \left(\sum_{n=0}^{N-1} x[n] \right)$$




Figure 5.2 RBLs argument for MVU estimator

is an unbiased estimator of A . By inspection this is $g(x) = x/N$, which yields

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

as the MVU estimator. This alternative method is much easier to apply, and therefore in practice, it is the one we generally employ. $\diamond$

We now formally state the RBLs theorem.

Theorem 5.2 (Rao-Blackwell-Lehmann-Scheffe) *If $\hat{\theta}$ is an unbiased estimator of θ and $T(\mathbf{x})$ is a sufficient statistic for θ , then $\hat{\theta} = E(\hat{\theta}|T(\mathbf{x}))$ is*

1. *a valid estimator for θ (not dependent on θ)*
2. *unbiased*
3. *of lesser or equal variance than that of $\hat{\theta}$, for all θ .*

Additionally, if the sufficient statistic is complete, then $\hat{\theta}$ is the MVU estimator.

A proof is given in Appendix 5B. In the previous example we saw that $E(x[0] | \sum_{n=0}^{N-1} x[n]) = \bar{x}$ did not depend on A , making it a valid estimator, was unbiased, and had less variance than $x[0]$. That there is no other estimator with less variance, as Theorem 5.2 asserts, follows from the property that the sufficient statistic $\sum_{n=0}^{N-1} x[n]$ is a *complete* sufficient statistic. In essence, a statistic is complete if *there is only one function of the statistic that is unbiased*. The argument that $\hat{\theta} = E(\hat{\theta}|T(\mathbf{x}))$ is the MVU estimator is now given. Consider all possible unbiased estimators of θ , as depicted in Figure 5.2. By determining $E(\hat{\theta}|T(\mathbf{x}))$ we can lower the variance of the estimator (property 3 of Theorem 5.2) and still remain within the class (property 2 of Theorem 5.2). But $E(\hat{\theta}|T(\mathbf{x}))$ is solely a function of the sufficient statistic $T(\mathbf{x})$ since

$$\begin{aligned} \hat{\theta} &= E(\hat{\theta}|T(\mathbf{x})) = \int \hat{\theta} p(\hat{\theta}|T(\mathbf{x})) d\hat{\theta} \\ &= g(T(\mathbf{x})). \end{aligned} \tag{5.6}$$

If $T(\mathbf{x})$ is complete, there is but *one* function of T that is an unbiased estimator. Hence, $\hat{\theta}$ is *unique*, independent of the $\tilde{\theta}$ we choose from the class shown in Figure 5.2. Every $\tilde{\theta}$ maps into the *same estimator* $\hat{\theta}$. Because the variance of $\hat{\theta}$ must be less than $\tilde{\theta}$ for any $\tilde{\theta}$ within the class (property 3 of Theorem 5.2), we conclude that $\hat{\theta}$ must be the MVU estimator. In summary, the MVU estimator can be found by taking *any* unbiased estimator and carrying out the operations in (5.6). Alternatively, since there is only one function of the sufficient statistic that leads to an unbiased estimator, we need only find the unique g to make the sufficient statistic unbiased. For this latter approach we found in Example 5.5 that $g(\sum_{n=0}^{N-1} x[n]) = \sum_{n=0}^{N-1} x[n]/N$.

The property of completeness depends on the PDF of $\mathbf{x}$, which in turn determines the PDF of the sufficient statistic. For many practical cases of interest it holds. In particular, for the exponential family of PDFs (see Problems 5.14 and 5.15) this condition is satisfied. To validate that a sufficient statistic is complete is in general quite difficult, and we refer the reader to the discussions in [Kendall and Stuart 1979]. A flavor for the concept of completeness is provided by the next two examples.

Example 5.6 - Completeness of a Sufficient Statistic

For the estimation of A , the sufficient statistic $\sum_{n=0}^{N-1} x[n]$ is complete or there is but one function g for which $E[g(\sum_{n=0}^{N-1} x[n])] = A$. Suppose, however, that there exists a second function h for which $E[h(\sum_{n=0}^{N-1} x[n])] = A$. Then, it would follow that with $T = \sum_{n=0}^{N-1} x[n]$,

$$E[g(T) - h(T)] = A - A = 0 \quad \text{for all } A$$

or since $T \sim \mathcal{N}(NA, N\sigma^2)$

$$\int_{-\infty}^{\infty} v(T) \frac{1}{\sqrt{2\pi N\sigma^2}} \exp\left[-\frac{1}{2N\sigma^2}(T - NA)^2\right] dT = 0 \quad \text{for all } A$$

where $v(T) = g(T) - h(T)$. Letting $\tau = T/N$ and $v'(\tau) = v(N\tau)$, we have

$$\int_{-\infty}^{\infty} v'(\tau) \frac{N}{\sqrt{2\pi N\sigma^2}} \exp\left[-\frac{N}{2\sigma^2}(A - \tau)^2\right] d\tau = 0 \quad \text{for all } A \quad (5.7)$$

which may be recognized as the convolution of a function $v'(\tau)$ with a Gaussian pulse $w(\tau)$ (see Figure 5.3). For the result to be zero for all A , $v'(\tau)$ must be identically zero. To see this recall that a signal is zero if and only if its Fourier transform is identically zero, resulting in the condition

$$V'(f)W(f) = 0 \quad \text{for all } f$$

where $V'(f) = \mathcal{F}\{v'(\tau)\}$ and $W(f)$ is the Fourier transform of the Gaussian pulse in (5.7). Since $W(f)$ is also Gaussian and therefore positive for all f , we have that the condition is satisfied if and only if $V'(f) = 0$ for all f . Hence, we must have that $v'(\tau) = 0$ for all τ . This implies that $g = h$ or that the function g is unique. $\diamond$




Figure 5.3 Completeness condition for sufficient statistic (satisfied)

Example 5.7 - Incomplete Sufficient Statistic

Consider the estimation of A for the datum

$$x[0] = A + w[0]$$

where $w[0] \sim \mathcal{U}[-\frac{1}{2}, \frac{1}{2}]$. A sufficient statistic is $x[0]$, being the only available data, and furthermore, $x[0]$ is an unbiased estimator of A . We may conclude that $g(x[0]) = x[0]$ is a viable candidate for the MVU estimator. That it is actually the MVU estimator still requires us to verify that it is a complete sufficient statistic. As in the previous example, we suppose that there exists another function h with the unbiased property $h(x[0]) = A$ and attempt to prove that $h = g$. Again letting $v(T) = g(T) - h(T)$, we examine the possible solutions for v of the equation

$$\int_{-\infty}^{\infty} v(T)p(\mathbf{x}; A) d\mathbf{x} = 0 \quad \text{for all } A.$$

For this problem, however, $\mathbf{x} = x[0] = T$, so that

$$\int_{-\infty}^{\infty} v(T)p(T; A) dT = 0 \quad \text{for all } A.$$

But

$$p(T; A) = \begin{cases} 1 & A - \frac{1}{2} \leq T \leq A + \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}$$

so that the condition reduces to

$$\int_{A-\frac{1}{2}}^{A+\frac{1}{2}} v(T) dT = 0 \quad \text{for all } A.$$




Figure 5.4 Completeness condition for sufficient statistic (not satisfied)

The nonzero function $v(T) = \sin 2\pi T$ will satisfy this condition as illustrated in Figure 5.4. Hence a solution is

$$v(T) = g(T) - h(T) = \sin 2\pi T$$

or

$$h(T) = T - \sin 2\pi T.$$

As a result, the estimator

$$\hat{A} = x[0] - \sin 2\pi x[0]$$

is also based on the sufficient statistic and is unbiased for A . Having found at least one other unbiased estimator that is also a function of the sufficient statistic, we may conclude that the sufficient statistic is not complete. The RBLS theorem no longer holds, and it is not possible to assert that $\hat{A} = x[0]$ is the MVU estimator. $\diamond$

To summarize the completeness condition, we say that a sufficient statistic is complete if the condition

$$\int_{-\infty}^{\infty} v(T)p(T; \theta) dT = 0 \quad \text{for all } \theta \quad (5.8)$$

is satisfied only by the zero function or by $v(T) = 0$ for all T .

At this point it is worthwhile to review our results and then apply them to an estimation problem for which we do not know the MVU estimator. The procedure is as follows (see also Figure 5.5):

1. Find a single sufficient statistic for θ , that is, $T(\mathbf{x})$, by using the Neyman-Fisher factorization theorem.
2. Determine if the sufficient statistic is complete and, if so, proceed; if not, this approach cannot be used.




Figure 5.5 Procedure for finding MVU estimator (scalar parameter)

3. Find a function g of the sufficient statistic that yields an unbiased estimator $\hat{\theta} = g(T(\mathbf{x}))$. The MVU estimator is then $\hat{\theta}$.

As an alternative implementation of step 3 we may

- 3'. Evaluate $\hat{\theta} = E(\tilde{\theta}|T(\mathbf{x}))$, where $\tilde{\theta}$ is any unbiased estimator.

However, in practice the conditional expectation evaluation is usually too tedious. The next example illustrates the overall procedure.

Example 5.8 - Mean of Uniform Noise

We observe the data

$$x[n] = w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is IID noise with PDF $\mathcal{U}[0, \beta]$ for $\beta > 0$. We wish to find the MVU estimator for the mean $\theta = \beta/2$. Our initial approach of using the CRLB for finding an efficient and hence MVU estimator cannot even be tried for this problem. This is because the PDF does not satisfy the required regularity conditions (see Problem 3.1). A natural estimator of θ is the sample mean or

$$\hat{\theta} = \frac{1}{N} \sum_{n=0}^{N-1} x[n].$$

The sample mean is easily shown to be unbiased and to have variance

$$\begin{aligned} \text{var}(\hat{\theta}) &= \frac{1}{N} \text{var}(x[n]) \\ &= \frac{\beta^2}{12N}. \end{aligned} \tag{5.9}$$

To determine if the sample mean is the MVU estimator for this problem we will follow the procedure outlined in Figure 5.5. Define the unit step function as

$$u(x) = \begin{cases} 1 & \text{for } x > 0 \\ 0 & \text{for } x < 0. \end{cases}$$

Then,

$$p(x[n]; \theta) = \frac{1}{\beta} [u(x[n]) - u(x[n] - \beta)]$$

where $\beta = 2\theta$, and therefore the PDF of the data is

$$p(\mathbf{x}; \theta) = \frac{1}{\beta^N} \prod_{n=0}^{N-1} [u(x[n]) - u(x[n] - \beta)].$$

This PDF will be nonzero only if $0 < x[n] < \beta$ for all $x[n]$, so that

$$p(\mathbf{x}; \theta) = \begin{cases} \frac{1}{\beta^N} & 0 < x[n] < \beta \quad n = 0, 1, \dots, N-1 \\ 0 & \text{otherwise.} \end{cases}$$

Alternatively, we can write

$$p(\mathbf{x}; \theta) = \begin{cases} \frac{1}{\beta^N} & \max x[n] < \beta, \min x[n] > 0 \\ 0 & \text{otherwise} \end{cases}$$

so that

$$p(\mathbf{x}; \theta) = \frac{1}{\beta^N} \underbrace{u(\beta - \max x[n])}_{g(T(\mathbf{x}), \theta)} \underbrace{u(\min x[n])}_{h(\mathbf{x})}.$$

By the Neyman-Fisher factorization theorem $T(\mathbf{x}) = \max x[n]$ is a sufficient statistic for θ . Furthermore, it can be shown that the sufficient statistic is complete. We omit the proof.

Next we need to determine a function of the sufficient statistic to make it unbiased. To do so requires us to determine the expected value of $T = \max x[n]$. The statistic T is known as an *order statistic*. Its PDF is now derived. Evaluating first the cumulative distribution function

$$\begin{aligned} \Pr\{T \leq \xi\} &= \Pr\{x[0] \leq \xi, x[1] \leq \xi, \dots, x[N-1] \leq \xi\} \\ &= \prod_{n=0}^{N-1} \Pr\{x[n] \leq \xi\} \\ &= \Pr\{x[n] \leq \xi\}^N \end{aligned}$$

since the random variables are IID. The PDF follows as

$$\begin{aligned} p_T(\xi) &= \frac{d \Pr\{T \leq \xi\}}{d\xi} \\ &= N \Pr\{x[n] \leq \xi\}^{N-1} \frac{d \Pr\{x[n] \leq \xi\}}{d\xi}. \end{aligned}$$

But $d \Pr\{x[n] \leq \xi\}/d\xi$ is the PDF of $x[n]$ or

$$p_{x[n]}(\xi; \theta) = \begin{cases} \frac{1}{\beta} & 0 < \xi < \beta \\ 0 & \text{otherwise.} \end{cases}$$

Integrating, we obtain

$$\Pr\{x[n] \leq \xi\} = \begin{cases} 0 & \xi < 0 \\ \frac{\xi}{\beta} & 0 < \xi < \beta \\ 1 & \xi > \beta \end{cases}$$

which finally yields

$$p_T(\xi) = \begin{cases} 0 & \xi < 0 \\ N \left(\frac{\xi}{\beta}\right)^{N-1} \frac{1}{\beta} & 0 < \xi < \beta \\ 0 & \xi > \beta. \end{cases}$$

We now have

$$\begin{aligned} E(T) &= \int_{-\infty}^{\infty} \xi p_T(\xi) d\xi \\ &= \int_0^{\beta} \xi N \left(\frac{\xi}{\beta}\right)^{N-1} \frac{1}{\beta} d\xi \\ &= \frac{N}{N+1} \beta \\ &= \frac{2N}{N+1} \theta. \end{aligned}$$

To make this unbiased we let $\hat{\theta} = [(N+1)/2N]T$, so that finally the MVU estimator is

$$\hat{\theta} = \frac{N+1}{2N} \max x[n].$$

Somewhat contrary to intuition, the sample mean is not the MVU estimator of the mean for uniformly distributed noise! It is of interest to compare the minimum variance to that of the sample mean estimator variance. Noting that

$$\text{var}(\hat{\theta}) = \left(\frac{N+1}{2N}\right)^2 \text{var}(T)$$

and

$$\begin{aligned}\text{var}(T) &= \int_0^\beta \xi^2 \frac{N\xi^{N-1}}{\beta^N} d\xi - \left(\frac{N\beta}{N+1} \right)^2 \\ &= \frac{N\beta^2}{(N+1)^2(N+2)}\end{aligned}$$

we have that the minimum variance is

$$\text{var}(\hat{\theta}) = \frac{\beta^2}{4N(N+2)}. \quad (5.10)$$

This minimum variance is smaller than that of the sample mean estimator (see (5.9)) for $N \geq 2$ (for $N = 1$ the estimators are identical). The difference in the variances is substantial for large N since the MVU estimator variance goes down as $1/N^2$ as compared to the $1/N$ dependence for the sample mean estimator variance. $\diamond$

The success of the previous example hinges on being able to find a single sufficient statistic. If this is not possible, then this approach cannot be employed. For instance, recalling the phase estimation example (Example 5.4) we required two statistics, $T_1(\mathbf{x})$ and $T_2(\mathbf{x})$. Theorem 5.2 may be extended to address this case. However, the MVU estimator (assuming complete sufficient statistics or $E(g(T_1, T_2)) = 0$ for all ϕ implies that $g = 0$) would need to be determined as $\hat{\phi} = E(\check{\phi}|T_1, T_2)$, where $\check{\phi}$ is any unbiased estimator of the phase. Otherwise, we will need to guess at the form of the function g that combines T_1 and T_2 into a single unbiased estimator for ϕ (see Problem 5.10). As mentioned previously, the evaluation of the conditional expectation is difficult.

5.6 Extension to a Vector Parameter

We now extend our results to search for the MVU estimator for a $p \times 1$ vector parameter $\boldsymbol{\theta}$. In this case we are seeking an unbiased vector estimator (each element is unbiased), such that each element has minimum variance. Many of the results of the previous sections extend naturally. In the vector case we may have more sufficient statistics than the number of parameters ($r > p$), exactly the same number ($r = p$), or even fewer sufficient statistics than parameters to be estimated ($r < p$). The cases of $r > p$ and $r = p$ will be discussed in the following examples. For an example of $r < p$ see Problem 5.16. From our previous discussions a desirable situation would be to have $r = p$. This would allow us to determine the MVU estimator by transforming the sufficient statistics to be unbiased. The RBLS theorem may be extended to justify this approach for a vector parameter. Unfortunately, this approach is not always possible. Before presenting some examples to illustrate these possibilities, extensions of the basic definitions and theorems will be given.

A vector statistic $\mathbf{T}(\mathbf{x}) = [T_1(\mathbf{x}) T_2(\mathbf{x}) \dots T_r(\mathbf{x})]^T$ is said to be sufficient for the estimation of $\boldsymbol{\theta}$ if the PDF of the data conditioned on the statistic or $p(\mathbf{x}|\mathbf{T}(\mathbf{x}); \boldsymbol{\theta})$ does not depend on $\boldsymbol{\theta}$. This is to say that $p(\mathbf{x}|\mathbf{T}(\mathbf{x}); \boldsymbol{\theta}) = p(\mathbf{x}|\mathbf{T}(\mathbf{x}))$. In general, many such

$\mathbf{T}(\mathbf{x})$ exist (including the data $\mathbf{x}$), so that we are interested in the *minimal sufficient statistic* or the $\mathbf{T}$ of minimum dimension. To find the minimal sufficient statistic we can again resort to the Neyman-Fisher factorization theorem which for the vector parameter case is as follows.

Theorem 5.3 (Neyman-Fisher Factorization Theorem (Vector Parameter))

If we can factor the PDF $p(\mathbf{x}; \boldsymbol{\theta})$ as

$$p(\mathbf{x}; \boldsymbol{\theta}) = g(\mathbf{T}(\mathbf{x}), \boldsymbol{\theta})h(\mathbf{x}) \quad (5.11)$$

where g is a function depending only on $\mathbf{x}$ through $\mathbf{T}(\mathbf{x})$, an $r \times 1$ statistic, and also on $\boldsymbol{\theta}$, and h is a function depending only on $\mathbf{x}$, then $\mathbf{T}(\mathbf{x})$ is a sufficient statistic for $\boldsymbol{\theta}$. Conversely, if $\mathbf{T}(\mathbf{x})$ is a sufficient statistic for $\boldsymbol{\theta}$, then the PDF can be factored as in (5.11).

A proof can be found in [Kendall and Stuart 1979]. Note that there always exists a set of sufficient statistics, the data set $\mathbf{x}$ itself being sufficient. However, it is the dimensionality of the sufficient statistic that is in question, and that is what the factorization theorem will tell us. Some examples now follow.

Example 5.9 - Sinusoidal Parameter Estimation

Assume that a sinusoidal signal is embedded in WGN

$$x[n] = A \cos 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where the amplitude A , frequency f_0 , and noise variance σ^2 are unknown. The unknown parameter vector is therefore $\boldsymbol{\theta} = [A f_0 \sigma^2]^T$. The PDF is

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A \cos 2\pi f_0 n)^2 \right].$$

Expanding the exponent, we have

$$\sum_{n=0}^{N-1} (x[n] - A \cos 2\pi f_0 n)^2 = \sum_{n=0}^{N-1} x^2[n] - 2A \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n + A^2 \sum_{n=0}^{N-1} \cos^2 2\pi f_0 n.$$

Now because of the term $\sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n$, where f_0 is *unknown*, we cannot reduce the PDF to have the form expressed in (5.11). If, on the other hand, the frequency were known, the unknown parameter vector would be $\boldsymbol{\theta} = [A \sigma^2]^T$ and the PDF could be expressed as

$$p(\mathbf{x}; \boldsymbol{\theta}) = \underbrace{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \left(\sum_{n=0}^{N-1} x^2[n] - 2A \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n + A^2 \sum_{n=0}^{N-1} \cos^2 2\pi f_0 n \right) \right]}_{g(\mathbf{T}(\mathbf{x}), \boldsymbol{\theta})} \cdot \underbrace{1}_{h(\mathbf{x})}$$

where

$$\mathbf{T}(\mathbf{x}) = \begin{bmatrix} \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \\ \sum_{n=0}^{N-1} x^2[n] \end{bmatrix}.$$

Hence, $\mathbf{T}(\mathbf{x})$ is a sufficient statistic for A and σ^2 , but only if f_0 is known (see also Problem 5.17). The next example continues this problem. $\diamond$

Example 5.10 - DC Level in White Noise with Unknown Noise Power

If

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 and the unknown parameters are A and σ^2 , we can use the results of the previous example to find a sufficient statistic of dimension 2. Letting the unknown vector parameter be $\theta = [A \sigma^2]^T$ and letting $f_0 = 0$ in the previous example, we have as our sufficient statistic

$$\mathbf{T}(\mathbf{x}) = \begin{bmatrix} \sum_{n=0}^{N-1} x[n] \\ \sum_{n=0}^{N-1} x^2[n] \end{bmatrix}.$$

Note that we had already observed that when σ^2 is known, $\sum_{n=0}^{N-1} x[n]$ is a sufficient statistic for A (see Example 5.2), and that when A is known ($A = 0$), $\sum_{n=0}^{N-1} x^2[n]$ is a sufficient statistic for σ^2 (see Example 5.3). In this example, the same statistics are jointly sufficient. This will not always be true, however. Since we have two parameters to be estimated and we have found two sufficient statistics, we should be able to find the MVU estimator for this example. $\diamond$

Before actually finding the MVU estimator for the previous example, we need to state the RBLs theorem for vector parameters.

Theorem 5.4 (Rao-Blackwell-Lehmann-Scheffe (Vector Parameter)) *If $\tilde{\theta}$ is an unbiased estimator of θ and $\mathbf{T}(\theta)$ is an $r \times 1$ sufficient statistic for θ , then $\hat{\theta} = E(\tilde{\theta}|\mathbf{T}(\mathbf{x}))$ is*

1. a valid estimator for θ (not dependent on θ)
2. unbiased
3. of lesser or equal variance than that of $\tilde{\theta}$ (each element of $\hat{\theta}$ has lesser or equal variance)

Additionally, if the sufficient statistic is complete, then $\hat{\theta}$ is the MVU estimator.

In the vector parameter case completeness means that for $\mathbf{v}(\mathbf{T})$, an arbitrary $r \times 1$ function of $\mathbf{T}$, if

$$E(\mathbf{v}(\mathbf{T})) = \int \mathbf{v}(\mathbf{T})p(\mathbf{T}; \theta) d\mathbf{T} = \mathbf{0} \quad \text{for all } \theta \quad (5.12)$$

then it must be true that

$$\mathbf{v}(\mathbf{T}) = \mathbf{0} \quad \text{for all } \mathbf{T}.$$

As before, this can be difficult to verify. Also, to be able to determine the MVU estimator directly from $\mathbf{T}(\mathbf{x})$ without having to evaluate $E(\hat{\theta}|\mathbf{T}(\mathbf{x}))$, the dimension of the sufficient statistic should be equal to the dimension of the unknown parameter or $r = p$. If this is satisfied, then we need only find an p -dimensional function $\mathbf{g}$ such that

$$E(\mathbf{g}(\mathbf{T})) = \theta$$

assuming that $\mathbf{T}$ is a complete sufficient statistic. We illustrate this approach by continuing Example 5.10.

Example 5.11 - DC Level in WGN with Unknown Noise Power (continued)

We first need to find the mean of the sufficient statistic

$$\mathbf{T}(\mathbf{x}) = \begin{bmatrix} T_1(\mathbf{x}) \\ T_2(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} \sum_{n=0}^{N-1} x[n] \\ \sum_{n=0}^{N-1} x^2[n] \end{bmatrix}.$$

Taking the expected value produces

$$E(\mathbf{T}(\mathbf{x})) = \begin{bmatrix} NA \\ NE(x^2[n]) \end{bmatrix} = \begin{bmatrix} NA \\ N(\sigma^2 + A^2) \end{bmatrix}.$$

Clearly, we could remove the bias of the first component of $\mathbf{T}(\mathbf{x})$ by dividing the statistic by N . However, this would not help the second component. It should be obvious that $T_2(\mathbf{x}) = \sum_{n=0}^{N-1} x^2[n]$ estimates the second moment and not the variance as desired. If we transform $\mathbf{T}(\mathbf{x})$ as

$$\begin{aligned} \mathbf{g}(\mathbf{T}(\mathbf{x})) &= \begin{bmatrix} \frac{1}{N}T_1(\mathbf{x}) \\ \frac{1}{N}T_2(\mathbf{x}) - \left[\frac{1}{N}T_1(\mathbf{x})\right]^2 \end{bmatrix} \\ &= \begin{bmatrix} \bar{x} \\ \frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \bar{x}^2 \end{bmatrix} \end{aligned}$$

then

$$E(\bar{x}) = A$$

and

$$E\left(\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \bar{x}^2\right) = \sigma^2 + A^2 - E(\bar{x}^2).$$

But we know that $\bar{x} \sim \mathcal{N}(A, \sigma^2/N)$, so that

$$E(\bar{x}^2) = A^2 + \sigma^2/N$$

and we have that

$$E\left(\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \bar{x}^2\right) = \sigma^2(1 - \frac{1}{N}) = \frac{N-1}{N}\sigma^2.$$

If we multiply this statistic by $N/(N-1)$, it will then be unbiased for σ^2 . Finally, the appropriate transformation is

$$\begin{aligned} \mathbf{g}(\mathbf{T}(\mathbf{x})) &= \begin{bmatrix} \frac{1}{N}T_1(\mathbf{x}) \\ \frac{1}{N-1} [T_2(\mathbf{x}) - N(\frac{1}{N}T_1(\mathbf{x}))^2] \end{bmatrix} \\ &= \begin{bmatrix} \bar{x} \\ \frac{1}{N-1} \left[\sum_{n=0}^{N-1} x^2[n] - N\bar{x}^2 \right] \end{bmatrix}. \end{aligned}$$

However, since

$$\begin{aligned} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 &= \sum_{n=0}^{N-1} x^2[n] - 2 \sum_{n=0}^{N-1} x[n]\bar{x} + N\bar{x}^2 \\ &= \sum_{n=0}^{N-1} x^2[n] - N\bar{x}^2, \end{aligned}$$

this may be written as

$$\hat{\theta} = \begin{bmatrix} \bar{x} \\ \frac{1}{N-1} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 \end{bmatrix}$$

which is the MVU estimator of $\theta = [A \sigma^2]^T$. The normalizing factor of $1/(N-1)$ for $\hat{\sigma}^2$ (the sample variance) is due to the one degree of freedom lost in estimating the mean. In this example $\hat{\theta}$ is not efficient. It can be shown [Hoel, Port, and Stone 1971] that $\bar{x}$

and $\hat{\sigma}^2$ are independent and that

$$\begin{aligned}\bar{x} &\sim \mathcal{N}\left(A, \frac{\sigma^2}{N}\right) \\ \frac{(N-1)\hat{\sigma}^2}{\sigma^2} &\sim \chi_{N-1}^2.\end{aligned}$$

Hence, the covariance matrix is

$$\mathbf{C}_{\hat{\theta}} = \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{2\sigma^4}{N-1} \end{bmatrix}$$

while the CRLB is (see Example 3.6)

$$\mathbf{I}^{-1}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{2\sigma^4}{N} \end{bmatrix}.$$

Thus, the MVU estimator could not have been obtained by examining the CRLB.

Finally, we should observe that if we had known beforehand or suspected that the sample mean and sample variance were the MVU estimators, then we could have reduced our work. The PDF

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right]$$

can be factored more directly by noting that

$$\begin{aligned}\sum_{n=0}^{N-1} (x[n] - A)^2 &= \sum_{n=0}^{N-1} (x[n] - \bar{x} + \bar{x} - A)^2 \\ &= \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 + 2(\bar{x} - A) \sum_{n=0}^{N-1} (x[n] - \bar{x}) + N(\bar{x} - A)^2.\end{aligned}$$

The middle term is zero, so that we have the factorization

$$p(\mathbf{x}; \boldsymbol{\theta}) = \underbrace{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \left[\sum_{n=0}^{N-1} (x[n] - \bar{x})^2 + N(\bar{x} - A)^2 \right] \right\}}_{g(\mathbf{T}'(\mathbf{x}), \boldsymbol{\theta})} \cdot \underbrace{1}_{h(\mathbf{x})}$$

where

$$\mathbf{T}'(\mathbf{x}) = \begin{bmatrix} \bar{x} \\ \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 \end{bmatrix}.$$

Dividing the second component by $(N-1)$ would produce $\hat{\theta}$. Of course $\mathbf{T}'(\mathbf{x})$ and $\mathbf{T}(\mathbf{x})$ are related by a one-to-one transformation, illustrating once again that the sufficient statistic is unique to within these transformations.

In asserting that $\hat{\theta}$ is the MVU estimator we have not verified the completeness of the statistic. Completeness follows because the Gaussian PDF is a special case of the vector exponential family of PDFs, which are known to be complete [Kendall and Stuart 1979] (see Problem 5.14 for the definition of the scalar exponential family of PDFs). $\diamond$

References

- Hoel, P.G., S.C. Port, C.J. Stone, *Introduction to Statistical Theory*, Houghton Mifflin, Boston, 1971.
- Hoskins, R.F., *Generalized Functions*, J. Wiley, New York, 1979.
- Kendall, Sir M., A. Stuart, *The Advanced Theory of Statistics*, Vol. 2, Macmillan, New York, 1979.
- Papoulis, A., *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, 1965.

Problems

- 5.1 For Example 5.3 prove that $\sum_{n=0}^{N-1} x^2[n]$ is a sufficient statistic for σ^2 by using the definition that $p(\mathbf{x} | \sum_{n=0}^{N-1} x^2[n] = T_0; \sigma^2)$ does not depend on σ^2 . Hint: Note that the PDF of $s = \sum_{n=0}^{N-1} x^2[n] / \sigma^2$ is a chi-squared distribution with N degrees of freedom or

$$p(s) = \begin{cases} \frac{1}{2^{\frac{N}{2}} \Gamma(\frac{N}{2})} \exp\left(-\frac{s}{2}\right) s^{\frac{N}{2}-1} & s > 0 \\ 0 & s < 0. \end{cases}$$

- 5.2 The IID observations $x[n]$ for $n = 0, 1, \dots, N-1$ have the Rayleigh PDF

$$p(x[n]; \sigma^2) = \begin{cases} \frac{x[n]}{\sigma^2} \exp\left(-\frac{1}{2} \frac{x^2[n]}{\sigma^2}\right) & x[n] > 0 \\ 0 & x[n] < 0. \end{cases}$$

Find a sufficient statistic for σ^2 .

- 5.3 The IID observations $x[n]$ for $n = 0, 1, \dots, N-1$ have the exponential PDF

$$p(x[n]; \lambda) = \begin{cases} \lambda \exp(-\lambda x[n]) & x[n] > 0 \\ 0 & x[n] < 0. \end{cases}$$

Find a sufficient statistic for λ .

- 5.4 The IID observations $x[n]$ for $n = 0, 1, \dots, N-1$ are distributed according to $\mathcal{N}(\theta, \theta)$, where $\theta > 0$. Find a sufficient statistic for θ .

- 5.5 The IID observations $x[n]$ for $n = 0, 1, \dots, N-1$ are distributed according to $\mathcal{U}[-\theta, \theta]$, where $\theta > 0$. Find a sufficient statistic for θ .
- 5.6 If $x[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$ are observed, where $w[n]$ is WGN with variance σ^2 , find the MVU estimator for σ^2 assuming that A is known. You may assume that the sufficient statistic is complete.
- 5.7 Consider the frequency estimation of a sinusoid embedded in WGN or

$$x[n] = \cos 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN of known variance σ^2 . Show that it is not possible to find a sufficient statistic for the frequency f_0 .

- 5.8 In a similar fashion to Problem 5.7 consider the estimation of the damping constant r for the data

$$x[n] = r^n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with known variance σ^2 . Again show that a sufficient statistic does not exist for r .

- 5.9 Assume that $x[n]$ is the result of a Bernoulli trial (a coin toss) with

$$\begin{aligned} \Pr\{x[n] = 1\} &= \theta \\ \Pr\{x[n] = 0\} &= 1 - \theta \end{aligned}$$

and that N IID observations have been made. Assuming the Neyman-Fisher factorization theorem holds for discrete random variables, find a sufficient statistic for θ . Then, assuming completeness, find the MVU estimator of θ .

- 5.10 For Example 5.4 the following estimator is proposed

$$\hat{\phi} = -\arctan \left[\frac{T_2(\mathbf{x})}{T_1(\mathbf{x})} \right].$$

Show that at high SNR or $\sigma^2 \rightarrow 0$, for which

$$\begin{aligned} T_1(\mathbf{x}) &\approx \sum_{n=0}^{N-1} A \cos(2\pi f_0 n + \phi) \cos 2\pi f_0 n \\ T_2(\mathbf{x}) &\approx \sum_{n=0}^{N-1} A \cos(2\pi f_0 n + \phi) \sin 2\pi f_0 n \end{aligned}$$

the estimator satisfies $\hat{\phi} \approx \phi$. Use any approximations you deem appropriate. Is the proposed estimator the MVU estimator?

- 5.11 For Example 5.2 it is desired to estimate $\theta = 2A + 1$ instead of A . Find the MVU estimator of θ . What if the parameter $\theta = A^3$ is desired? Hint: Reparameterize the PDF in terms of θ and use the standard approach.

5.12 This problem examines the idea that sufficient statistics are unique to within one-to-one transformations. Once again, consider Example 5.2 but with the sufficient statistics

$$T_1(\mathbf{x}) = \sum_{n=0}^{N-1} x[n]$$

$$T_2(\mathbf{x}) = \left(\sum_{n=0}^{N-1} x[n] \right)^3.$$

Find functions g_1 and g_2 so that the unbiased condition is satisfied or

$$E[g_1(T_1(\mathbf{x}))] = A$$

$$E[g_2(T_2(\mathbf{x}))] = A$$

and hence the MVU estimator. Is the transformation from T_1 to T_2 one-to-one? What would happen if the statistic $T_3(\mathbf{x}) = (\sum_{n=0}^{N-1} x[n])^2$ were considered? Could you find a function g_3 to satisfy the unbiased constraint?

5.13 In this problem we derive the MVU estimator for an example that we have not encountered previously. If N IID observations are made according to the PDF

$$p(x[n]; \theta) = \begin{cases} \exp[-(x[n] - \theta)] & x[n] > \theta \\ 0 & x[n] < \theta, \end{cases}$$

find the MVU estimator for θ . Note that θ represents the minimum value that $x[n]$ may attain. Assume that the sufficient statistic is complete.

5.14 For the random variable x consider the PDF

$$p(x; \theta) = \exp [A(\theta)B(x) + C(x) + D(\theta)].$$

A PDF of this type is said to belong to the *scalar exponential family* of PDFs. Many useful properties are associated with this family, and in particular, that of sufficient statistics, which is explored in this and the next problem. Show that the following PDFs are within this class.

a. Gaussian

$$p(x; \mu) = \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2}(x - \mu)^2 \right].$$

b. Rayleigh

$$p(x; \sigma^2) = \begin{cases} \frac{x}{\sigma^2} \exp \left[-\frac{1}{2} \frac{x^2}{\sigma^2} \right] & x > 0 \\ 0 & x < 0. \end{cases}$$

c. Exponential

$$p(x; \lambda) = \begin{cases} \lambda \exp[-\lambda x] & x > 0 \\ 0 & x < 0 \end{cases}$$

where $\lambda > 0$.

- 5.15 If we observe $x[n]$ for $n = 0, 1, \dots, N-1$, which are IID and whose PDF belongs to the exponential family of PDFs, show that

$$T(\mathbf{x}) = \sum_{n=0}^{N-1} B(x[n])$$

is a sufficient statistic for θ . Next apply this result to determine the sufficient statistic for the Gaussian, Rayleigh, and exponential PDFs in Problem 5.14. Finally, assuming that the sufficient statistics are complete, find the MVU estimator in each case (if possible). (You may want to compare your results to Example 5.2 and Problems 5.2 and 5.3, respectively.)

- 5.16 In this problem an example is given where there are fewer sufficient statistics than parameters to be estimated [Kendall and Stuart 1979]. Assume that $x[n]$ for $n = 0, 1, \dots, N-1$ are observed, where $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C})$ and

$$\boldsymbol{\mu} = \begin{bmatrix} N\mu \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

$$\mathbf{C} = \begin{bmatrix} N-1 + \sigma^2 & -\mathbf{1}^T \\ -\mathbf{1} & \mathbf{I} \end{bmatrix}.$$

The covariance matrix $\mathbf{C}$ has the dimensions

$$\begin{bmatrix} 1 \times 1 & 1 \times (N-1) \\ (N-1) \times 1 & (N-1) \times (N-1) \end{bmatrix}.$$

Show that with $\boldsymbol{\theta} = [\mu \sigma^2]^T$

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}} \sigma} \exp \left\{ -\frac{1}{2} \left[\frac{N^2}{\sigma^2} (\bar{x} - \mu)^2 + \sum_{n=1}^{N-1} x^2[n] \right] \right\}.$$

What are the sufficient statistics for $\boldsymbol{\theta}$? Hint: You will need to find the inverse and determinant of a partitioned matrix and to use Woodbury's identity.

- 5.17 Consider a sinusoid of known frequency embedded in WGN or

$$x[n] = A \cos 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 . Find the MVU estimator of

- the amplitude A , assuming σ^2 is known
- the amplitude A and noise variance σ^2 .

You may assume that the sufficient statistics are complete. Hint: The results of Examples 5.9 and 5.11 may prove helpful.

5.18 If $x[n]$ for $n = 0, 1, \dots, N-1$ are observed, where the samples are IID and distributed according to $\mathcal{U}[\theta_1, \theta_2]$, find a sufficient statistic for $\boldsymbol{\theta} = [\theta_1 \ \theta_2]^T$.

5.19 Recall the linear model in Chapter 4 in which the data are modeled as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ with σ^2 known. Find the MVU estimator of $\boldsymbol{\theta}$ by using the Neyman-Fisher and RBLS theorems. Assume the sufficient statistic is complete. Compare your result with that described in Chapter 4. Hint: First prove that the following identity holds:

$$(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}) = (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}) + (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T \mathbf{H}^T \mathbf{H} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})$$

where

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

Appendix 5A

Proof of Neyman-Fisher Factorization Theorem (Scalar Parameter)

Consider the joint PDF $p(\mathbf{x}, T(\mathbf{x}); \theta)$. In evaluating this PDF it should be noted that $T(\mathbf{x})$ is functionally dependent on $\mathbf{x}$. Hence, the joint PDF must be zero when evaluated at $\mathbf{x} = \mathbf{x}_0$, $T(\mathbf{x}) = T_0$, unless $T(\mathbf{x}_0) = T_0$. This is analogous to the situation where we have two random variables x and y with $x = y$ always. The joint PDF $p(x, y)$ is $p(x)\delta(y - x)$, a line impulse in the two-dimensional plane, and is a degenerate two-dimensional PDF. In our case we have $p(\mathbf{x}, T(\mathbf{x}) = T_0; \theta) = p(\mathbf{x}; \theta)\delta(T(\mathbf{x}) - T_0)$, which we will use in our proof. A second result that we will need is a shorthand way of representing the PDF of a function of several random variables. If $y = g(\mathbf{x})$, for $\mathbf{x}$ a vector random variable, the PDF of y can be written as

$$p(y) = \int p(\mathbf{x})\delta(y - g(\mathbf{x})) d\mathbf{x}. \quad (5A.1)$$

Using properties of the Dirac delta function [Hoskins 1979], this can be shown to be equivalent to the usual formula for transformation of random variables [Papoulis 1965]. If, for example, $\mathbf{x}$ is a scalar random variable, then

$$\delta(y - g(x)) = \sum_{i=1}^k \frac{\delta(x - x_i)}{\left| \frac{dg}{dx} \right|_{x=x_i}}$$

where $\{x_1, x_2, \dots, x_k\}$ are all the solutions of $y = g(x)$. Substituting this into (5A.1) and evaluating produces the usual formula.

We begin by proving that $T(\mathbf{x})$ is a sufficient statistic when the factorization holds.

$$\begin{aligned} p(\mathbf{x}|T(\mathbf{x}) = T_0; \theta) &= \frac{p(\mathbf{x}, T(\mathbf{x}) = T_0; \theta)}{p(T(\mathbf{x}) = T_0; \theta)} \\ &= \frac{p(\mathbf{x}; \theta)\delta(T(\mathbf{x}) - T_0)}{p(T(\mathbf{x}) = T_0; \theta)}. \end{aligned}$$

Using the factorization

$$p(\mathbf{x}|T(\mathbf{x}) = T_0; \theta) = \frac{g(T(\mathbf{x}) = T_0, \theta)h(\mathbf{x})\delta(T(\mathbf{x}) - T_0)}{p(T(\mathbf{x}) = T_0; \theta)}. \quad (5A.2)$$

But, from (5A.1),

$$p(T(\mathbf{x}) = T_0; \theta) = \int p(\mathbf{x}; \theta)\delta(T(\mathbf{x}) - T_0) d\mathbf{x}.$$

Again using the factorization

$$\begin{aligned} p(T(\mathbf{x}) = T_0; \theta) &= \int g(T(\mathbf{x}) = T_0, \theta)h(\mathbf{x})\delta(T(\mathbf{x}) - T_0) d\mathbf{x} \\ &= g(T(\mathbf{x}) = T_0, \theta) \int h(\mathbf{x})\delta(T(\mathbf{x}) - T_0) d\mathbf{x}. \end{aligned}$$

The latter step is possible because the integral is zero except over the surface in R^N for which $T(\mathbf{x}) = T_0$. Over this surface g is a constant. Using this in (5A.2) produces

$$p(\mathbf{x}|T(\mathbf{x}) = T_0; \theta) = \frac{h(\mathbf{x})\delta(T(\mathbf{x}) - T_0)}{\int h(\mathbf{x})\delta(T(\mathbf{x}) - T_0) d\mathbf{x}}$$

which does not depend on θ . Hence, we conclude that $T(\mathbf{x})$ is a sufficient statistic.

Next we prove that if $T(\mathbf{x})$ is a sufficient statistic, then the factorization holds. Consider the joint PDF

$$p(\mathbf{x}, T(\mathbf{x}) = T_0; \theta) = p(\mathbf{x}|T(\mathbf{x}) = T_0; \theta)p(T(\mathbf{x}) = T_0; \theta) \quad (5A.3)$$

and note that

$$p(\mathbf{x}, T(\mathbf{x}) = T_0; \theta) = p(\mathbf{x}; \theta)\delta(T(\mathbf{x}) - T_0).$$

Because $T(\mathbf{x})$ is a sufficient statistic, the conditional PDF does not depend on θ . Hence, we write it as $p(\mathbf{x}|T(\mathbf{x}) = T_0)$. Now since we are given $T(\mathbf{x}) = T_0$, which defines a surface in R^N , the conditional PDF is nonzero only on that surface. We can let $p(\mathbf{x}|T(\mathbf{x}) = T_0) = w(\mathbf{x})\delta(T(\mathbf{x}) - T_0)$, where

$$\int w(\mathbf{x})\delta(T(\mathbf{x}) - T_0) d\mathbf{x} = 1. \quad (5A.4)$$

Thus, substituting in (5A.3) produces

$$p(\mathbf{x}; \theta)\delta(T(\mathbf{x}) - T_0) = w(\mathbf{x})\delta(T(\mathbf{x}) - T_0)p(T(\mathbf{x}) = T_0; \theta). \quad (5A.5)$$

We can let

$$w(\mathbf{x}) = \frac{h(\mathbf{x})}{\int h(\mathbf{x})\delta(T(\mathbf{x}) - T_0) d\mathbf{x}}$$

so that (5A.4) is satisfied. Then, (5A.5) becomes

$$p(\mathbf{x}; \theta) \delta(T(\mathbf{x}) - T_0) = \frac{h(\mathbf{x}) \delta(T(\mathbf{x}) - T_0)}{\int h(\mathbf{x}) \delta(T(\mathbf{x}) - T_0) d\mathbf{x}} p(T(\mathbf{x}) = T_0; \theta)$$

or

$$p(\mathbf{x}; \theta) = g(T(\mathbf{x}) = T_0; \theta) h(\mathbf{x})$$

where

$$g(T(\mathbf{x}) = T_0; \theta) = \frac{p(T(\mathbf{x}) = T_0; \theta)}{\int h(\mathbf{x}) \delta(T(\mathbf{x}) - T_0) d\mathbf{x}}.$$

In this form we see how, in principle, the PDF of the sufficient statistic can be found based on the factorization since

$$p(T(\mathbf{x}) = T_0; \theta) = g(T(\mathbf{x}) = T_0; \theta) \int h(\mathbf{x}) \delta(T(\mathbf{x}) - T_0) d\mathbf{x}.$$

Appendix 5B

Proof of Rao-Blackwell-Lehmann-Scheffe Theorem (Scalar Parameter)

To prove (1) of Theorem 5.2, that $\hat{\theta}$ is a valid estimator of θ (not a function of θ), note that

$$\begin{aligned}\hat{\theta} &= E(\tilde{\theta}|T(\mathbf{x})) \\ &= \int \tilde{\theta}(\mathbf{x})p(\mathbf{x}|T(\mathbf{x}); \theta) d\mathbf{x}.\end{aligned}\tag{5B.1}$$

By the definition of a sufficient statistic, $p(\mathbf{x}|T(\mathbf{x}); \theta)$ does not depend on θ but only on $\mathbf{x}$ and $T(\mathbf{x})$. Therefore, after the integration is performed over $\mathbf{x}$, the result will be solely a function of T and therefore of $\mathbf{x}$ only.

To prove (2), that $\hat{\theta}$ is unbiased, we use (5B.1) and the fact that $\hat{\theta}$ is a function of T only.

$$\begin{aligned}E(\hat{\theta}) &= \int \int \tilde{\theta}(\mathbf{x})p(\mathbf{x}|T(\mathbf{x}); \theta) d\mathbf{x} p(T(\mathbf{x}); \theta) dT \\ &= \int \tilde{\theta}(\mathbf{x}) \int p(\mathbf{x}|T(\mathbf{x}); \theta) p(T(\mathbf{x}); \theta) dT d\mathbf{x} \\ &= \int \tilde{\theta}(\mathbf{x}) p(\mathbf{x}; \theta) d\mathbf{x} \\ &= E(\tilde{\theta}).\end{aligned}$$

But $\tilde{\theta}$ is unbiased by assumption, and therefore

$$E(\hat{\theta}) = E(\tilde{\theta}) = \theta.$$

To prove (3), that

$$\text{var}(\hat{\theta}) \leq \text{var}(\tilde{\theta})$$

we have

$$\begin{aligned}\text{var}(\check{\theta}) &= E[(\check{\theta} - E(\check{\theta}))^2] \\ &= E[(\check{\theta} - \hat{\theta} + \hat{\theta} - \theta)^2] \\ &= E[(\check{\theta} - \hat{\theta})^2] + 2E[(\check{\theta} - \hat{\theta})(\hat{\theta} - \theta)] + E[(\hat{\theta} - \theta)^2].\end{aligned}$$

The cross-term in the above expression, $E[(\check{\theta} - \hat{\theta})(\hat{\theta} - \theta)]$, is now shown to be zero. We know that $\hat{\theta}$ is solely a function of T . Thus, the expectation of the cross-term is with respect to the joint PDF of T and $\check{\theta}$ or

$$E_{T, \check{\theta}}[(\check{\theta} - \hat{\theta})(\hat{\theta} - \theta)] = E_T E_{\check{\theta}|T}[(\check{\theta} - \hat{\theta})(\hat{\theta} - \theta)].$$

But

$$E_{\check{\theta}|T}[(\check{\theta} - \hat{\theta})(\hat{\theta} - \theta)] = E_{\check{\theta}|T}[\check{\theta} - \hat{\theta}](\hat{\theta} - \theta)$$

and also

$$E_{\check{\theta}|T}[\check{\theta} - \hat{\theta}] = E_{\check{\theta}|T}(\check{\theta}|T) - \hat{\theta} = \hat{\theta} - \hat{\theta} = 0.$$

Thus, we have the desired result

$$\begin{aligned}\text{var}(\check{\theta}) &= E[(\check{\theta} - \hat{\theta})^2] + \text{var}(\hat{\theta}) \\ &\geq \text{var}(\hat{\theta}).\end{aligned}$$

Finally, that $\hat{\theta}$ is the MVU estimator if it is complete follows from the discussion given in Section 5.5.

Chapter 6

Best Linear Unbiased Estimators

6.1 Introduction

It frequently occurs in practice that the MVU estimator, if it exists, cannot be found. For instance, we may not know the PDF of the data or even be willing to assume a model for it. In this case our previous methods, which rely on the CRLB and the theory of sufficient statistics, cannot be applied. Even if the PDF is known, the latter approaches are not guaranteed to produce the MVU estimator. Faced with our inability to determine the optimal MVU estimator, it is reasonable to resort to a suboptimal estimator. In doing so, we are never sure how much performance we may have lost (since the minimum variance of the MVU estimator is unknown). However, if the variance of the suboptimal estimator can be ascertained and if it meets our system specifications, then its use may be justified as being adequate for the problem at hand. If its variance is too large, then we will need to look at other suboptimal estimators, hoping to find one that meets our specifications. A common approach is to restrict the estimator to be *linear* in the data and find the linear estimator that is *unbiased* and has *minimum variance*. As we will see, this estimator, which is termed the *best linear unbiased estimator* (BLUE), can be determined with knowledge of only the first and second moments of the PDF. Since complete knowledge of the PDF is not necessary, the BLUE is frequently more suitable for practical implementation.

6.2 Summary

The BLUE is based on the linear estimator defined in (6.1). If we constrain this estimator to be unbiased as in (6.2) and to minimize the variance of (6.3), then the BLUE is given by (6.5). The minimum variance of the BLUE is (6.6). To determine the BLUE requires knowledge of only the mean and covariance of the data. In the vector parameter case the BLUE is given by (6.16) and the covariance by (6.17). In either the scalar parameter or vector parameter case, if the data are Gaussian, then the BLUE is also the MVU estimator.

6.3 Definition of the BLUE

We observe the data set $\{x[0], x[1], \dots, x[N-1]\}$ whose PDF $p(\mathbf{x}; \theta)$ depends on an unknown parameter θ . The BLUE restricts the estimator to be linear in the data or

$$\hat{\theta} = \sum_{n=0}^{N-1} a_n x[n] \quad (6.1)$$

where the a_n 's are constants yet to be determined. (See also Problem 6.6 for a similar definition except for the addition of a constant.) Depending on the a_n 's chosen, we may generate a large number of different estimators of θ . However, the best estimator or BLUE is defined to be the one that is unbiased and has minimum variance. Before determining the a_n 's that yield the BLUE, some comments about the optimality of the BLUE are in order. Since we are restricting the class of estimators to be linear, the BLUE will be optimal (that is to say, the MVU estimator) only when the MVU estimator turns out to be linear. For example, for the problem of estimating the value of a DC level in WGN (see Example 3.3) the MVU estimator is the sample mean

$$\hat{\theta} = \bar{x} = \sum_{n=0}^{N-1} \frac{1}{N} x[n]$$

which is clearly linear in the data. Hence, if we restrict our attention to only linear estimators, then we will lose nothing since the MVU estimator is within this class. Figure 6.1a depicts this idea. On the other hand, for the problem of estimating the mean of uniformly distributed noise (see Example 5.8), the MVU estimator was found to be

$$\hat{\theta} = \frac{N+1}{2N} \max x[n]$$

which is nonlinear in the data. If we restrict our estimator to be linear, then the BLUE is the sample mean, as we will see shortly. The BLUE for this problem is suboptimal, as illustrated in Figure 6.1b. As further shown in Example 5.8, the difference in performance is substantial. Unfortunately, without knowledge of the PDF there is no way to determine the loss in performance by resorting to a BLUE.

Finally, for some estimation problems the use of a BLUE can be totally inappropriate. Consider the estimation of the power of WGN. It is easily shown that the MVU estimator is (see Example 3.6)

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$$

which is nonlinear in the data. If we force the estimator to be linear as per (6.1), so that

$$\hat{\sigma}^2 = \sum_{n=0}^{N-1} a_n x[n],$$




Figure 6.1 Optimality of BLUE

the expected value of the estimator becomes

$$E(\hat{\sigma}^2) = \sum_{n=0}^{N-1} a_n E(x[n]) = 0$$

since $E(x[n]) = 0$ for all n . Thus, we cannot even find a single linear estimator that is unbiased, let alone one that has minimum variance. Although the BLUE is unsuitable for this problem, a BLUE utilizing the *transformed data* $y[n] = x^2[n]$ would produce a viable estimator since for

$$\hat{\sigma}^2 = \sum_{n=0}^{N-1} a_n y[n] = \sum_{n=0}^{N-1} a_n x^2[n]$$

the unbiased constraint yields

$$E(\hat{\sigma}^2) = \sum_{n=0}^{N-1} a_n \sigma^2 = \sigma^2.$$

There are many values of the a_n 's that could satisfy this constraint. Can you guess what the a_n 's should be to yield the BLUE? (See also Problem 6.5 for an example of this data transformation approach.) Hence, with enough ingenuity the BLUE may still be used if the data are first transformed suitably.

6.4 Finding the BLUE

To determine the BLUE we constrain $\hat{\theta}$ to be linear and unbiased and then find the a_n 's to minimize the variance. The unbiased constraint, is from (6.1),

$$E(\hat{\theta}) = \sum_{n=0}^{N-1} a_n E(x[n]) = \theta. \quad (6.2)$$

The variance of $\hat{\theta}$ is

$$\text{var}(\hat{\theta}) = E \left[\left(\sum_{n=0}^{N-1} a_n x[n] - E \left(\sum_{n=0}^{N-1} a_n x[n] \right) \right)^2 \right].$$

But by using (6.2) and letting $\mathbf{a} = [a_0 \ a_1 \ \dots \ a_{N-1}]^T$ we have

$$\begin{aligned} \text{var}(\hat{\theta}) &= E \left[(\mathbf{a}^T \mathbf{x} - \mathbf{a}^T E(\mathbf{x}))^2 \right] \\ &= E \left[(\mathbf{a}^T (\mathbf{x} - E(\mathbf{x})))^2 \right] \\ &= E \left[\mathbf{a}^T (\mathbf{x} - E(\mathbf{x})) (\mathbf{x} - E(\mathbf{x}))^T \mathbf{a} \right] \\ &= \mathbf{a}^T \mathbf{C} \mathbf{a}. \end{aligned} \quad (6.3)$$

The vector $\mathbf{a}$ of weights is found by minimizing (6.3) subject to the constraint of (6.2). Before proceeding we need to assume some form for $E(x[n])$. In order to satisfy the unbiased constraint, $E(x[n])$ must be linear in θ or

$$E(x[n]) = s[n]\theta \quad (6.4)$$

where the $s[n]$'s are known. Otherwise, it may be impossible to satisfy the constraint. As an example, if $E(x[n]) = \cos \theta$, then the unbiased constraint is $\sum_{n=0}^{N-1} a_n \cos \theta = \theta$ for all θ . Clearly, there do not exist a_n 's that will satisfy the unbiased constraint. Note that if we write $x[n]$ as

$$x[n] = E(x[n]) + [x[n] - E(x[n])]$$

then by viewing $x[n] - E(x[n])$ as noise or $w[n]$ we have

$$x[n] = \theta s[n] + w[n].$$

The assumption of (6.4) means that the BLUE is applicable to amplitude estimation of known signals in noise. To generalize its use we require a nonlinear transformation of the data as already described.

With the assumption given by (6.4) we now summarize the estimation problem. To find the BLUE we need to minimize the variance

$$\text{var}(\hat{\theta}) = \mathbf{a}^T \mathbf{C} \mathbf{a}$$

subject to the unbiased constraint, which from (6.2) and (6.4) becomes

$$\begin{aligned}\sum_{n=0}^{N-1} a_n E(x[n]) &= \theta \\ \sum_{n=0}^{N-1} a_n s[n] \theta &= \theta \\ \sum_{n=0}^{N-1} a_n s[n] &= 1\end{aligned}$$

or

$$\mathbf{a}^T \mathbf{s} = 1$$

where $\mathbf{s} = [s[0] \ s[1] \ \dots \ s[N-1]]^T$. The solution to this minimization problem is derived in Appendix 6A as

$$\mathbf{a}_{\text{opt}} = \frac{\mathbf{C}^{-1} \mathbf{s}}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}$$

so that the BLUE is

$$\hat{\theta} = \frac{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{x}}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} \quad (6.5)$$

and has minimum variance

$$\text{var}(\hat{\theta}) = \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}. \quad (6.6)$$

Note from (6.4) that since $E(\mathbf{x}) = \theta \mathbf{s}$, the BLUE is unbiased

$$\begin{aligned}E(\hat{\theta}) &= \frac{\mathbf{s}^T \mathbf{C}^{-1} E(\mathbf{x})}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} \\ &= \frac{\mathbf{s}^T \mathbf{C}^{-1} \theta \mathbf{s}}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} \\ &= \theta.\end{aligned}$$

Also, as we asserted earlier, to determine the BLUE we only require knowledge only of

1. $\mathbf{s}$ or the scaled mean

2. $\mathbf{C}$, the covariance

or the first two moments but not the entire PDF. Some examples now follow.

Example 6.1 - DC Level in White Noise

If we observe

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is white noise with variance σ^2 (and of unspecified PDF), then the problem is to estimate A . Since $w[n]$ is not necessarily Gaussian, the noise samples may be

statistically dependent even though they are uncorrelated. Because $E(x[n]) = A$, we have from (6.4) that $s[n] = 1$ and therefore $s = 1$. From (6.5) the BLUE is

$$\begin{aligned}\hat{A} &= \frac{\mathbf{1}^T \frac{1}{\sigma^2} \mathbf{I} \mathbf{x}}{\mathbf{1}^T \frac{1}{\sigma^2} \mathbf{I} \mathbf{1}} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] = \bar{x}\end{aligned}$$

and has minimum variance, from (6.6),

$$\begin{aligned}\text{var}(\hat{A}) &= \frac{1}{\mathbf{1}^T \frac{1}{\sigma^2} \mathbf{1}} \\ &= \frac{\sigma^2}{N}.\end{aligned}$$

Hence, the sample mean is the BLUE independent of the PDF of the data. In addition, as already discussed, it is the MVU estimator for Gaussian noise. $\diamond$

Example 6.2 - DC Level in Uncorrelated Noise

Now let $w[n]$ be zero mean *uncorrelated* noise with $\text{var}(w[n]) = \sigma_n^2$ and repeat Example 6.1. As before, $s = 1$, and from (6.5) and (6.6)

$$\begin{aligned}\hat{A} &= \frac{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{x}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}} \\ \text{var}(\hat{A}) &= \frac{1}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}.\end{aligned}$$

The covariance matrix is

$$\mathbf{C} = \begin{bmatrix} \sigma_0^2 & 0 & \dots & 0 \\ 0 & \sigma_1^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{N-1}^2 \end{bmatrix},$$

and its inverse is

$$\mathbf{C}^{-1} = \begin{bmatrix} \frac{1}{\sigma_0^2} & 0 & \dots & 0 \\ 0 & \frac{1}{\sigma_1^2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{1}{\sigma_{N-1}^2} \end{bmatrix}.$$

Therefore,

$$\hat{A} = \frac{\sum_{n=0}^{N-1} \frac{x[n]}{\sigma_n^2}}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2}} \quad (6.7)$$

$$\text{var}(\hat{A}) = \frac{1}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2}}. \quad (6.8)$$

The BLUE weights those samples most heavily with smallest variances in an attempt to equalize the noise contribution from each sample. The denominator in (6.7) is the scale factor needed to make the estimator unbiased. See also Problem 6.2 for some further results. $\diamond$

In general, the presence of $\mathbf{C}^{-1}$ in the BLUE acts to prewhiten the data prior to averaging. This was previously encountered in Example 4.4 which discussed estimation of a DC level in colored *Gaussian* noise. Also, in that example the identical estimator for A was obtained. Because of the Gaussian noise assumption, however, the estimator could be said to be efficient and hence MVU. It is not just coincidental that the MVU estimator for Gaussian noise and the BLUE are identical for this problem. This is a general result which says that for estimation of the parameters of a *linear model* (see Chapter 4) the BLUE is identical to the MVU estimator for Gaussian noise. We will say more about this in the next section.

6.5 Extension to a Vector Parameter

If the parameter to be estimated is a $p \times 1$ vector parameter, then for the estimator to be linear in the data we require

$$\hat{\theta}_i = \sum_{n=0}^{N-1} a_{in} x[n] \quad i = 1, 2, \dots, p \quad (6.9)$$

where the a_{in} 's are weighting coefficients. In matrix form this is

$$\hat{\boldsymbol{\theta}} = \mathbf{A} \mathbf{x}$$

where $\mathbf{A}$ is a $p \times N$ matrix. In order for $\hat{\boldsymbol{\theta}}$ to be unbiased we require

$$E(\hat{\theta}_i) = \sum_{n=0}^{N-1} a_{in} E(x[n]) = \theta_i \quad i = 1, 2, \dots, p \quad (6.10)$$

or in matrix form

$$E(\hat{\theta}) = \mathbf{A}E(\mathbf{x}) = \theta. \quad (6.11)$$

Keep in mind that a linear estimator will be appropriate only for problems in which the unbiased constraint can be satisfied. From (6.11) we must have that

$$E(\mathbf{x}) = \mathbf{H}\theta \quad (6.12)$$

for $\mathbf{H}$ a known $N \times p$ matrix. In the scalar parameter case (see (6.4))

$$E(\mathbf{x}) = \underbrace{\begin{bmatrix} s[0] \\ s[1] \\ \vdots \\ s[N-1] \end{bmatrix}}_{\mathbf{H}} \theta$$

so that (6.12) is truly the vector parameter generalization. Now, substitution of (6.12) into (6.11) produces the unbiased constraint

$$\mathbf{A}\mathbf{H} = \mathbf{I}. \quad (6.13)$$

If we define $\mathbf{a}_i = [a_{i0} \ a_{i1} \ \dots \ a_{i(N-1)}]^T$, so that $\hat{\theta}_i = \mathbf{a}_i^T \mathbf{x}$, the unbiased constraint may be rewritten for each $\mathbf{a}_i$ by noting that

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \vdots \\ \mathbf{a}_p^T \end{bmatrix}$$

and letting $\mathbf{h}_i$ denote the i th column of $\mathbf{H}$, so that

$$\mathbf{H} = [\mathbf{h}_1 \ \mathbf{h}_2 \ \dots \ \mathbf{h}_p].$$

With these definitions the unbiased constraint of (6.13) reduces to

$$\mathbf{a}_i^T \mathbf{h}_j = \delta_{ij} \quad i = 1, 2, \dots, p; j = 1, 2, \dots, p. \quad (6.14)$$

Also, the variance is

$$\text{var}(\hat{\theta}_i) = \mathbf{a}_i^T \mathbf{C} \mathbf{a}_i \quad (6.15)$$

using the results for the scalar case (see (6.3)). The BLUE for a vector parameter is found by minimizing (6.15) subject to the constraints of (6.14), repeating the minimization for each i . In Appendix 6B this minimization is carried out to yield the BLUE as

$$\hat{\theta} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (6.16)$$

and a covariance matrix of

$$\mathbf{C}_{\hat{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}. \quad (6.17)$$

The form of the BLUE is identical to the MVU estimator for the general linear model (see (4.25)). This extends the observations of Examples 6.1 and 6.2. That this should be the case follows from the definition of the general linear model

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$. With this model we are attempting to estimate $\boldsymbol{\theta}$, where $E(\mathbf{x}) = \mathbf{H}\boldsymbol{\theta}$. This is exactly the assumption made in (6.12). But from (4.25) the MVU estimator for Gaussian data is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$$

which is clearly linear in $\mathbf{x}$. Hence, restricting our estimator to be linear does not lead to a suboptimal estimator since the MVU estimator is within the linear class. We may conclude that *if the data are truly Gaussian, then the BLUE is also the MVU estimator.*

To summarize our discussions we now state the general BLUE for a vector parameter of a general linear model. In the present context the general linear model does not assume Gaussian noise. We refer to the data as having the general linear model *form*. The following theorem is termed the *Gauss-Markov theorem*.

Theorem 6.1 (Gauss-Markov Theorem) *If the data are of the general linear model form*

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \quad (6.18)$$

where $\mathbf{H}$ is a known $N \times p$ matrix, $\boldsymbol{\theta}$ is a $p \times 1$ vector of parameters to be estimated, and $\mathbf{w}$ is a $N \times 1$ noise vector with zero mean and covariance $\mathbf{C}$ (the PDF of $\mathbf{w}$ is otherwise arbitrary), then the BLUE of $\boldsymbol{\theta}$ is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (6.19)$$

and the minimum variance of $\hat{\theta}_i$ is

$$\text{var}(\hat{\theta}_i) = [(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}]_{ii}. \quad (6.20)$$

In addition, the covariance matrix of $\hat{\boldsymbol{\theta}}$ is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}. \quad (6.21)$$

Some properties of BLUEs are given in Problems 6.12 and 6.13. Further examples of the computation of BLUEs are given in Chapter 4.

6.6 Signal Processing Example

In the design of many signal processing systems we are given measurements that may not correspond to the “raw data.” For instance, in optical interferometry the input to the system is light, and the output is the photon count as measured by a photodetector [Chamberlain 1979]. The system is designed to mix the incoming light signal with a delayed version of itself (through spatially separated mirrors and a nonlinear device).




Figure 6.2 Localization of aircraft by time difference of arrival measurements

The data then is more nearly modeled as an autocorrelation function. The raw data, that is, the light intensity, is unavailable. From the given information it is desired to estimate the spectrum of the incoming light. To assume that the measurements are Gaussian would probably be unrealistic. This is because if a random process is Gaussian, then the autocorrelation data are certainly non-Gaussian [Anderson 1971]. Furthermore, the exact PDF is impossible to find in general, being mathematically intractable. Such a situation lends itself to the use of the BLUE. A second example is explored in some detail next.

Example 6.3 - Source Localization

A problem of considerable interest (especially to anyone who has traveled by airplane) is that of determining the position of a source based on the emitted signal of the source. In order to track the position of an aircraft we typically use several antennas. A common method as shown in Figure 6.2 is to base our estimate on the time difference of arrival (TDOA) measurements. In the figure the signal emitted by the aircraft is received at the same time at all three antennas, so that $t_1 = t_2 = t_3$. The TDOA $t_2 - t_1$ between antennas 1 and 2 is zero, as is the TDOA $t_3 - t_2$ between antennas 2 and 3. Hence, we may conclude that the ranges R_i to all three antennas are the same. The dashed line between antennas 1 and 2 indicates the possible positions of the aircraft in order for $R_1 = R_2$, and likewise for the dashed line between antennas 2 and 3. The intersection of these lines gives the aircraft position. In the more general situation for which the TDOAs are not zero, the dashed lines become hyperbolas, the intersection again giving the position. It is necessary to have at least three antennas for localization, and when noise is present, it is desirable to have more.

We now examine the localization problem using estimation theory [Lee 1975]. To do so we assume that N antennas have been placed at known locations and that the time of arrival measurements t_i for $i = 0, 1, \dots, N - 1$ are available. The problem is to estimate the source position (x_s, y_s) as illustrated in Figure 6.3. The arrival times are assumed to be corrupted by noise with zero mean and known covariance but otherwise




Figure 6.3 Source localization geometry

unknown PDF. For a signal emitted by the source at time $t = T_0$, the measurements are modeled by

$$t_i = T_0 + R_i/c + \epsilon_i \quad i = 0, 1, \dots, N-1 \quad (6.22)$$

where the ϵ_i 's are measurement noises and c denotes the propagation speed. The noise samples are assumed to be zero mean with variance σ^2 and uncorrelated with each other. *No assumptions are made about the PDF of the noise.* To proceed further we must relate the range from each antenna R_i to the unknown position $\theta = [x_s, y_s]^T$. Letting the position of the i th antenna be (x_i, y_i) , which is assumed to be known, we have that

$$R_i = \sqrt{(x_s - x_i)^2 + (y_s - y_i)^2}. \quad (6.23)$$

Substituting (6.23) into (6.22), we see that the model is nonlinear in the unknown parameters x_s and y_s . To apply the Gauss-Markov theorem we will assume that a nominal source position (x_n, y_n) is available. This nominal position, which is close to the true source position, could have been obtained from previous measurements. Such a situation is typical when the source is being tracked. Hence, to estimate the *new* source position we require an estimate of $\theta = [(x_s - x_n) \ (y_s - y_n)]^T = [\delta x_s \ \delta y_s]^T$, as shown in Figure 6.3. Denote the nominal ranges by R_{n_i} . We use a first-order Taylor expansion of R_i (considered as a two-dimensional function of x_s and y_s) about the nominal position $x_s = x_n, y_s = y_n$

$$R_i \approx R_{n_i} + \frac{x_n - x_i}{R_{n_i}} \delta x_s + \frac{y_n - y_i}{R_{n_i}} \delta y_s. \quad (6.24)$$

See also Chapter 8 for a more complete description of the linearization of a nonlinear model. With this approximation we have upon substitution in (6.22)

$$t_i = T_0 + \frac{R_{n_i}}{c} + \frac{x_n - x_i}{R_{n_i}c} \delta x_s + \frac{y_n - y_i}{R_{n_i}c} \delta y_s + \epsilon_i$$

which is now linear in the unknown parameters δx_s and δy_s . Alternatively, since as defined in Figure 6.3

$$\frac{x_n - x_i}{R_{n_i}} = \cos \alpha_i$$

$$\frac{y_n - y_i}{R_{n_i}} = \sin \alpha_i,$$

the model simplifies to

$$t_i = T_0 + \frac{R_{n_i}}{c} + \frac{\cos \alpha_i}{c} \delta x_s + \frac{\sin \alpha_i}{c} \delta y_s + \epsilon_i.$$

The term R_{n_i}/c is a *known* constant which can be incorporated into the measurement by letting (see also Problem 6.6)

$$\tau_i = t_i - \frac{R_{n_i}}{c}.$$

Hence, we have for our linear model

$$\tau_i = T_0 + \frac{\cos \alpha_i}{c} \delta x_s + \frac{\sin \alpha_i}{c} \delta y_s + \epsilon_i, \quad (6.25)$$

where the unknown parameters are $T_0, \delta x_s, \delta y_s$. By assuming the time T_0 when the source emits a signal is unknown, we are bowing to practical considerations. Knowledge of T_0 would require accurate clock synchronization between the source and the receiver, which for economic reasons is avoided. It is customary to consider time *difference* of arrivals or TDOA measurements to eliminate T_0 in (6.25). We generate the TDOAs as

$$\begin{aligned} \xi_1 &= \tau_1 - \tau_0 \\ \xi_2 &= \tau_2 - \tau_1 \\ &\vdots \\ \xi_{N-1} &= \tau_{N-1} - \tau_{N-2}. \end{aligned}$$

Then, from (6.25) we have as our final linear model

$$\xi_i = \frac{1}{c} (\cos \alpha_i - \cos \alpha_{i-1}) \delta x_s + \frac{1}{c} (\sin \alpha_i - \sin \alpha_{i-1}) \delta y_s + \epsilon_i - \epsilon_{i-1} \quad (6.26)$$

for $i = 1, 2, \dots, N-1$. An alternative but equivalent approach is to use (6.25) and estimate T_0 as well as the source position. Now we have reduced the estimation problem to that described by the Gauss-Markov theorem where

$$\begin{aligned} \boldsymbol{\theta} &= [\delta x_s \quad \delta y_s]^T \\ \mathbf{H} &= \frac{1}{c} \begin{bmatrix} \cos \alpha_1 - \cos \alpha_0 & \sin \alpha_1 - \sin \alpha_0 \\ \cos \alpha_2 - \cos \alpha_1 & \sin \alpha_2 - \sin \alpha_1 \\ \vdots & \vdots \\ \cos \alpha_{N-1} - \cos \alpha_{N-2} & \sin \alpha_{N-1} - \sin \alpha_{N-2} \end{bmatrix} \\ \mathbf{w} &= \begin{bmatrix} \epsilon_1 - \epsilon_0 \\ \epsilon_2 - \epsilon_1 \\ \vdots \\ \epsilon_{N-1} - \epsilon_{N-2} \end{bmatrix}. \end{aligned}$$




Figure 6.4 Example of source localization geometry for line array with minimum number of antennas

The noise vector $\mathbf{w}$ is zero mean but is no longer composed of uncorrelated random variables. To find the covariance matrix note that

$$\mathbf{w} = \underbrace{\begin{bmatrix} -1 & 1 & 0 & 0 & \dots & 0 \\ 0 & -1 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & -1 & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \epsilon_0 \\ \epsilon_1 \\ \vdots \\ \epsilon_{N-1} \end{bmatrix}}_{\boldsymbol{\epsilon}},$$

where $\mathbf{A}$ has dimension $(N-1) \times N$. Since the covariance matrix of $\boldsymbol{\epsilon}$ is $\sigma^2 \mathbf{I}$, we have that

$$\mathbf{C} = E[\mathbf{A}\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T\mathbf{A}^T] = \sigma^2\mathbf{A}\mathbf{A}^T.$$

From (6.19) the BLUE of the source position parameters is

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= (\mathbf{H}^T\mathbf{C}^{-1}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{C}^{-1}\boldsymbol{\xi} \\ &= [\mathbf{H}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{H}]^{-1}\mathbf{H}^T(\mathbf{A}\mathbf{A}^T)^{-1}\boldsymbol{\xi} \end{aligned} \quad (6.27)$$

and the minimum variance, is from (6.20),

$$\text{var}(\hat{\theta}_i) = \sigma^2 \left[\{ \mathbf{H}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{H} \}^{-1} \right]_{ii} \quad (6.28)$$

or the covariance matrix, is from (6.21),

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \sigma^2 [\mathbf{H}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{H}]^{-1}. \quad (6.29)$$

As an example, for the three antenna line array (the minimum number of antennas required) shown in Figure 6.4 we have that

$$\begin{aligned} \mathbf{H} &= \frac{1}{c} \begin{bmatrix} -\cos \alpha & 1 - \sin \alpha \\ -\cos \alpha & -(1 - \sin \alpha) \end{bmatrix} \\ \mathbf{A} &= \begin{bmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix}. \end{aligned}$$

After substituting in (6.29) we have for the covariance matrix

$$\mathbf{C}_{\hat{\theta}} = \sigma^2 c^2 \begin{bmatrix} \frac{1}{2 \cos^2 \alpha} & 0 \\ 0 & \frac{3/2}{(1 - \sin \alpha)^2} \end{bmatrix}.$$

For the best localization we would like α to be small. This is accomplished by making the spacing between antennas d large, so that the *baseline* of the array, which is the total length, is large. Note also that the localization accuracy is range dependent, with the best accuracy for short ranges or for α small.

References

- Anderson, T.W., *The Statistical Analysis of Time Series*, J. Wiley, New York, 1971.
 Chamberlain, J., *The Principles of Interferometric Spectroscopy*, J. Wiley, New York, 1979.
 Lee, H.B., "A Novel Procedure for Assessing the Accuracy of Hyperbolic Multilateration Systems,"
IEEE Trans. Aerosp. Electron. Syst., Vol. 11, pp. 2-15, Jan. 1975.

Problems

- 6.1 If $x[n] = Ar^n + w[n]$ for $n = 0, 1, \dots, N-1$, where A is an unknown parameter, r is a known constant, and $w[n]$ is zero mean white noise with variance σ^2 , find the BLUE of A and the minimum variance. Does the minimum variance approach zero as $N \rightarrow \infty$?
- 6.2 In Example 6.2 let the noise variances be given by $\sigma_n^2 = n+1$ and examine what happens to the variance of the BLUE as $N \rightarrow \infty$. Repeat for $\sigma_n^2 = (n+1)^2$ and explain your results.
- 6.3 Consider the estimation problem described in Example 6.2 but now assume that the noise samples are correlated with the covariance matrix

$$\mathbf{C} = \sigma^2 \begin{bmatrix} 1 & \rho & 0 & \cdots & 0 \\ \rho & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \rho & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & \rho \\ & & & & \rho & 1 \end{bmatrix}$$

where $|\rho| < 1$ and N , the dimension of the matrix, is assumed to be even. $\mathbf{C}$ is a block-diagonal matrix and so is easily inverted (see (Appendix 1)). Find the BLUE and its variance and interpret your results.

- 6.4 The observed samples $\{x[0], x[1], \dots, x[N-1]\}$ are IID according to the following PDFs:

a. Laplacian

$$p(x[n]; \mu) = \frac{1}{2} \exp[-|x[n] - \mu|]$$

b. Gaussian

$$p(x[n]; \mu) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(x[n] - \mu)^2\right].$$

Find the BLUE of the mean μ in both cases. What can you say about the MVU estimator for μ ?

- 6.5 The observed samples $\{x[0], x[1], \dots, x[N-1]\}$ are IID according to the lognormal PDF

$$p(x[n]; \theta) = \begin{cases} \frac{1}{\sqrt{2\pi}x[n]} \exp\left[-\frac{1}{2}(\ln x[n] - \theta)^2\right] & x[n] > 0 \\ 0 & x[n] < 0. \end{cases}$$

Prove that the mean is $\exp(\theta + 1/2)$ and therefore that the unbiased constraint cannot be satisfied. Using the transformation of random variables approach with $y[n] = \ln x[n]$, find the BLUE for θ .

- 6.6 In this problem we extend the scalar BLUE results. Assume that $E(x[n]) = \theta s[n] + \beta$, where θ is the unknown parameter to be estimated and β is a *known* constant. The data vector $\mathbf{x}$ has covariance matrix $\mathbf{C}$. We define a modified linear (actually affine) estimator for this problem as

$$\hat{\theta} = \sum_{n=0}^{N-1} a_n x[n] + b.$$

Prove that the BLUE is given by

$$\hat{\theta} = \frac{\mathbf{s}^T \mathbf{C}^{-1} (\mathbf{x} - \beta \mathbf{1})}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}.$$

Also, find the minimum variance.

- 6.7 Assume that $x[n] = As[n] + w[n]$ for $n = 0, 1, \dots, N-1$ are observed, where $w[n]$ is zero mean noise with covariance matrix $\mathbf{C}$ and $s[n]$ is a known signal. The amplitude A is to be estimated using a BLUE. Find the BLUE and discuss what happens if $\mathbf{s} = [s[0] s[1] \dots s[N-1]]^T$ is an eigenvector of $\mathbf{C}$. Also, find the minimum variance.
- 6.8 Continuing Problem 6.7, we can always represent the signal vector $\mathbf{s}$ as a linear combination of the eigenvectors of $\mathbf{C}$. This is because we can always find N eigenvectors that are linearly independent. Furthermore, it can be shown that

these eigenvectors are orthonormal due to the symmetric nature of $\mathbf{C}$. Hence, an orthogonal representation of the signal is

$$\mathbf{s} = \sum_{i=0}^{N-1} \alpha_i \mathbf{v}_i$$

where $\{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{N-1}\}$ are the orthonormal eigenvectors of $\mathbf{C}$. Prove that the minimum variance of the BLUE is

$$\text{var}(\hat{A}) = \frac{1}{\sum_{i=0}^{N-1} \frac{\alpha_i^2}{\lambda_i}}$$

where λ_i is the eigenvalue of $\mathbf{C}$ corresponding to the eigenvector $\mathbf{v}_i$. Next show that the signal energy $\mathcal{E} = \mathbf{s}^T \mathbf{s}$ is given by $\sum_{i=0}^{N-1} \alpha_i^2$. Finally, prove that if the energy is constrained to be some value $\mathcal{E}_0$, then the smallest possible variance of the BLUE is obtained by choosing the signal

$$\mathbf{s} = c \mathbf{v}_{\min}$$

where $\mathbf{v}_{\min}$ is the eigenvector of $\mathbf{C}$ corresponding to the minimum eigenvalue and c is chosen to satisfy the energy constraint $\mathcal{E}_0 = \mathbf{s}^T \mathbf{s} = c^2$. Assume that the eigenvalues of $\mathbf{C}$ are distinct. Explain why this result makes sense.

6.9 In an on-off keyed (OOK) communication system we transmit one of two signals

$$s_0(t) = 0 \quad 0 \leq t \leq T$$

to represent a binary 0 or

$$s_1(t) = A \cos 2\pi f_1 t \quad 0 \leq t \leq T$$


to represent a binary 1. The amplitude A is assumed to be positive. To determine which bit was transmitted we sample the received waveform to produce

$$x[n] = s_i[n] + w[n] \quad n = 0, 1, \dots, N-1.$$

If a 0 was transmitted, we would have $s_i[n] = 0$, and if a 1 was transmitted, we would have $s_i[n] = A \cos 2\pi f_1 n$ (assuming a sampling rate of 1 sample/s). The noise samples given by $w[n]$ are used to model channel noise. A simple receiver might estimate the sinusoidal amplitude A and decide a 1 was sent if $\hat{A} > \gamma$, and a 0 if $\hat{A} < \gamma$. Hence, an equivalent problem is to estimate the amplitude A for the received data

$$x[n] = A \cos 2\pi f_1 n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is zero mean noise with covariance matrix $\mathbf{C} = \sigma^2 \mathbf{I}$. Find the BLUE for this problem and interpret the resultant detector. Find the best frequency in the range $0 \leq f_1 < 1/2$ to use at the transmitter.



$$H(z) = \sum_{k=0}^{N-1} h[k]z^{-k}$$

Figure 6.5 FIR filter estimator for Problem 6.11

- 6.10** We continue Problem 6.9 by examining the question of signal selection for an OOK system in the presence of *colored* noise. Let the noise $w[n]$ be a zero mean WSS random process with ACF

$$r_{ww}[k] = \begin{cases} 1.81 & k = 0 \\ 0 & k = 1 \\ 0.9 & k = 2 \\ 0 & k \geq 3. \end{cases}$$

Find the PSD and plot it for frequencies $0 \leq f \leq 1/2$. As in the previous problem, find the frequency which yields the smallest value of the BLUE variance for $N = 50$. Explain your results. Hint: You will need a computer to do this.

- 6.11** Consider the problem of estimating a DC level in WSS noise or given

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is a zero mean WSS random process with ACF $r_{ww}[k]$, estimate A . It is proposed to estimate A by using the output of the FIR filter shown in Figure 6.5 at time $n = N-1$. Note that the estimator is given by

$$\hat{A} = \sum_{k=0}^{N-1} h[k]x[N-1-k].$$

The input $x[n]$ is assumed to be zero for $n < 0$. To obtain a good estimator we wish to have the filter pass the DC signal A but block the noise $w[n]$. Hence, we choose $H(\exp(j0)) = 1$ as a constraint and minimize the noise power at the output at time $n = N-1$ by our choice of the FIR filter coefficients $h[k]$. Find the optimal filter coefficients and the minimum noise power at the output of the optimal filter. Explain your results.

- 6.12** Prove that the BLUE commutes over linear (actually affine) transformations of θ . Thus, if we wish to estimate

$$\alpha = B\theta + b$$

where B is a known $p \times p$ invertible matrix and b is a known $p \times 1$ vector, prove that the BLUE is given by

$$\hat{\alpha} = B\hat{\theta} + b$$

where $\hat{\theta}$ is the BLUE for θ . Assume that $\mathbf{x} = \mathbf{H}\theta + \mathbf{w}$, where $E(\mathbf{w}) = \mathbf{0}$ and $E(\mathbf{w}\mathbf{w}^T) = \mathbf{C}$. Hint: Replace θ in the data model by α .

- 6.13** In this problem it is shown that the BLUE is identical to the *weighted least squares estimator* to be discussed in Chapter 8. In particular, the weighted least squares estimator is found by minimizing the criterion

$$J = (\mathbf{x} - \mathbf{H}\theta)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta).$$

Prove that the θ that minimizes J is the BLUE.

- 6.14** A noise process is composed of IID zero mean random variables with PDF

$$p(w[n]) = \frac{1-\epsilon}{\sqrt{2\pi\sigma_B^2}} \exp\left[-\frac{1}{2}\left(\frac{w^2[n]}{\sigma_B^2}\right)\right] + \frac{\epsilon}{\sqrt{2\pi\sigma_I^2}} \exp\left[-\frac{1}{2}\left(\frac{w^2[n]}{\sigma_I^2}\right)\right]$$

where $0 < \epsilon < 1$. Such a PDF is called a Gaussian mixture. It is used to model noise that is Gaussian with variance σ_B^2 for $100(1-\epsilon)\%$ of the time and Gaussian with variance σ_I^2 the remaining time. Typically, $\sigma_I^2 \gg \sigma_B^2$ and $\epsilon \ll 1$, so that the noise is predominately due to *background* noise with variance σ_B^2 but also contains occasional high level events or *interferences* modeled as Gaussian noise with variance σ_I^2 . Show that the variance of this PDF is

$$\sigma^2 = (1-\epsilon)\sigma_B^2 + \epsilon\sigma_I^2.$$

If $\{w^2[0], w^2[1], \dots, w^2[N-1]\}$ is considered to be the data and σ_B^2, ϵ are assumed known, find the BLUE of σ_I^2 . Hint: Use the results of Problem 6.12.

- 6.15** For the general linear model

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{s} + \mathbf{w}$$

where $\mathbf{s}$ is a known $N \times 1$ vector and $E(\mathbf{w}) = \mathbf{0}$, $E(\mathbf{w}\mathbf{w}^T) = \mathbf{C}$, find the BLUE.

- 6.16** In implementing the BLUE for the linear model described in Theorem 6.1, we may unknowingly use the wrong covariance matrix to form the estimate

$$\hat{\theta} = (\mathbf{H}^T \hat{\mathbf{C}}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \hat{\mathbf{C}}^{-1} \mathbf{x}.$$

To examine the effect of this modeling error assume that we wish to estimate A as in Example 6.2. Find the variance of this estimator and that of the BLUE for the case where $N = 2$ and

$$\mathbf{C} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$\hat{\mathbf{C}} = \begin{bmatrix} 1 & 0 \\ 0 & \alpha \end{bmatrix}.$$

Compare the variances for $\alpha \rightarrow 0$, $\alpha = 1$, and $\alpha \rightarrow \infty$.

Appendix 6A

Derivation of Scalar BLUE

To minimize $\text{var}(\hat{\theta}) = \mathbf{a}^T \mathbf{C} \mathbf{a}$ subject to the constraint $\mathbf{a}^T \mathbf{s} = 1$ we use the method of Lagrangian multipliers. The Lagrangian function J becomes

$$J = \mathbf{a}^T \mathbf{C} \mathbf{a} + \lambda(\mathbf{a}^T \mathbf{s} - 1).$$

Using (4.3), the gradient with respect to $\mathbf{a}$ is

$$\frac{\partial J}{\partial \mathbf{a}} = 2\mathbf{C}\mathbf{a} + \lambda\mathbf{s}.$$

Setting this equal to the zero vector and solving produces

$$\mathbf{a} = -\frac{\lambda}{2}\mathbf{C}^{-1}\mathbf{s}.$$

The Lagrangian multiplier λ is found using the constraint equation as

$$\mathbf{a}^T \mathbf{s} = -\frac{\lambda}{2}\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s} = 1$$

or

$$-\frac{\lambda}{2} = \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}$$

so that the gradient is zero with the constraint satisfied for

$$\mathbf{a}_{\text{opt}} = \frac{\mathbf{C}^{-1} \mathbf{s}}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}.$$

The variance for this value of $\mathbf{a}$ is found from

$$\begin{aligned} \text{var}(\hat{\theta}) &= \mathbf{a}_{\text{opt}}^T \mathbf{C} \mathbf{a}_{\text{opt}} \\ &= \frac{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{C} \mathbf{C}^{-1} \mathbf{s}}{(\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s})^2} \\ &= \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}}. \end{aligned}$$

That $\mathbf{a}_{\text{opt}}$ is indeed the global minimum and not just a local minimum is verified by considering the function

$$G = (\mathbf{a} - \mathbf{a}_{\text{opt}})^T \mathbf{C} (\mathbf{a} - \mathbf{a}_{\text{opt}}).$$

Expanding G , we have

$$\begin{aligned} G &= \mathbf{a}^T \mathbf{C} \mathbf{a} - 2\mathbf{a}_{\text{opt}}^T \mathbf{C} \mathbf{a} + \mathbf{a}_{\text{opt}}^T \mathbf{C} \mathbf{a}_{\text{opt}} \\ &= \mathbf{a}^T \mathbf{C} \mathbf{a} - 2 \frac{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{C} \mathbf{a}}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} + \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} \\ &= \mathbf{a}^T \mathbf{C} \mathbf{a} - \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}} \end{aligned}$$

where we have used the constraint equation. Hence,

$$\mathbf{a}^T \mathbf{C} \mathbf{a} = (\mathbf{a} - \mathbf{a}_{\text{opt}})^T \mathbf{C} (\mathbf{a} - \mathbf{a}_{\text{opt}}) + \frac{1}{\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}},$$

which is uniquely minimized for $\mathbf{a} = \mathbf{a}_{\text{opt}}$. This is because the positive definite property of $\mathbf{C}$ makes the first term on the right-hand side greater than zero for all $(\mathbf{a} - \mathbf{a}_{\text{opt}}) \neq \mathbf{0}$.

Appendix 6B

Derivation of Vector BLUE

To derive the BLUE for a vector parameter we need to minimize

$$\text{var}(\hat{\theta}_i) = \mathbf{a}_i^T \mathbf{C} \mathbf{a}_i$$

for $i = 1, 2, \dots, p$ subject to the constraints

$$\mathbf{a}_i^T \mathbf{h}_j = \delta_{ij} \quad i = 1, 2, \dots, p; j = 1, 2, \dots, p.$$

This problem is much the same as that of the scalar BLUE (see Appendix 6A) except for the additional constraints on $\mathbf{a}_i$. Now we have p constraints on $\mathbf{a}_i$ instead of just one. Since each $\mathbf{a}_i$ is free to assume any value, independently of the others, we actually have p separate minimization problems linked only by the constraints. Proceeding in a similar manner to that of Appendix 6A, we consider the Lagrangian function for $\mathbf{a}_i$.

$$J_i = \mathbf{a}_i^T \mathbf{C} \mathbf{a}_i + \sum_{j=1}^p \lambda_j^{(i)} (\mathbf{a}_i^T \mathbf{h}_j - \delta_{ij}).$$

Taking the gradient of the Lagrangian using (4.3)

$$\frac{\partial J_i}{\partial \mathbf{a}_i} = 2\mathbf{C} \mathbf{a}_i + \sum_{j=1}^p \lambda_j^{(i)} \mathbf{h}_j.$$

We now let $\boldsymbol{\lambda}_i = [\lambda_1^{(i)} \lambda_2^{(i)} \dots \lambda_p^{(i)}]^T$ and note that $\mathbf{H} = [\mathbf{h}_1 \mathbf{h}_2 \dots \mathbf{h}_p]$ to yield

$$\frac{\partial J_i}{\partial \mathbf{a}_i} = 2\mathbf{C} \mathbf{a}_i + \mathbf{H} \boldsymbol{\lambda}_i.$$

Setting the gradient equal to zero yields

$$\mathbf{a}_i = -\frac{1}{2} \mathbf{C}^{-1} \mathbf{H} \boldsymbol{\lambda}_i. \quad (6B.1)$$

To find the vector of Lagrangian multipliers we use the constraint equations

$$\mathbf{a}_i^T \mathbf{h}_j = \delta_{ij} \quad j = 1, 2, \dots, p$$

or in combined form

$$\begin{bmatrix} \mathbf{h}_1^T \\ \vdots \\ \mathbf{h}_{i-1}^T \\ \mathbf{h}_i^T \\ \mathbf{h}_{i+1}^T \\ \vdots \\ \mathbf{h}_p^T \end{bmatrix} \mathbf{a}_i = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Letting $\mathbf{e}_i$ denote the vector of all zeros except in the i th place, we have the constraint equation

$$\mathbf{H}^T \mathbf{a}_i = \mathbf{e}_i.$$

Using (6B.1), the constraint equations become

$$\mathbf{H}^T \mathbf{a}_i = -\frac{1}{2} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H} \boldsymbol{\lambda}_i = \mathbf{e}_i$$

so that assuming invertibility of $\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}$, the Lagrangian multiplier vector is

$$-\frac{1}{2} \boldsymbol{\lambda}_i = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{e}_i.$$

Using this in (6B.1) yields the final result

$$\mathbf{a}_{i_{\text{opt}}} = \mathbf{C}^{-1} \mathbf{H} (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{e}_i \quad (6B.2)$$

and a corresponding variance of

$$\begin{aligned} \text{var}(\hat{\theta}_i) &= \mathbf{a}_{i_{\text{opt}}}^T \mathbf{C} \mathbf{a}_{i_{\text{opt}}} \\ &= \mathbf{e}_i^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{C} \mathbf{C}^{-1} \mathbf{H} (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{e}_i \\ &= \mathbf{e}_i^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{e}_i. \end{aligned}$$

In like fashion to Appendix 6A, $\mathbf{a}_{i_{\text{opt}}}$ can be shown to produce the global minimum. We can express the vector BLUE in a more compact form by letting

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= \begin{bmatrix} \mathbf{a}_{1_{\text{opt}}}^T \mathbf{x} \\ \mathbf{a}_{2_{\text{opt}}}^T \mathbf{x} \\ \vdots \\ \mathbf{a}_{p_{\text{opt}}}^T \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{e}_1^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \\ \mathbf{e}_2^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \\ \vdots \\ \mathbf{e}_p^T (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{e}_1^T \\ \mathbf{e}_2^T \\ \vdots \\ \mathbf{e}_p^T \end{bmatrix} (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \\ &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \end{aligned}$$

since $[\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_p]^T$ is the identity matrix. Also, the covariance matrix of $\hat{\boldsymbol{\theta}}$ is

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = E \left[(\hat{\boldsymbol{\theta}} - E(\hat{\boldsymbol{\theta}}))(\hat{\boldsymbol{\theta}} - E(\hat{\boldsymbol{\theta}}))^T \right]$$

where

$$\begin{aligned} \hat{\boldsymbol{\theta}} - E(\hat{\boldsymbol{\theta}}) &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} (\mathbf{H} \boldsymbol{\theta} + \mathbf{w}) - E(\hat{\boldsymbol{\theta}}) \\ &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{w}. \end{aligned}$$

Thus,

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = E[(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{w} \mathbf{w}^T \mathbf{C}^{-1} \mathbf{H} (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}].$$

Since the covariance matrix of $\mathbf{w}$ is $\mathbf{C}$, we have

$$\begin{aligned} \mathbf{C}_{\hat{\boldsymbol{\theta}}} &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{C} \mathbf{C}^{-1} \mathbf{H} (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \\ &= (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \end{aligned}$$

with the minimum variances given by the diagonal elements of $\mathbf{C}_{\hat{\boldsymbol{\theta}}}$ or

$$\text{var}(\hat{\theta}_i) = \mathbf{e}_i^T \mathbf{C}_{\hat{\boldsymbol{\theta}}} \mathbf{e}_i = [(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}]_{ii}$$

as previously derived.

Chapter 7

Maximum Likelihood Estimation

7.1 Introduction

We now investigate an alternative to the MVU estimator, which is desirable in situations where the MVU estimator does not exist or cannot be found even if it does exist. This estimator, which is based on the maximum likelihood principle, is overwhelmingly the most popular approach to obtaining *practical* estimators. It has the distinct advantage of being a “turn-the-crank” procedure, allowing it to be implemented for complicated estimation problems. Additionally, for most cases of practical interest its performance is optimal for large enough data records. Specifically, it is approximately the MVU estimator due to its approximate efficiency. For these reasons almost all practical estimators are based on the maximum likelihood principle.

7.2 Summary

In Section 7.3 we discuss an estimation problem for which the CRLB cannot be achieved and hence an efficient estimator does not exist, and for which the sufficient statistic approach in Chapter 5 to finding the MVU estimator cannot be implemented. An estimator is examined, however, that is asymptotically (for large data records) efficient. This is the **maximum likelihood estimator (MLE)**, defined as the value of θ that maximizes the likelihood function. In general, as stated in Theorem 7.1, the MLE has the asymptotic properties of being unbiased, achieving the CRLB, and having a Gaussian PDF. Thus, it can be said to be asymptotically optimal. In Example 7.6 it is also shown that for signal in noise problems, the MLE achieves the CRLB for high SNRs. The MLE of θ can be used to find the MLE of a function of θ by invoking the invariance property as described by Theorem 7.2. When a closed form expression cannot be found for the MLE, a numerical approach employs either a grid search or an iterative maximization of the likelihood function. The iterative approaches, which are used only when a grid search is not practical, are the Newton-Raphson (7.29) and scoring (7.33) methods. Convergence to the MLE, however, is not guaranteed. The MLE for a vector parameter θ is the value maximizing the likelihood function, which is now a

function of the components of θ . The asymptotic properties of the vector parameter MLE are summarized in Theorem 7.3, and the invariance property in Theorem 7.4. For the linear model the MLE achieves the CRLB for finite data records. The resultant estimator is thus the usual MVU estimator as summarized in Theorem 7.5. Numerical determination of the vector parameter MLE is based on the Newton-Raphson (7.48) or scoring (7.50) approach. Another technique is the expectation-maximization (EM) algorithm of (7.56) and (7.57), which is iterative in nature. Finally, for WSS Gaussian random processes the MLE can be found approximately by maximizing the asymptotic log-likelihood function of (7.60). This approach is computationally less intensive than determination of the exact MLE.

7.3 An Example

As a rationale for why we might be interested in an approximately optimal estimator, we now examine an estimation problem for which there is no obvious way to find the MVU estimator. By resorting to the estimator based on the maximum likelihood principle, termed the *maximum likelihood estimator* (MLE), we can obtain an estimator that is approximately the MVU estimator. The nature of the approximation relies on the property that the MLE is asymptotically (for large data records) efficient.

Example 7.1 - DC Level in White Gaussian Noise - Modified

Consider the observed data set

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where A is an unknown level, which is assumed to be positive ($A > 0$), and $w[n]$ is WGN with unknown variance A . This problem differs from our usual one (see Example 3.3) in that the unknown parameter A is reflected in the mean and the variance. In searching for the MVU estimator of A we first determine the CRLB (see Chapter 3) to see if it is satisfied with equality. The PDF is

$$p(\mathbf{x}; A) = \frac{1}{(2\pi A)^{\frac{N}{2}}} \exp \left[-\frac{1}{2A} \sum_{n=0}^{N-1} (x[n] - A)^2 \right]. \quad (7.1)$$

Taking the derivative of the log-likelihood function (the logarithm of the PDF considered as a function of A), we have

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; A)}{\partial A} &= -\frac{N}{2A} + \frac{1}{A} \sum_{n=0}^{N-1} (x[n] - A) + \frac{1}{2A^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \\ &\stackrel{?}{=} I(A)(\hat{A} - A). \end{aligned}$$

It is certainly not obvious if the derivative of the log-likelihood function can be put in the required form. It appears from a casual observation that it cannot, and therefore,

an efficient estimator does not exist. We can still determine the CRLB for this problem to find that (see Problem 7.1)

$$\text{var}(\hat{A}) \geq \frac{A^2}{N(A + \frac{1}{2})}. \quad (7.2)$$

We next try to find the MVU estimator by resorting to the theory of sufficient statistics (see Chapter 5). Attempting to factor (7.1) into the form of (5.3), we note that

$$\frac{1}{A} \sum_{n=0}^{N-1} (x[n] - A)^2 = \frac{1}{A} \sum_{n=0}^{N-1} x^2[n] - 2N\bar{x} + NA$$

so that the PDF factors as

$$p(\mathbf{x}; A) = \underbrace{\frac{1}{(2\pi A)^{\frac{N}{2}}} \exp\left[-\frac{1}{2} \left(\frac{1}{A} \sum_{n=0}^{N-1} x^2[n] + NA\right)\right]}_{g\left(\sum_{n=0}^{N-1} x^2[n], A\right)} \underbrace{\exp(N\bar{x})}_{h(\mathbf{x})}.$$

Based on the Neyman-Fisher factorization theorem a single sufficient statistic for A is $T(\mathbf{x}) = \sum_{n=0}^{N-1} x^2[n]$. The next step is to find a function of the sufficient statistic that produces an unbiased estimator, assuming that $T(\mathbf{x})$ is a *complete* sufficient statistic. To do so we need to find a function g such that

$$E\left[g\left(\sum_{n=0}^{N-1} x^2[n]\right)\right] = A \quad \text{for all } A > 0.$$

Since

$$\begin{aligned} E\left[\sum_{n=0}^{N-1} x^2[n]\right] &= NE[x^2[n]] \\ &= N[\text{var}(x[n]) + E^2(x[n])] \\ &= N(A + A^2) \end{aligned}$$

it is not obvious how to choose g . We cannot simply scale the sufficient statistic to make it unbiased as we did in Example 5.8.

A second approach would be to determine the conditional expectation $E(\hat{A} | \sum_{n=0}^{N-1} x^2[n])$ where $\hat{A}$ is any unbiased estimator. As an example, if we were to choose the unbiased estimator $\hat{A} = x[0]$, then the MVU estimator would take the form

$$E\left(x[0] \left| \sum_{n=0}^{N-1} x^2[n] \right.\right). \quad (7.3)$$

Unfortunately, the evaluation of the conditional expectation appears to be a formidable task.

We have now exhausted our possible optimal approaches. This is not to say that we could not propose some estimators. One possibility considers A to be the mean, so that

$$\hat{A}_1 = \begin{cases} \bar{x} & \text{if } \bar{x} > 0 \\ 0 & \text{if } \bar{x} \leq 0 \end{cases}$$

since we know that $A > 0$. Another estimator considers A to be the variance, to yield

$$\hat{A}_2 = \frac{1}{N-1} \sum_{n=0}^{N-1} (x[n] - \hat{A}_1)^2$$

(see also Problem 7.2). However, these estimators cannot be claimed to be optimal in any sense. $\diamond$

Faced with our inability to find the MVU estimator we will propose an estimator that is *approximately optimal*. We claim that for large data records or as $N \rightarrow \infty$, the proposed estimator is efficient. This means that as $N \rightarrow \infty$

$$E(\hat{A}) \rightarrow A \quad (7.4)$$

$$\text{var}(\hat{A}) \rightarrow \text{CRLB}, \quad (7.5)$$

the CRLB being given by (7.2). An estimator $\hat{A}$ that satisfies (7.4) is said to be *asymptotically unbiased*. If, in addition, the estimator satisfies (7.5), then it is said to be *asymptotically efficient*. For finite data records, however, we can say nothing about its optimality (see Problem 7.4). Better estimators may exist, but finding them may not be easy!

Example 7.2 - DC Level in White Gaussian Noise - Modified (continued)

We propose the following estimator:

$$\hat{A} = -\frac{1}{2} + \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{4}}. \quad (7.6)$$

This estimator is biased since

$$\begin{aligned} E(\hat{A}) &= E\left(-\frac{1}{2} + \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{4}}\right) \\ &\neq -\frac{1}{2} + \sqrt{E\left(\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]\right) + \frac{1}{4}} \quad \text{for all } A \\ &= -\frac{1}{2} + \sqrt{A + A^2 + \frac{1}{4}} \\ &= A. \end{aligned}$$

It is nonetheless a reasonable estimator in that as $N \rightarrow \infty$, we have by the law of large numbers

$$\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] \rightarrow E(x^2[n]) = A + A^2$$

and therefore from (7.6)

$$\hat{A} \rightarrow A.$$

The estimator $\hat{A}$ is said to be a *consistent* estimator (see Problems 7.5 and 7.6). To find the mean and variance of $\hat{A}$ as $N \rightarrow \infty$ we use the statistical linearization argument described in Section 3.6. For this example the PDF of $\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$ will be concentrated about its mean, $A + A^2$, for large data records. This allows us to linearize the function given in (7.6) that transforms $\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$ into $\hat{A}$. Let g be that function, so that

$$\hat{A} = g(u)$$

where $u = \frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$, and therefore,

$$g(u) = -\frac{1}{2} + \sqrt{u + \frac{1}{4}}.$$

Linearizing about $u_0 = E(u) = A + A^2$, we have

$$g(u) \approx g(u_0) + \left. \frac{dg(u)}{du} \right|_{u=u_0} (u - u_0)$$

or

$$\hat{A} \approx A + \frac{\frac{1}{2}}{A + \frac{1}{2}} \left[\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - (A + A^2) \right]. \quad (7.7)$$

It now follows that the asymptotic mean is

$$E(\hat{A}) = A$$

so that $\hat{A}$ is asymptotically unbiased. Additionally, the asymptotic variance becomes, from (7.7),

$$\begin{aligned} \text{var}(\hat{A}) &= \left(\frac{\frac{1}{2}}{A + \frac{1}{2}} \right)^2 \text{var} \left[\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] \right] \\ &= \frac{\frac{1}{4}}{N(A + \frac{1}{2})^2} \text{var}(x^2[n]). \end{aligned}$$

But $\text{var}(x^2[n])$ can be shown to be $4A^3 + 2A^2$ (see Section 3.6), so that

$$\begin{aligned} \text{var}(\hat{A}) &= \frac{\frac{1}{4}}{N(A + \frac{1}{2})^2} 4A^2(A + \frac{1}{2}) \\ &= \frac{A^2}{N(A + \frac{1}{2})} \end{aligned}$$




Figure 7.1 Rationale for maximum likelihood estimator

which from (7.2) is the CRLB!

Summarizing our results, the proposed estimator given by (7.6) is asymptotically unbiased and asymptotically achieves the CRLB. Hence, it is *asymptotically efficient*. Furthermore, by the central limit theorem the random variable $\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$ is Gaussian as $N \rightarrow \infty$. Because $\hat{A}$ is a linear function of this Gaussian random variable for large data records (as per (7.7)), it too will have a Gaussian PDF. $\diamond$

The proposed estimator is termed the MLE. How it was found is discussed in the next section.

7.4 Finding the MLE

The MLE for a scalar parameter is defined to be the value of θ that maximizes $p(\mathbf{x}; \theta)$ for $\mathbf{x}$ fixed, i.e., the value that maximizes the likelihood function. The maximization is performed over the allowable range of θ . In the previous example this was $A > 0$. Since $p(\mathbf{x}; \theta)$ will also be a function of $\mathbf{x}$, the maximization produces a $\hat{\theta}$ that is a function of $\mathbf{x}$. The rationale for the MLE hinges on the observation that $p(\mathbf{x}; \theta) d\mathbf{x}$ gives the probability of observing $\mathbf{x}$ in a small volume for a given θ . In Figure 7.1 the PDF is evaluated for $\mathbf{x} = \mathbf{x}_0$ and then plotted versus θ . The value of $p(\mathbf{x} = \mathbf{x}_0; \theta) d\mathbf{x}$ for each θ tells us the probability of observing $\mathbf{x}$ in the region in R^N centered around $\mathbf{x}_0$ with volume $d\mathbf{x}$, assuming the given value of θ . If $\mathbf{x} = \mathbf{x}_0$ had indeed been observed, then inferring that $\theta = \theta_1$ would be unreasonable. Because if $\theta = \theta_1$, the probability of actually observing $\mathbf{x} = \mathbf{x}_0$ would be small. It is more “likely” that $\theta = \theta_2$ is the true value. It yields a high probability of observing $\mathbf{x} = \mathbf{x}_0$, the data that were *actually observed*. Thus, we choose $\hat{\theta} = \theta_2$ as our estimate or the value that maximizes $p(\mathbf{x} = \mathbf{x}_0; \theta)$ over the allowable range of θ . We now continue with our example.

Example 7.3 - DC Level in White Gaussian Noise - Modified (continued)

To actually find the MLE for this problem we first write the PDF from (7.1) as

$$p(\mathbf{x}; A) = \frac{1}{(2\pi A)^{\frac{N}{2}}} \exp \left[-\frac{1}{2A} \sum_{n=0}^{N-1} (x[n] - A)^2 \right].$$

Considering this as a function of A , it becomes the likelihood function. Differentiating the log-likelihood function, we have

$$\frac{\partial \ln p(\mathbf{x}; A)}{\partial A} = -\frac{N}{2A} + \frac{1}{A} \sum_{n=0}^{N-1} (x[n] - A) + \frac{1}{2A^2} \sum_{n=0}^{N-1} (x[n] - A)^2$$

and setting it equal to zero produces

$$\hat{A}^2 + \hat{A} - \frac{1}{N} \sum_{n=0}^{N-1} x^2[n] = 0.$$

Solving for $\hat{A}$ produces the two solutions

$$\hat{A} = -\frac{1}{2} \pm \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{4}}.$$

We choose the solution

$$\hat{A} = -\frac{1}{2} + \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{4}}$$

to correspond to the permissible range of A or $A > 0$. Note that $\hat{A} > 0$ for all possible values of $\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$. Finally, that $\hat{A}$ indeed maximizes the log-likelihood function is verified by examining the second derivative. $\diamond$

Not only does the maximum likelihood procedure yield an estimator that is asymptotically efficient, it also sometimes yields an efficient estimator for *finite* data records. This is illustrated in the following example.

Example 7.4 - DC Level in White Gaussian Noise

For the received data

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where A is the unknown level to be estimated and $w[n]$ is WGN with known variance σ^2 , the PDF is

$$p(\mathbf{x}; A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right].$$

Taking the derivative of the log-likelihood function produces

$$\frac{\partial \ln p(\mathbf{x}; A)}{\partial A} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)$$

which being set equal to zero yields the MLE

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n].$$

But we have already seen that the sample mean is an efficient estimator (see Example 3.3). Hence, the MLE is efficient. $\diamond$

This result is true in general. *If an efficient estimator exists, the maximum likelihood procedure will produce it.* See Problem 7.12 for an outline of the proof.

7.5 Properties of the MLE

The example discussed in Section 7.3 led to an estimator that for large data records (or asymptotically) was unbiased, achieved the CRLB, and had a Gaussian PDF. In summary, the MLE was distributed as

$$\hat{\theta} \stackrel{a}{\sim} \mathcal{N}(\theta, I^{-1}(\theta)) \quad (7.8)$$

where $\stackrel{a}{\sim}$ denotes “asymptotically distributed according to.” This result is quite general and forms the basis for claiming optimality of the MLE. Of course, in practice it is seldom known in advance how large N must be in order for (7.8) to hold. An analytical expression for the PDF of the MLE is usually impossible to derive. As an alternative means of assessing performance, a computer simulation is usually required, as discussed next.

Example 7.5 - DC Level in White Gaussian Noise - Modified (continued)

A computer simulation was performed to determine how large the data record had to be for the asymptotic results to apply. In principle the exact PDF of $\hat{A}$ (see (7.6)) could be found but would be extremely tedious. Using the Monte Carlo method (see Appendix 7A), $M = 1000$ realizations of $\hat{A}$ were generated for various data record lengths. The mean $E(\hat{A})$ and variance $\text{var}(\hat{A})$ were *estimated* by

$$\widehat{E(\hat{A})} = \frac{1}{M} \sum_{i=1}^M \hat{A}_i \quad (7.9)$$

$$\widehat{\text{var}(\hat{A})} = \frac{1}{M} \sum_{i=1}^M \left(\hat{A}_i - \widehat{E(\hat{A})} \right)^2. \quad (7.10)$$

For a value of A equal to 1 the results are shown in Table 7.1 for various data record lengths. Instead of the asymptotic variance or the CRLB of (7.2), we tabulate

$$N\text{var}(\hat{A}) = \frac{A^2}{A + \frac{1}{2}}$$

TABLE 7.1 Theoretical Asymptotic and Actual Mean and Variance for Estimator in Example 7.2

Data Record Length, N	Mean, $E(\hat{A})$	$N \times$ Variance, $N \text{var}(\hat{A})$
5	0.954	0.624
10	0.976	0.648
15	0.991	0.696
20	0.996 (0.987)	0.707 (0.669)
25	0.994	0.656
Theoretical asymptotic value	1	0.667

since this is independent of N , allowing us to check the convergence more readily. The theoretical asymptotic values of the mean and normalized variance are

$$\begin{aligned} E(\hat{A}) &= A = 1 \\ N\text{var}(\hat{A}) &= \frac{2}{3}. \end{aligned}$$

It is observed from Table 7.1 that the mean converges at about $N = 20$ samples, while the variance jumps around somewhat for $N \geq 15$ samples. The latter is due to the statistical fluctuations in estimating the variance via a computer simulation, as well as possible inaccuracies in the random number generator. To check this the number of realizations was increased to $M = 5000$ for a data record length of $N = 20$. This resulted in the mean and normalized variance shown in parentheses. The normalized variance is now nearly identical to its asymptotic value, whereas for some unknown reason (presumably the random number generator) the mean is off slightly from its asymptotic value. (See also Problem 9.8 for a more accurate formula for $E(\hat{A})$.)



Next, the PDF of $\hat{A}$ was determined using a Monte Carlo computer simulation. This was done for data record lengths of $N = 5$ and $N = 20$. According to (7.8), the asymptotic PDF is

$$\hat{A} \stackrel{a}{\sim} \mathcal{N}(A, I^{-1}(A)),$$

which for $A = 1$ becomes, upon using (7.2),

$$\hat{A} \stackrel{a}{\sim} \mathcal{N}\left(1, \frac{2/3}{N}\right).$$

In Figure 7.2 the theoretical PDF and estimated PDF or *histogram* (see Appendix 7A) are shown. To construct the histogram we used $M = 5000$ realizations of $\hat{A}$ and divided the horizontal axis into 100 cells or divisions. Note that for $N = 5$ the estimated PDF is somewhat displaced to the left, in accordance with the mean being too small (see Table 7.1). For $N = 20$, however, the match is better, although the estimated PDF still appears to be skewed to the left. Presumably for larger data records the asymptotic PDF will more closely match the true one. $\diamond$

(a) $N = 5$ (b) $N = 20$ **Figure 7.2** Theoretical PDF and histogram

In describing the performance of estimators whose PDFs cannot be determined analytically, we must frequently resort to computer simulations, as in the previous example. In practice, the computer has become the dominant tool for analyzing the performance of nonlinear estimators. For this reason it is important to be able to carry out such a simulation. In Appendix 7A a description of the computer methods used in the previous example is given. Of course, the subject of Monte Carlo computer methods for statistical evaluation warrants a more complete discussion. The interested reader should consult the following references: [Bendat and Piersol 1971, Schwartz and Shaw

1975]. Some practice in performing these simulations is provided by Problems 7.13 and 7.14. We now summarize the asymptotic properties of the MLE in a theorem.

Theorem 7.1 (Asymptotic Properties of the MLE) If the PDF $p(\mathbf{x}; \theta)$ of the data $\mathbf{x}$ satisfies some “regularity” conditions, then the MLE of the unknown parameter θ is asymptotically distributed (for large data records) according to

$$\hat{\theta} \approx \mathcal{N}(\theta, I^{-1}(\theta)) \quad (7.11)$$

where $I(\theta)$ is the Fisher information evaluated at the true value of the unknown parameter.

The regularity conditions require the existence of the derivatives of the log-likelihood function, as well as the Fisher information being nonzero, as described more fully in Appendix 7B. An outline of the proof for IID observations is also given there.

From the asymptotic distribution, the MLE is seen to be asymptotically unbiased and asymptotically attains the CRLB. It is therefore asymptotically efficient, and hence asymptotically optimal. Of course, in practice the key question is always How large does N have to be for the asymptotic properties to apply? Fortunately, for many cases of interest the data record lengths are not excessive, as illustrated by Example 7.5. Another example follows.

Example 7.6 - MLE of the Sinusoidal Phase

We now reconsider the problem in Example 3.4 in which we wish to estimate the phase ϕ of a sinusoid embedded in noise or

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 and the amplitude A and frequency f_0 are assumed to be known. We saw in Chapter 5 that no *single* sufficient statistic exists for this problem. The sufficient statistics were

$$\begin{aligned} T_1(\mathbf{x}) &= \sum_{n=0}^{N-1} x[n] \cos(2\pi f_0 n) \\ T_2(\mathbf{x}) &= \sum_{n=0}^{N-1} x[n] \sin(2\pi f_0 n). \end{aligned} \quad (7.12)$$

The MLE is found by maximizing $p(\mathbf{x}; \phi)$ or

$$p(\mathbf{x}; \phi) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi))^2 \right]$$

or, equivalently, by minimizing

$$J(\phi) = \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi))^2. \quad (7.13)$$

Differentiating with respect to ϕ produces

$$\frac{\partial J(\phi)}{\partial \phi} = 2 \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi)) A \sin(2\pi f_0 n + \phi)$$

and setting it equal to zero yields

$$\sum_{n=0}^{N-1} x[n] \sin(2\pi f_0 n + \hat{\phi}) = A \sum_{n=0}^{N-1} \sin(2\pi f_0 n + \hat{\phi}) \cos(2\pi f_0 n + \hat{\phi}). \quad (7.14)$$

But the right-hand side may be approximated since (see Problem 3.7)

$$\frac{1}{N} \sum_{n=0}^{N-1} \sin(2\pi f_0 n + \hat{\phi}) \cos(2\pi f_0 n + \hat{\phi}) = \frac{1}{2N} \sum_{n=0}^{N-1} \sin(4\pi f_0 n + 2\hat{\phi}) \approx 0 \quad (7.15)$$

for f_0 not near 0 or $1/2$. Thus, the left-hand side of (7.14) when divided by N and set equal to zero will produce an approximate MLE, which satisfies

$$\sum_{n=0}^{N-1} x[n] \sin(2\pi f_0 n + \hat{\phi}) = 0. \quad (7.16)$$

Upon expanding this we have

$$\sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n \cos \hat{\phi} = - \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \sin \hat{\phi}$$

or finally the MLE of phase is given approximately as

$$\hat{\phi} = - \arctan \frac{\sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n}{\sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n}. \quad (7.17)$$

It is interesting to note that the MLE is a function of the sufficient statistics. In hindsight, this should not be surprising if we keep in mind the Neyman-Fisher factorization theorem. In this example there are two sufficient statistics, effecting a factorization as

$$p(\mathbf{x}; \phi) = g(T_1(\mathbf{x}), T_2(\mathbf{x}), \phi) h(\mathbf{x}).$$

Clearly, maximizing $p(\mathbf{x}; \phi)$ is equivalent to maximizing g ($h(\mathbf{x})$ can always be chosen so that $h(\mathbf{x}) > 0$), and thus $\hat{\phi}$ must be a function of $T_1(\mathbf{x})$ and $T_2(\mathbf{x})$.

According to Theorem 7.1, the asymptotic PDF of the phase estimator is

$$\hat{\phi} \stackrel{a}{\sim} \mathcal{N}(\phi, I^{-1}(\phi)). \quad (7.18)$$

TABLE 7.2 Theoretical Asymptotic and Actual Mean and Variance for Estimator in Example 7.6

Data Record Length, N	Mean, $E(\hat{\phi})$	$N \times$ Variance, $N \text{var}(\hat{\phi})$
20	0.732	0.0978
40	0.746	0.108
60	0.774	0.110
80	0.789	0.0990
Theoretical asymptotic value	$\phi = 0.785$	$1/\eta = 0.1$

From Example 3.4

$$I(\phi) = \frac{NA^2}{2\sigma^2}$$

so that the asymptotic variance is

$$\text{var}(\hat{\phi}) = \frac{1}{N \frac{A^2}{2\sigma^2}} = \frac{1}{N\eta} \quad (7.19)$$

where $\eta = (A^2/2)/\sigma^2$ is the SNR. To determine the data record length for the asymptotic mean and variance to apply we performed a computer simulation using $A = 1$, $f_0 = 0.08$, $\phi = \pi/4$, and $\sigma^2 = 0.05$. The results are listed in Table 7.2. It is seen that the asymptotic mean and normalized variance are attained for $N = 80$. For shorter data records the estimator is considerably biased. Part of the bias is due to the assumption made in (7.15). The MLE given by (7.17) is actually valid only for large N . To find the exact MLE we would have to minimize J as given by (7.13) by evaluating it for all ϕ . This could be done by using a grid search to find the minimum. Also, observe from the table that the variance for $N = 20$ is below the CRLB. This is possible due to the bias of the estimator, thereby invalidating the CRLB which assumes an unbiased estimator.

Next we fixed the data record length at $N = 80$ and varied the SNR. Plots of the mean and variance versus the SNR, as well as the asymptotic values, are shown in Figures 7.3 and 7.4. As shown in Figure 7.3, the estimator attains the asymptotic mean above about -10 dB. In Figure 7.4 we have plotted $10 \log_{10} \text{var}(\hat{\phi})$. This has the desirable effect of causing the CRLB to be a straight line when plotted versus the SNR in dB. In particular, from (7.19) the asymptotic variance or CRLB is

$$\begin{aligned} 10 \log_{10} \text{var}(\hat{\phi}) &= 10 \log_{10} \frac{1}{N\eta} \\ &= -10 \log_{10} N - 10 \log_{10} \eta. \end{aligned}$$

Examining the results, we see a peculiar trend. For low SNRs the variance is higher than the CRLB. Only at higher SNRs is the CRLB attained. Hence, the required data record length for the asymptotic results to apply also depends on the SNR. To




Figure 7.3 Actual vs. asymptotic mean for phase estimator




Figure 7.4 Actual vs. asymptotic variance for phase estimator

understand why this occurs we plot typical realizations of the log-likelihood function for different SNRs. As seen in Figure 7.5 for a high SNR, the maximum is relatively stable from realization to realization, and hence the MLE exhibits a low variance. For lower SNRs, however, the effect of the increased noise is to cause other peaks to occur. Occasionally these peaks are larger than the peak near the true value, causing a large estimation error and ultimately a larger variance. These large error estimates are said to be *outliers* and cause the *threshold effect* seen in Figures 7.3 and 7.4. Nonlinear estimators nearly always exhibit this effect. $\diamond$




Figure 7.5 Typical realizations of log-likelihood function for phase

In summary, the asymptotic PDF of the MLE is valid for large enough data records. For signal in noise problems the CRLB may be attained even for short data records if the SNR is high enough. To see why this is so the phase estimator can be written from (7.17) as

$$\hat{\phi} = -\arctan \frac{\sum_{n=0}^{N-1} [A \cos(2\pi f_0 n + \phi) + w[n]] \sin 2\pi f_0 n}{\sum_{n=0}^{N-1} [A \cos(2\pi f_0 n + \phi) + w[n]] \cos 2\pi f_0 n}$$




Figure 7.6 Non-Gaussian PDF for Example 7.7

$$\approx -\arctan \frac{-\frac{NA}{2} \sin \phi + \sum_{n=0}^{N-1} w[n] \sin 2\pi f_0 n}{\frac{NA}{2} \cos \phi + \sum_{n=0}^{N-1} w[n] \cos 2\pi f_0 n}$$

where we have used the same type of approximation as in (7.15) and some standard trigonometric identities. Simplifying, we have

$$\hat{\phi} \approx \arctan \frac{\sin \phi - \frac{2}{NA} \sum_{n=0}^{N-1} w[n] \sin 2\pi f_0 n}{\cos \phi + \frac{2}{NA} \sum_{n=0}^{N-1} w[n] \cos 2\pi f_0 n}. \quad (7.20)$$

If the data record is large and/or the sinusoidal power is large, the noise terms will be small. *It is this condition, that the estimation error is small, that allows the MLE to attain its asymptotic distribution.* See also Problem 7.15 for a further discussion of this point.

In some cases the asymptotic distribution does not hold, no matter how large the data record and/or the SNR becomes. This tends to occur when the estimation error cannot be reduced due to a lack of averaging in the estimator. An example follows.

Example 7.7 - DC Level in Nonindependent Non-Gaussian Noise

Consider the observations

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where each sample of $w[n]$ has the PDF $p(w[n])$ as shown in Figure 7.6a. The PDF is symmetric about $w[n] = 0$ and has a maximum at $w[n] = 0$. Furthermore, we assume all the noise samples are equal or $w[0] = w[1] = \dots = w[N-1]$. In estimating A we need to consider only a single observation since all observations are identical. Utilizing

only $x[0]$, we first note that the PDF of $x[0]$ is a shifted version of $p(w[n])$, where the shift is A as shown in Figure 7.6b. This is because $p_{x[0]}(x[0]; A) = p_{w[0]}(x[0] - A)$. The MLE of A is the value that maximizes $p_{w[0]}(x[0] - A)$, which because the PDF of $w[0]$ has a maximum at $w[0] = 0$ becomes

$$\hat{A} = x[0].$$

This estimator has the mean

$$E(\hat{A}) = E(x[0]) = A$$

since the noise PDF is symmetric about $w[0] = 0$. The variance of $\hat{A}$ is the same as the variance of $x[0]$ or of $w[0]$. Hence,

$$\text{var}(\hat{A}) = \int_{-\infty}^{\infty} u^2 p_{w[0]}(u) du$$

while the CRLB, is from Problem 3.2,

$$\text{var}(\hat{A}) \geq \left[\int_{-\infty}^{\infty} \frac{\left(\frac{dp_{w[0]}(u)}{du} \right)^2}{p_{w[0]}(u)} du \right]^{-1}$$

and the two are not in general equal (see Problem 7.16). In this example, then, the estimation error does not decrease as the data record length increases but remains the same. Furthermore, the PDF of $\hat{A} = x[0]$ as shown in Figure 7.6b is a shifted version of $p(w[n])$, clearly not Gaussian. Finally, $\hat{A}$ is not even consistent as $N \rightarrow \infty$. $\diamond$

7.6 MLE for Transformed Parameters

In many instances we wish to estimate a function of θ , the parameter characterizing the PDF. For example, we may not be interested in the value of a DC level A in WGN but only in the power A^2 . In such a situation the MLE of A^2 is easily found from the MLE of A . Some examples illustrate how this is done.

Example 7.8 - Transformed DC Level in WGN

Consider the data

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 . We wish to find the MLE of $\alpha = \exp(A)$. The PDF is given as

$$p(\mathbf{x}; A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right] \quad -\infty < A < \infty \quad (7.21)$$

and is parameterized by the parameter $\theta = A$. However, since α is a one-to-one transformation of A , we can equivalently parameterize the PDF as

$$p_T(\mathbf{x}; \alpha) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - \ln \alpha)^2 \right] \quad \alpha > 0 \quad (7.22)$$

where the subscript T indicates that the PDF is parameterized according to the *transformed* parameter. Clearly, $p_T(\mathbf{x}; \alpha)$ is the PDF of the data set

$$x[n] = \ln \alpha + w[n] \quad n = 0, 1, \dots, N-1$$

and (7.21) and (7.22) are entirely equivalent. The MLE of α is found by maximizing (7.22) over α . Setting the derivative of $p_T(\mathbf{x}; \alpha)$ with respect to α equal to zero yields

$$\sum_{n=0}^{N-1} (x[n] - \ln \hat{\alpha}) \frac{1}{\hat{\alpha}} = 0$$

or

$$\hat{\alpha} = \exp(\bar{x}).$$

But $\bar{x}$ is just the MLE of A , so that

$$\hat{\alpha} = \exp(\hat{A}) = \exp(\hat{\theta}).$$

The MLE of the transformed parameter is found by substituting the MLE of the original parameter into the transformation. This property of the MLE is termed the ***invariance property***.

Example 7.9 - Transformed DC Level in WGN (Another Example)

Now consider the transformation $\alpha = A^2$ for the data set in the previous example. If we try to repeat the steps, we soon encounter a problem. Attempting to parameterize $p(\mathbf{x}; A)$ with respect to α , we find that

$$A = \pm\sqrt{\alpha}$$

since the transformation is not one-to-one. If we choose $A = \sqrt{\alpha}$, then some of the possible PDFs of (7.21) will be missing. We actually require two sets of PDFs

$$\begin{aligned} p_{T_1}(\mathbf{x}; \alpha) &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - \sqrt{\alpha})^2 \right] & \alpha \geq 0 \\ p_{T_2}(\mathbf{x}; \alpha) &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] + \sqrt{\alpha})^2 \right] & \alpha > 0 \end{aligned} \quad (7.23)$$

to characterize all possible PDFs. It is possible to find the MLE of α as the value of α that yields the maximum of $p_{T_1}(\mathbf{x}; \alpha)$ and $p_{T_2}(\mathbf{x}; \alpha)$ or

$$\hat{\alpha} = \arg \max_{\alpha} \{p_{T_1}(\mathbf{x}; \alpha), p_{T_2}(\mathbf{x}; \alpha)\}. \quad (7.24)$$

Alternatively, we can find the maximum in two steps as




Figure 7.7 Construction of modified likelihood function

1. For a given value of α , say α_0 , determine whether $p_{T_1}(\mathbf{x}; \alpha)$ or $p_{T_2}(\mathbf{x}; \alpha)$ is larger. If, for example,

$$p_{T_1}(\mathbf{x}; \alpha_0) > p_{T_2}(\mathbf{x}; \alpha_0),$$

then denote the value of $p_{T_1}(\mathbf{x}; \alpha_0)$ as $\bar{p}_T(\mathbf{x}; \alpha_0)$. Repeat for all $\alpha > 0$ to form $\bar{p}_T(\mathbf{x}; \alpha)$. (Note that $\bar{p}_T(\mathbf{x}; \alpha = 0) = p(\mathbf{x}; A = 0)$.)

2. The MLE is given as the α that maximizes $\bar{p}_T(\mathbf{x}; \alpha)$ over $\alpha \geq 0$.

This procedure is illustrated in Figure 7.7. The function $\bar{p}_T(\mathbf{x}; \alpha)$ can be thought of as a *modified likelihood function*, having been derived from the original likelihood function by transforming the value of A that yields the maximum value for a given α . In this example, for each α the possible values of A are $\pm\sqrt{\alpha}$. Now, from (7.24) the MLE $\hat{\alpha}$ is

$$\hat{\alpha} = \arg \max_{\alpha > 0} \{p(\mathbf{x}; \sqrt{\alpha}), p(\mathbf{x}; -\sqrt{\alpha})\}$$

$$\begin{aligned}
&= \left[\arg \max_{\sqrt{\alpha} \geq 0} \{p(\mathbf{x}; \sqrt{\alpha}), p(\mathbf{x}; -\sqrt{\alpha})\} \right]^2 \\
&= \left[\arg \max_{-\infty < A < \infty} p(\mathbf{x}; A) \right]^2 \\
&= \hat{A}^2 \\
&= \bar{x}^2
\end{aligned}$$

so that again the invariance property holds. The understanding is that $\hat{\alpha}$ maximizes the *modified* likelihood function $\bar{p}_T(\mathbf{x}; \alpha)$ since the standard likelihood function with α as a parameter cannot be defined. $\diamond$

We summarize the preceding discussion in the following theorem.

Theorem 7.2 (Invariance Property of the MLE) The MLE of the parameter $\alpha = g(\theta)$, where the PDF $p(\mathbf{x}; \theta)$ is parameterized by θ , is given by

$$\hat{\alpha} = g(\hat{\theta})$$

where $\hat{\theta}$ is the MLE of θ . The MLE of $\hat{\theta}$ is obtained by maximizing $p(\mathbf{x}; \theta)$. If g is not a one-to-one function, then $\hat{\alpha}$ maximizes the modified likelihood function $\bar{p}_T(\mathbf{x}; \alpha)$, defined as

$$\bar{p}_T(\mathbf{x}; \alpha) = \max_{\{\theta: \alpha = g(\theta)\}} p(\mathbf{x}; \theta).$$

We complete our discussion by giving another example.

Example 7.10 - Power of WGN in dB

We observe N samples of WGN with variance σ^2 whose power in dB is to be estimated. To do so we first find the MLE of σ^2 . Then, we use the invariance principle to find the power P in dB, which is defined as

$$P = 10 \log_{10} \sigma^2.$$

The PDF is given by

$$p(\mathbf{x}; \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right].$$

Differentiating the log-likelihood function produces

$$\begin{aligned}
\frac{\partial \ln p(\mathbf{x}; \sigma^2)}{\partial \sigma^2} &= \frac{\partial}{\partial \sigma^2} \left[-\frac{N}{2} \ln 2\pi - \frac{N}{2} \ln \sigma^2 - \frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right] \\
&= -\frac{N}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{n=0}^{N-1} x^2[n]
\end{aligned}$$




Figure 7.8 Grid search for MLE

and upon setting it equal to zero yields the MLE

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} x^2[n].$$

The MLE of the power in dB readily follows as

$$\begin{aligned} \hat{P} &= 10 \log_{10} \hat{\sigma}^2 \\ &= 10 \log_{10} \frac{1}{N} \sum_{n=0}^{N-1} x^2[n]. \end{aligned}$$

◇

7.7 Numerical Determination of the MLE

A distinct advantage of the MLE is that we can always find it for a given data set *numerically*. This is because the MLE is determined as the maximum of a known function, namely, the likelihood function. If, for example, the allowable values of θ lie in the interval $[a, b]$, then we need only maximize $p(\mathbf{x}; \theta)$ over that interval. The “safest” way to do this is to perform a grid search over the $[a, b]$ interval as shown in Figure 7.8. As long as the spacing between θ values is small enough, we are guaranteed to find the MLE for the given set of data. Of course, for a new set of data we will have to repeat the search since the likelihood function will undoubtedly change. If, however, the range of θ is not confined to a finite interval, as in estimating the variance of a noise process for which $\sigma^2 > 0$, then a grid search may not be computationally feasible. In such a case, we are forced to resort to iterative maximization procedures. Some typical ones are the Newton-Raphson method, the scoring approach, and the expectation-maximization algorithm. In general, these methods will produce the MLE if the initial guess is close to the true maximum. If not, convergence may not be attained, or only convergence to a local maximum. The difficulty with the use of these iterative methods is that in general we do not know beforehand if they will converge and, even if convergence is attained, whether the value produced is the MLE. An important distinction of our estimation problem which sets it apart from other maximization problems is that the function to be maximized is *not known a priori*. The likelihood function changes for each data set,

requiring the maximization of a *random function*. Nevertheless, these methods can at times produce good results. We now describe some of the more common ones. The interested reader is referred to [Bard 1974] for a more complete description of methods for nonlinear optimization as applied to estimation problems.

As a means of comparison, we will apply the methods to the following example.

Example 7.11 - Exponential in WGN

Consider the data

$$x[n] = r^n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 . The parameter r , the exponential factor, is to be estimated. Allowable values are $r > 0$. The MLE of r is the value that maximizes the likelihood function

$$p(\mathbf{x}; r) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - r^n)^2 \right]$$

or, equivalently, the value that minimizes

$$J(r) = \sum_{n=0}^{N-1} (x[n] - r^n)^2.$$

Differentiating $J(r)$ and setting it equal to zero produces

$$\sum_{n=0}^{N-1} (x[n] - r^n) nr^{n-1} = 0. \quad (7.25)$$

This is a nonlinear equation in r and cannot be solved directly. We will now consider the iterative methods of Newton-Raphson and scoring. $\diamond$

The iterative methods attempt to maximize the log-likelihood function by finding a zero of the derivative function. To do so the derivative is taken and set equal to zero, yielding

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = 0. \quad (7.26)$$

Then, the methods attempt to solve this equation iteratively. Let

$$g(\theta) = \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta}$$

and assume that we have an initial guess for the solution to (7.26). Call this guess θ_0 . Then, if $g(\theta)$ is approximately linear near θ_0 , we can approximate it by

$$g(\theta) \approx g(\theta_0) + \left. \frac{dg(\theta)}{d\theta} \right|_{\theta=\theta_0} (\theta - \theta_0) \quad (7.27)$$




Figure 7.9 Newton-Raphson method for finding zero of function

as shown in Figure 7.9. Next, we use (7.27) to solve for the zero θ_1 , so that upon setting $g(\theta_1)$ equal to zero and solving for θ_1 we have

$$\theta_1 = \theta_0 - \frac{g(\theta_0)}{\left. \frac{dg(\theta)}{d\theta} \right|_{\theta=\theta_0}}.$$

Again we linearize g but use the new guess, θ_1 , as our point of linearization and repeat the previous procedure to find the new zero. As shown in Figure 7.9, the sequence of guesses will converge to the true zero of $g(\theta)$. In general, the Newton-Raphson iteration finds the new guess, θ_{k+1} , based on the previous one, θ_k , using

$$\theta_{k+1} = \theta_k - \frac{g(\theta_k)}{\left. \frac{dg(\theta)}{d\theta} \right|_{\theta=\theta_k}}. \quad (7.28)$$

Note that at convergence $\theta_{k+1} = \theta_k$, and from (7.28) $g(\theta_k) = 0$, as desired. Since $g(\theta)$ is the derivative of the log-likelihood function, we find the MLE as

$$\theta_{k+1} = \theta_k - \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right]^{-1} \left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\theta_k}. \quad (7.29)$$

Several points need to be raised concerning the Newton-Raphson iterative procedure.

1. The iteration may not converge. This will be particularly evident when the second derivative of the log-likelihood function is small. In this case it is seen from (7.29) that the correction term may fluctuate wildly from iteration to iteration.
2. Even if the iteration converges, the point found may not be the global maximum but possibly only a local maximum or even a local minimum. Hence, to avoid these possibilities it is best to use several starting points and at convergence choose the one that yields the maximum. Generally, if the initial point is close to the global maximum, the iteration will converge to it. *The importance of a good initial guess cannot be overemphasized.* An illustration is given in Problem 7.18.

Applying the Newton-Raphson iteration to the problem in Example 7.11, we have

$$\begin{aligned}
 \frac{\partial \ln p(\mathbf{x}; r)}{\partial r} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - r^n) n r^{n-1} \\
 \frac{\partial^2 \ln p(\mathbf{x}; r)}{\partial r^2} &= \frac{1}{\sigma^2} \left[\sum_{n=0}^{N-1} n(n-1) x[n] r^{n-2} - \sum_{n=0}^{N-1} n(2n-1) r^{2n-2} \right] \\
 &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} n r^{n-2} [(n-1)x[n] - (2n-1)r^n]
 \end{aligned} \tag{7.30}$$

and thus the Newton-Raphson iteration becomes

$$r_{k+1} = r_k - \frac{\sum_{n=0}^{N-1} (x[n] - r_k^n) n r_k^{n-1}}{\sum_{n=0}^{N-1} n r_k^{n-2} [(n-1)x[n] - (2n-1)r_k^n]}. \tag{7.31}$$

A second common iterative procedure is the method of scoring. It recognizes that

$$\left. \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\theta_k} \approx -I(\theta_k) \tag{7.32}$$

where $I(\theta)$ is the Fisher information. Indeed, for IID samples we have

$$\begin{aligned}
 \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} &= \sum_{n=0}^{N-1} \frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \\
 &= N \frac{1}{N} \sum_{n=0}^{N-1} \frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \\
 &\approx NE \left[\frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right] \\
 &= -Ni(\theta) \\
 &= -I(\theta)
 \end{aligned}$$

by the law of large numbers. Presumably, replacing the second derivative by its expected value will increase the stability of the iteration. The method then becomes

$$\theta_{k+1} = \theta_k + I^{-1}(\theta) \left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\theta_k}. \tag{7.33}$$

This approach is termed the *method of scoring*. It too suffers from the same convergence problems as the Newton-Raphson iteration. As applied to the problem in Example 7.11,

TABLE 7.3 Sequence of Iterates for Newton-Raphson Method

Iteration	Initial Guess, r_0		
	0.8	0.2	1.2
1	0.723	0.799	1.187
2	0.638	0.722	1.174
3	0.561	0.637	1.161
4	0.510	0.560	1.148
5	0.494	0.510	1.136
6	0.493	0.494	1.123
7		0.493	1.111
8			1.098
9			1.086
10			1.074
$\vdots$			$\vdots$
29			0.493

it produces upon using (7.30)

$$I(\theta) = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} n^2 r^{2n-2}$$

so that

$$r_{k+1} = r_k + \frac{\sum_{n=0}^{N-1} (x[n] - r_k^n) n r_k^{n-1}}{\sum_{n=0}^{N-1} n^2 r_k^{2n-2}}. \quad (7.34)$$

As an example of the iterative approach, we implement the Newton-Raphson method for Example 7.11 using a computer simulation. Using $N = 50$, $r = 0.5$, and $\sigma^2 = 0.01$, we generated a realization of the process. For a particular outcome $-J(r)$ is plotted in Figure 7.10 for $0 < r < 1$. The peak of the function, which is the MLE, occurs at $r = 0.493$. It is seen that the function is fairly broad for $r < 0.5$ but rather sharp for larger values of r . In fact, for $r > 1$ it was difficult to plot due to the exponential signal r^n , which causes the function to become very large and negative. We applied the Newton-Raphson method as given by (7.31) using several initial guesses. The resulting iterates are listed in Table 7.3. For $r_0 = 0.8$ or $r_0 = 0.2$ the iteration quickly converged to the true maximum. However, for $r_0 = 1.2$ the convergence was much slower, although the true maximum was attained after 29 iterations. If the initial guess was less than about 0.18 or greater than about 1.2, the succeeding iterates exceeded 1 and kept




Figure 7.10 Function to be maximized for MLE

increasing, causing a computer overflow condition in the algorithm. In this case the Newton-Raphson method fails to converge. As expected, the method works well when the initial guess is close to the true MLE but may produce poor results otherwise.

A third method that has recently been proposed is the expectation-maximization algorithm. Unlike the first two methods, it requires some ingenuity on the part of the designer. Its advantage, however, is its guaranteed convergence to at least a local maximum under some mild conditions. Convergence to the global maximum is, as with all iterative schemes, not assured. The expectation-maximization algorithm, not being a “turn-the-crank” procedure, does not lend itself to all estimation problems. A particular class of problems for which it is useful involves a vector parameter, and so we defer our discussion until the next section.

7.8 Extension to a Vector Parameter

The MLE for a vector parameter θ is defined to be the value that maximizes the likelihood function $p(\mathbf{x}; \theta)$ over the allowable domain for θ . Assuming a differentiable likelihood function, the MLE is found from

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = 0. \quad (7.35)$$

If multiple solutions exist, then the one that maximizes the likelihood function is the MLE. We now find the MLE for the problem in Example 3.6.

Example 7.12 - DC Level in WGN

Consider the data

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 and the vector parameter $\theta = [A \sigma^2]^T$ is to be estimated. Then, from Example 3.6

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial A} &= \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n] - A) \\ \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \sigma^2} &= -\frac{N}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{n=0}^{N-1} (x[n] - A)^2. \end{aligned}$$

Solving for A from the first equation, we have the MLE

$$\hat{A} = \bar{x}.$$

Solving for σ^2 in the second equation and using $\hat{A} = \bar{x}$ produces

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2.$$

The MLE is

$$\hat{\theta} = \begin{bmatrix} \bar{x} \\ \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 \end{bmatrix}.$$

◇

For a vector parameter the asymptotic properties of the MLE are summarized in the following theorem.

Theorem 7.3 (Asymptotic Properties of the MLE (Vector Parameter)) *If the PDF $p(\mathbf{x}; \theta)$ of the data $\mathbf{x}$ satisfies some “regularity” conditions, then the MLE of the unknown parameter θ is asymptotically distributed according to*

$$\hat{\theta} \stackrel{\Delta}{\sim} \mathcal{N}(\theta, \mathbf{I}^{-1}(\theta)) \quad (7.36)$$

where $\mathbf{I}(\theta)$ is the Fisher information matrix evaluated at the true value of the unknown parameter.

The regularity conditions are similar to those discussed in Appendix 7B for the scalar parameter case. As an illustration, consider the previous example. It can be shown that [Hoel, Port, and Stone 1971]

$$\begin{aligned} \bar{x} &\sim \mathcal{N}(A, \sigma^2/N) \\ T = \sum_{n=0}^{N-1} \frac{(x[n] - \bar{x})^2}{\sigma^2} &\sim \chi_{N-1}^2 \end{aligned} \quad (7.37)$$

and that the two statistics are independent. For large N it is well known from the central limit theorem that

$$\chi_N^2 \stackrel{a}{\sim} \mathcal{N}(N, 2N)$$

so that

$$T \stackrel{a}{\sim} \mathcal{N}(N-1, 2(N-1))$$

or

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 \stackrel{a}{\sim} \mathcal{N}\left(\frac{(N-1)\sigma^2}{N}, \frac{2(N-1)\sigma^4}{N^2}\right). \quad (7.38)$$

The asymptotic PDF of $\hat{\theta}$ is jointly Gaussian since the statistics are individually Gaussian and independent. Furthermore, from (7.37) and (7.38) for large N

$$\begin{aligned} E(\hat{\theta}) &= \begin{bmatrix} A \\ \frac{(N-1)\sigma^2}{N} \end{bmatrix} \rightarrow \begin{bmatrix} A \\ \sigma^2 \end{bmatrix} = \theta \\ \mathbf{C}(\theta) &= \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{2(N-1)\sigma^4}{N^2} \end{bmatrix} \rightarrow \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{2\sigma^4}{N} \end{bmatrix} = \mathbf{I}^{-1}(\theta) \end{aligned}$$

where the inverse of the Fisher information matrix has previously been given in Example 3.6. Hence, the asymptotic distribution is described by (7.36).

In some instances, the asymptotic PDF of the MLE is not given by (7.36). These situations generally occur when the number of parameters to be estimated is too large relative to the number of data samples available. As in Example 7.7, this situation restricts the averaging in the estimator.

Example 7.13 - Signal in Non-Gaussian Noise

Consider the data

$$x[n] = s[n] + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is zero mean IID noise with the Laplacian PDF

$$p(w[n]) = \frac{1}{4} \exp\left[-\frac{1}{2}|w[n]|\right].$$

The signal samples $\{s[0], s[1], \dots, s[N-1]\}$ are to be estimated. The PDF of the data is

$$p(\mathbf{x}; \theta) = \prod_{n=0}^{N-1} \frac{1}{4} \exp\left[-\frac{1}{2}|x[n] - s[n]|\right]. \quad (7.39)$$

The MLE of $\theta = [s[0] s[1] \dots s[N-1]]^T$ is easily seen to be

$$\hat{s}[n] = x[n] \quad n = 0, 1, \dots, N-1$$

or, equivalently,

$$\hat{\theta} = \mathbf{x}.$$

It is clear that the MLE is *not* Gaussian even as $N \rightarrow \infty$. The PDF of $\hat{\theta}$ is given by (7.39) with $\mathbf{x}$ replaced by $\hat{\theta}$. The difficulty is that no averaging is possible since we have chosen to estimate as many parameters as data points. Therefore, the central limit theorem, which accounts for the asymptotic Gaussian PDF of the MLE (see Appendix 7B), does not apply. $\diamond$

As in the scalar case, the invariance property holds, as summarized by the following theorem.

Theorem 7.4 (Invariance Property of MLE (Vector Parameter)) *The MLE of the parameter $\alpha = \mathbf{g}(\theta)$, where $\mathbf{g}$ is an r -dimensional function of the $p \times 1$ parameter θ , and the PDF $p(\mathbf{x}; \theta)$ is parameterized by θ , is given by*

$$\hat{\alpha} = \mathbf{g}(\hat{\theta})$$

for $\hat{\theta}$, the MLE of θ . If $\mathbf{g}$ is not an invertible function, then $\hat{\alpha}$ maximizes the modified likelihood function $\bar{p}_T(\mathbf{x}; \alpha)$, defined as

$$\bar{p}_T(\mathbf{x}; \alpha) = \max_{\{\theta: \alpha = \mathbf{g}(\theta)\}} p(\mathbf{x}; \theta).$$

An example of the use of this theorem can be found in Problem 7.21.

In Chapter 3 we discussed the computation of the CRLB for the general Gaussian case in which the observed data vector had the PDF

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\theta), \mathbf{C}(\theta)).$$

In Appendix 3C (see (3C.5)) it was shown that the partial derivatives of the log-likelihood function were

$$\begin{aligned} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_k} &= -\frac{1}{2} \text{tr} \left(\mathbf{C}^{-1}(\theta) \frac{\partial \mathbf{C}(\theta)}{\partial \theta_k} \right) + \frac{\partial \boldsymbol{\mu}(\theta)^T}{\partial \theta_k} \mathbf{C}^{-1}(\theta) (\mathbf{x} - \boldsymbol{\mu}(\theta)) \\ &\quad - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}(\theta))^T \frac{\partial \mathbf{C}^{-1}(\theta)}{\partial \theta_k} (\mathbf{x} - \boldsymbol{\mu}(\theta)) \end{aligned} \quad (7.40)$$

for $k = 1, 2, \dots, p$. By setting (7.40) equal to zero we can obtain necessary conditions for the MLE and on occasion, if solvable, the exact MLE. A somewhat more convenient form uses (3C.2) in Appendix 3C to replace $\partial \mathbf{C}^{-1}(\theta) / \partial \theta_k$ in (7.40). An important example of the latter is the linear model described in detail in Chapter 4. Recall that the general linear data model is

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{w} \quad (7.41)$$

where $\mathbf{H}$ is a known $N \times p$ matrix and $\mathbf{w}$ is a noise vector of dimension $N \times 1$ with PDF $\mathcal{N}(\mathbf{0}, \mathbf{C})$. Under these conditions the PDF is

$$p(\mathbf{x}; \theta) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{H}\theta)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta) \right]$$

so that the MLE of θ is found by minimizing

$$J(\theta) = (\mathbf{x} - \mathbf{H}\theta)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta). \quad (7.42)$$

Since this is a quadratic function of the elements of θ and $\mathbf{C}^{-1}$ is a positive definite matrix, differentiation will produce the global minimum. Now, using (7.40) and noting that

$$\mu(\theta) = \mathbf{H}\theta$$

and that the covariance matrix does not depend on θ , we have

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta_k} = \frac{\partial (\mathbf{H}\theta)^T}{\partial \theta_k} \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta).$$

Combining the partial derivatives to form the gradient, we have

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = \frac{\partial (\mathbf{H}\theta)^T}{\partial \theta} \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta)$$

so that upon setting the gradient equal to zero we have

$$\mathbf{H}^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\hat{\theta}) = 0.$$

Solving for $\hat{\theta}$ produces the MLE

$$\hat{\theta} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}. \quad (7.43)$$

But this estimator was shown in Chapter 4 to be the MVU estimator as well as an efficient estimator. As such, $\hat{\theta}$ is unbiased and has a covariance of

$$\mathbf{C}_{\hat{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}. \quad (7.44)$$

Finally, the PDF of the MLE is Gaussian (being a linear function of $\mathbf{x}$), so that the asymptotic properties of Theorem 7.3 are attained even for finite data records. For the linear model it can be said that *the MLE is optimal*. These results are summarized by the following theorem.

Theorem 7.5 (Optimality of the MLE for the Linear Model) *If the observed data $\mathbf{x}$ are described by the general linear model*

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{w} \quad (7.45)$$

where $\mathbf{H}$ is a known $N \times p$ matrix with $N > p$ and of rank p , θ is a $p \times 1$ parameter vector to be estimated, and $\mathbf{w}$ is a noise vector with PDF $\mathcal{N}(\mathbf{0}, \mathbf{C})$, then the MLE of θ is

$$\hat{\theta} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}. \quad (7.46)$$

$\hat{\theta}$ is also an efficient estimator in that it attains the CRLB and hence is the MVU estimator. The PDF of $\hat{\theta}$ is

$$\hat{\theta} \sim \mathcal{N}(\theta, (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}). \quad (7.47)$$

Many examples are given in Chapter 4. The preceding result can be generalized to assert that *if an efficient estimator exists, it is given by the MLE* (see Problem 7.12).

In finding the MLE it is quite common to have to resort to numerical techniques of maximization. The Newton-Raphson and scoring methods were described in Section 7.7. They are easily extended to the vector parameter case. The Newton-Raphson iteration becomes

$$\theta_{k+1} = \theta_k - \left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta \partial \theta^T} \right]^{-1} \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \bigg|_{\theta = \theta_k} \quad (7.48)$$

where

$$\left[\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta \partial \theta^T} \right]_{ij} = \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta_i \partial \theta_j} \quad \begin{matrix} i = 1, 2, \dots, p \\ j = 1, 2, \dots, p \end{matrix}$$

is the Hessian of the log-likelihood function and $\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta}$ is the $p \times 1$ gradient vector. In implementing (7.48) inversion of the Hessian is not required. Rewriting (7.48) as

$$\frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta \partial \theta^T} \bigg|_{\theta = \theta_k} \theta_{k+1} = \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta \partial \theta^T} \bigg|_{\theta = \theta_k} \theta_k - \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \bigg|_{\theta = \theta_k} \quad (7.49)$$

we see that the new iterate, θ_{k+1} , can be found from the previous iterate, θ_k , by solving a set of p simultaneous linear equations.

The scoring method is obtained from the Newton-Raphson method by replacing the Hessian by the negative of the Fisher information matrix to yield

$$\theta_{k+1} = \theta_k + \mathbf{I}^{-1}(\theta) \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \bigg|_{\theta = \theta_k} \quad (7.50)$$

As explained previously, the inversion may be avoided by writing (7.50) in the form of (7.49). As in the scalar case, the Newton-Raphson and scoring methods may suffer from convergence problems (see Section 7.7). Care must be exercised in using them. Typically, as the data record becomes large, the log-likelihood function becomes more nearly quadratic near the maximum, and the iterative procedure will produce the MLE.

A third method of numerically determining the MLE is the expectation-maximization (EM) algorithm [Dempster, Laird, and Rubin 1977]. This method, although iterative in nature, is guaranteed under certain mild conditions to converge, and at convergence to produce at least a local maximum. It has the desirable property of increasing the likelihood at each step. The EM algorithm exploits the observation that some data sets may allow easier determination of the MLE than the given one. As an example, consider

$$x[n] = \sum_{i=1}^p \cos 2\pi f_i n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 and the frequencies $\mathbf{f} = [f_1 f_2 \dots f_p]^T$ are to be estimated. The MLE would require a multidimensional minimization of

$$J(\mathbf{f}) = \sum_{n=0}^{N-1} \left(x[n] - \sum_{i=1}^p \cos 2\pi f_i n \right)^2.$$

On the other hand, if the original data could be replaced by the independent data sets

$$y_i[n] = \cos 2\pi f_i n + w_i[n] \quad \begin{matrix} i = 1, 2, \dots, p \\ n = 0, 1, \dots, N-1 \end{matrix} \quad (7.51)$$

where $w_i[n]$ is WGN with variance σ_i^2 , then the problem would be decoupled. It is easily shown that due to the independence assumption the PDF would factor into the PDF for each data set. Consequently, the MLE of f_i could be obtained from a minimization of

$$J(f_i) = \sum_{n=0}^{N-1} (y_i[n] - \cos 2\pi f_i n)^2 \quad i = 1, 2, \dots, p.$$

The original p -dimensional minimization has been reduced to p separate one-dimensional minimizations, in general, an easier problem. The new data set $\{y_1[n], y_2[n], \dots, y_p[n]\}$ is termed the *complete data* and can be related to the original data as

$$x[n] = \sum_{i=1}^p y_i[n] \quad (7.52)$$

if the noise is decomposed as

$$w[n] = \sum_{i=1}^p w_i[n].$$

For this decomposition to hold we will assume that the $w_i[n]$ noise processes are independent of each other and

$$\sigma^2 = \sum_{i=1}^p \sigma_i^2. \quad (7.53)$$

The key question remains as to how to obtain the complete data from the original or *incomplete* data. The reader should also note that the decomposition is not unique. We could have just as easily hypothesized the complete data to be

$$\begin{aligned} y_1[n] &= \sum_{i=1}^p \cos 2\pi f_i n \\ y_2[n] &= w[n] \end{aligned}$$

and thus

$$x[n] = y_1[n] + y_2[n]$$

as a signal and noise decomposition.

In general, we suppose that there is a complete to incomplete data transformation given as

$$\mathbf{x} = \mathbf{g}(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M) = \mathbf{g}(\mathbf{y}) \quad (7.54)$$

where in the previous example $M = p$ and the elements of $\mathbf{y}_i$ are given by (7.51). The function $\mathbf{g}$ is a many-to-one transformation. We wish to find the MLE of θ by

maximizing $\ln p_x(\mathbf{x}; \boldsymbol{\theta})$. Finding this too difficult, we instead maximize $\ln p_y(\mathbf{y}; \boldsymbol{\theta})$. Since $\mathbf{y}$ is unavailable, we replace the log-likelihood function by its conditional expectation or

$$E_{\mathbf{y}|\mathbf{x}}[\ln p_y(\mathbf{y}; \boldsymbol{\theta})] = \int \ln p_y(\mathbf{y}; \boldsymbol{\theta}) p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) d\mathbf{y}. \quad (7.55)$$

Finally, since we need to know $\boldsymbol{\theta}$ to determine $p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$ and hence the expected log-likelihood function, we use the current guess. Letting $\boldsymbol{\theta}_k$ denote the k th guess of the MLE of $\boldsymbol{\theta}$, we then have the following iterative algorithm:

Expectation (E): Determine the average log-likelihood of the complete data

$$U(\boldsymbol{\theta}, \boldsymbol{\theta}_k) = \int \ln p_y(\mathbf{y}; \boldsymbol{\theta}) p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_k) d\mathbf{y}. \quad (7.56)$$

Maximization (M): Maximize the average log-likelihood function of the complete data

$$\boldsymbol{\theta}_{k+1} = \arg \max_{\boldsymbol{\theta}} U(\boldsymbol{\theta}, \boldsymbol{\theta}_k). \quad (7.57)$$

At convergence we hopefully will have the MLE. This approach is termed the *EM algorithm*. Applying it to our frequency estimation problem, we have the following iteration as derived in Appendix 7C.

E Step: For $i = 1, 2, \dots, p$

$$\hat{y}_i[n] = \cos 2\pi f_{i_k} n + \beta_i(x[n] - \sum_{i=1}^p \cos 2\pi f_{i_k} n). \quad (7.58)$$

M Step: For $i = 1, 2, \dots, p$

$$f_{i_{k+1}} = \arg \max_{f_i} \sum_{n=0}^{N-1} \hat{y}_i[n] \cos 2\pi f_i n \quad (7.59)$$

where the β_i 's can be arbitrarily chosen as long as $\sum_{i=1}^p \beta_i = 1$.

Note that $\beta_i(x[n] - \sum_{i=1}^p \cos 2\pi f_{i_k} n)$ in (7.58) is an estimate of $w_i[n]$. Hence, the algorithm iteratively decouples the original data set into p separate data sets, with each one consisting of a *single* sinusoid in WGN. The maximization given above corresponds to the MLE of a single sinusoid with the data set given by the *estimated* complete data (see Problem 7.19). Good results have been obtained with this approach. Its disadvantages are the difficulty of determining the conditional expectation in closed form and the arbitrariness in the choice of the complete data. Nonetheless, for the Gaussian problem this method can easily be applied. The reader should consult [Feder and Weinstein 1988] for further details.

7.9 Asymptotic MLE

In many cases it is difficult to evaluate the MLE of a parameter whose PDF is Gaussian due to the need to invert a large dimension covariance matrix. For example, if $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}(\boldsymbol{\theta}))$, the MLE of $\boldsymbol{\theta}$ is obtained by maximizing

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C}(\boldsymbol{\theta}))} \exp \left[-\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1}(\boldsymbol{\theta}) \mathbf{x} \right].$$

If the covariance matrix cannot be inverted in closed form, then a search technique will require inversion of the $N \times N$ matrix for each value of $\boldsymbol{\theta}$ to be searched. An alternative *approximate* method can be applied when $\mathbf{x}$ is data from a zero mean WSS random process, so that the covariance matrix is Toeplitz (see Appendix 1). In such a case, it is shown in Appendix 3D that the asymptotic (for large data records) log-likelihood function is given by (see (3D.6))

$$\ln p(\mathbf{x}; \boldsymbol{\theta}) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\ln P_{xx}(f) + \frac{I(f)}{P_{xx}(f)} \right] df \quad (7.60)$$

where

$$I(f) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x[n] \exp(-j2\pi f n) \right|^2$$

is the periodogram of the data and $P_{xx}(f)$ is the PSD. The dependence of the log-likelihood function on $\boldsymbol{\theta}$ is through the PSD. Differentiation of (7.60) produces the necessary conditions for the MLE

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i} = -\frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\frac{1}{P_{xx}(f)} - \frac{I(f)}{P_{xx}^2(f)} \right] \frac{\partial P_{xx}(f)}{\partial \theta_i} df$$

or

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\frac{1}{P_{xx}(f)} - \frac{I(f)}{P_{xx}^2(f)} \right] \frac{\partial P_{xx}(f)}{\partial \theta_i} df = 0. \quad (7.61)$$

The second derivative is given by (3D.7) in Appendix 3D, so that the Newton-Raphson or scoring method may be implemented using the asymptotic likelihood function. This leads to simpler iterative procedures and is commonly used in practice. An example follows.

Example 7.14 - Gaussian Moving Average Process

A common WSS random process has the ACF

$$r_{xx}[k] = \begin{cases} 1 + b^2[1] + b^2[2] & k = 0 \\ b[1] + b[1]b[2] & k = 1 \\ b[2] & k = 2 \\ 0 & k \geq 3. \end{cases}$$

The PSD is easily shown to be

$$P_{xx}(f) = |1 + b[1]\exp(-j2\pi f) + b[2]\exp(-j4\pi f)|^2.$$

The process $x[n]$ is obtained by passing WGN of variance 1 through a filter with system function $B(z) = 1 + b[1]z^{-1} + b[2]z^{-2}$. It is usually assumed that the filter $B(z)$ is minimum-phase or that the zeros z_1, z_2 are within the unit circle. This is termed a moving average (MA) process of order 2. In finding the MLE of the MA filter parameters $b[1], b[2]$, we will need to invert the covariance matrix. Instead, we can use (7.60) to obtain an approximate MLE. Noting that the filter is minimum-phase

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln P_{xx}(f) df = 0$$

for this process (see Problem 7.22), the approximate MLE is found by minimizing

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{I(f)}{|1 + b[1]\exp(-j2\pi f) + b[2]\exp(-j4\pi f)|^2} df.$$

However,

$$B(z) = (1 - z_1 z^{-1})(1 - z_2 z^{-1})$$

with $|z_1| < 1, |z_2| < 1$. The MLE is found by minimizing

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{I(f)}{|1 - z_1 \exp(-j2\pi f)|^2 |1 - z_2 \exp(-j2\pi f)|^2} df$$

over the allowable zero values and then converting to the MA parameters by

$$\begin{aligned} b[1] &= -(z_1 + z_2) \\ b[2] &= z_1 z_2. \end{aligned}$$

For values of z_1 that are complex we would employ the constraint $z_2 = z_1^*$. For real values of z_1 we must ensure that z_2 is also real. These constraints are necessitated by the coefficients of $B(z)$ being real. Either a grid search or one of the iterative techniques could then be used. $\diamond$

Another example of the asymptotic MLE is given in the next section.

7.10 Signal Processing Examples

In Section 3.11 we discussed four signal processing estimation examples. The CRLB was determined for each one. By the asymptotic properties of the MLE the asymptotic PDF of the MLE is given by $\mathcal{N}(\boldsymbol{\theta}, \mathbf{I}^{-1}(\boldsymbol{\theta}))$, where $\mathbf{I}^{-1}(\boldsymbol{\theta})$ is given in Section 3.11. We now determine the explicit form for the MLEs. The reader should consult Section 3.11 for the problem statements and notation.

Example 7.15 - Range Estimation (Example 3.13)

From (3.36) the PDF is seen to be

$$p(\mathbf{x}; n_0) = \prod_{n=0}^{n_0-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} x^2[n] \right] \\ \cdot \prod_{n=n_0}^{n_0+M-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} (x[n] - s[n - n_0])^2 \right] \\ \cdot \prod_{n=n_0+M}^{N-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} x^2[n] \right].$$

In this case the continuous parameter τ_0 has been discretized as $n_0 = \tau_0/\Delta$. The likelihood function simplifies to

$$p(\mathbf{x}; n_0) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] \right] \\ \cdot \prod_{n=n_0}^{n_0+M-1} \exp \left[-\frac{1}{2\sigma^2} (-2x[n]s[n - n_0] + s^2[n - n_0]) \right].$$

The MLE of n_0 is found by maximizing

$$\exp \left[-\frac{1}{2\sigma^2} \sum_{n=n_0}^{n_0+M-1} (-2x[n]s[n - n_0] + s^2[n - n_0]) \right]$$

or, equivalently, by minimizing

$$\sum_{n=n_0}^{n_0+M-1} (-2x[n]s[n - n_0] + s^2[n - n_0]).$$

But $\sum_{n=n_0}^{n_0+M-1} s^2[n - n_0] = \sum_{n=0}^{N-1} s^2[n]$ and is not a function of n_0 . Hence, the MLE of n_0 is found by maximizing

$$\sum_{n=n_0}^{n_0+M-1} x[n]s[n - n_0]. \quad (7.62)$$

By the invariance principle, since $R = c\tau_0/2 = cn_0\Delta/2$, the MLE of range is $\hat{R} = (c\Delta/2)\hat{n}_0$. Note that the MLE of the delay n_0 is found by *correlating* the data with all possible received signals and then choosing the maximum. $\diamond$

Example 7.16 - Sinusoidal Parameter Estimation (Example 3.14)

The PDF is given as

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi))^2 \right]$$

where $A > 0$ and $0 < f_0 < 1/2$. The MLE of amplitude A , frequency f_0 , and phase ϕ is found by minimizing

$$J(A, f_0, \phi) = \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi))^2.$$

We first expand the cosine to yield

$$J(A, f_0, \phi) = \sum_{n=0}^{N-1} (x[n] - A \cos \phi \cos 2\pi f_0 n + A \sin \phi \sin 2\pi f_0 n)^2.$$

Although J is nonquadratic in A and ϕ , we may transform it to a quadratic function by letting

$$\begin{aligned} \alpha_1 &= A \cos \phi \\ \alpha_2 &= -A \sin \phi \end{aligned}$$

which is a one-to-one transformation. The inverse transformation is given by

$$\begin{aligned} A &= \sqrt{\alpha_1^2 + \alpha_2^2} \\ \phi &= \arctan \left(\frac{-\alpha_2}{\alpha_1} \right). \end{aligned} \quad (7.63)$$

Also, let

$$\begin{aligned} \mathbf{c} &= [1 \cos 2\pi f_0 \dots \cos 2\pi f_0 (N-1)]^T \\ \mathbf{s} &= [0 \sin 2\pi f_0 \dots \sin 2\pi f_0 (N-1)]^T. \end{aligned}$$

Then, we have

$$\begin{aligned} J'(\alpha_1, \alpha_2, f_0) &= (\mathbf{x} - \alpha_1 \mathbf{c} - \alpha_2 \mathbf{s})^T (\mathbf{x} - \alpha_1 \mathbf{c} - \alpha_2 \mathbf{s}) \\ &= (\mathbf{x} - \mathbf{H}\boldsymbol{\alpha})^T (\mathbf{x} - \mathbf{H}\boldsymbol{\alpha}) \end{aligned}$$

where $\boldsymbol{\alpha} = [\alpha_1 \alpha_2]^T$ and $\mathbf{H} = [\mathbf{c} \mathbf{s}]$. Hence, the function to be minimized over $\boldsymbol{\alpha}$ is exactly that encountered in the linear model in (7.42) with $\mathbf{C} = \mathbf{I}$. The minimizing solution is, from (7.43),

$$\hat{\alpha} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \quad (7.64)$$

so that

$$\begin{aligned} J'(\hat{\alpha}_1, \hat{\alpha}_2, f_0) &= (\mathbf{x} - \mathbf{H}\hat{\alpha})^T (\mathbf{x} - \mathbf{H}\hat{\alpha}) \\ &= (\mathbf{x} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x})^T (\mathbf{x} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}) \\ &= \mathbf{x}^T (\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T) \mathbf{x} \end{aligned}$$

since $\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T$ is an idempotent matrix. (A matrix is idempotent if $\mathbf{A}^2 = \mathbf{A}$.) Hence, to find $\hat{f}_0$ we need to minimize J' over f_0 or, equivalently, maximize

$$\mathbf{x}^T \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

Using the definition of $\mathbf{H}$, we have the MLE for frequency as the value that maximizes

$$\begin{bmatrix} \mathbf{c}^T \mathbf{x} \\ s^T \mathbf{x} \end{bmatrix}^T \begin{bmatrix} \mathbf{c}^T \mathbf{c} & \mathbf{c}^T s \\ s^T \mathbf{c} & s^T s \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{c}^T \mathbf{x} \\ s^T \mathbf{x} \end{bmatrix}. \quad (7.65)$$

Once $\hat{f}_0$ is found from the above expression, $\hat{\alpha}$ can be found from (7.64), and then $\hat{A}, \hat{\phi}$ from (7.63). An *approximate* MLE can be obtained if f_0 is not near 0 or 1/2 (see also Problem 7.24) since in this case

$$\begin{aligned} \frac{1}{N} \mathbf{c}^T s &= \frac{1}{N} \sum_{n=0}^{N-1} \cos 2\pi f_0 n \sin 2\pi f_0 n \\ &= \frac{1}{2N} \sum_{n=0}^{N-1} \sin 4\pi f_0 n \\ &\approx 0 \end{aligned}$$

and, similarly, $\mathbf{c}^T \mathbf{c}/N \approx 1/2$ and $s^T s/N \approx 1/2$ (see Problem 3.7). Then, (7.65) becomes approximately

$$\begin{aligned} &\begin{bmatrix} \mathbf{c}^T \mathbf{x} \\ s^T \mathbf{x} \end{bmatrix}^T \begin{bmatrix} \frac{N}{2} & 0 \\ 0 & \frac{N}{2} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{c}^T \mathbf{x} \\ s^T \mathbf{x} \end{bmatrix} \\ &= \frac{2}{N} [(\mathbf{c}^T \mathbf{x})^2 + (s^T \mathbf{x})^2] \\ &= \frac{2}{N} \left[\left(\sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \right)^2 + \left(\sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n \right)^2 \right] \\ &= \frac{2}{N} \left| \sum_{n=0}^{N-1} x[n] \exp(-j2\pi f_0 n) \right|^2 \end{aligned}$$

or the MLE of frequency is obtained by maximizing the periodogram

$$I(f) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x[n] \exp(-j2\pi fn) \right|^2 \quad (7.66)$$

over f . Then, from (7.64)

$$\begin{aligned} \hat{\alpha} &\approx \frac{2}{N} \begin{bmatrix} \hat{\mathbf{c}}^T \mathbf{x} \\ \hat{\mathbf{s}}^T \mathbf{x} \end{bmatrix} \\ &= \begin{bmatrix} \frac{2}{N} \sum_{n=0}^{N-1} x[n] \cos 2\pi \hat{f}_0 n \\ \frac{2}{N} \sum_{n=0}^{N-1} x[n] \sin 2\pi \hat{f}_0 n \end{bmatrix} \end{aligned}$$

and thus finally we have

$$\begin{aligned} \hat{A} &= \sqrt{\hat{\alpha}_1^2 + \hat{\alpha}_2^2} = \frac{2}{N} \left| \sum_{n=0}^{N-1} x[n] \exp(-j2\pi \hat{f}_0 n) \right| \\ \hat{\phi} &= \arctan \frac{-\sum_{n=0}^{N-1} x[n] \sin 2\pi \hat{f}_0 n}{\sum_{n=0}^{N-1} x[n] \cos 2\pi \hat{f}_0 n}. \end{aligned}$$

◇

Example 7.17 - Bearing Estimation (Example 3.15)

The bearing β of a sinusoidal signal impinging on a line array is related to the spatial frequency f_s as

$$f_s = \frac{F_0 d}{c} \cos \beta.$$

Assuming the amplitude, phase, and spatial frequency are to be estimated, we have the sinusoidal parameter estimation problem just described. It is assumed, however, that the temporal frequency F_0 is known, as well as d and c . Once f_s is estimated using an MLE, the MLE of bearing follows from the invariance property as

$$\hat{\beta} = \arccos \left(\frac{\hat{f}_s c}{F_0 d} \right).$$

To find $\hat{f}_s$ it is common practice to use the approximate MLE obtained from the peak of the periodogram (7.66) or

$$I(f_s) = \frac{1}{M} \left| \sum_{n=0}^{M-1} x[n] \exp(-j2\pi f_s n) \right|^2$$

for M sensors. Here, $x[n]$ is the sample obtained from sensor n . Alternatively, the approximate MLE of bearing can be found directly by maximizing

$$I'(\beta) = \frac{1}{M} \left| \sum_{n=0}^{M-1} x[n] \exp \left(-j2\pi F_0 \frac{d}{c} n \cos \beta \right) \right|^2$$

over β . ◇

Example 7.18 - Autoregressive Parameter Estimation (Example 3.16)

To find the MLE we will use the asymptotic form of the log-likelihood given by (7.60). Then, since the PSD is

$$P_{xx}(f) = \frac{\sigma_u^2}{|A(f)|^2}$$

we have

$$\ln p(\mathbf{x}; \mathbf{a}, \sigma_u^2) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\ln \frac{\sigma_u^2}{|A(f)|^2} + \frac{I(f)}{\frac{\sigma_u^2}{|A(f)|^2}} \right] df.$$

As shown in Problem 7.22, since $A(z)$ is minimum-phase (required for the stability of $1/A(z)$), then

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |A(f)|^2 df = 0$$

and therefore

$$\ln p(\mathbf{x}; \mathbf{a}, \sigma_u^2) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \ln \sigma_u^2 - \frac{N}{2\sigma_u^2} \int_{-\frac{1}{2}}^{\frac{1}{2}} |A(f)|^2 I(f) df. \quad (7.67)$$

Differentiating with respect to σ_u^2 and setting the result equal to zero produces

$$-\frac{N}{2\sigma_u^2} + \frac{N}{2\sigma_u^4} \int_{-\frac{1}{2}}^{\frac{1}{2}} |A(f)|^2 I(f) df = 0$$

and thus

$$\hat{\sigma}_u^2 = \int_{-\frac{1}{2}}^{\frac{1}{2}} |A(f)|^2 I(f) df.$$

Substituting this into (7.67) results in

$$\ln p(\mathbf{x}; \mathbf{a}, \hat{\sigma}_u^2) = -\frac{N}{2} \ln 2\pi - \frac{N}{2} \ln \hat{\sigma}_u^2 - \frac{N}{2}.$$

To find $\hat{\mathbf{a}}$ we must minimize $\hat{\sigma}_u^2$ or

$$J(\mathbf{a}) = \int_{-\frac{1}{2}}^{\frac{1}{2}} |A(f)|^2 I(f) df.$$

Note that this function is quadratic in $\mathbf{a}$, resulting in the global minimum upon differentiation. For $k = 1, 2, \dots, p$ we have

$$\begin{aligned} \frac{\partial J(\mathbf{a})}{\partial a[k]} &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[A(f) \frac{\partial A^*(f)}{\partial a[k]} + \frac{\partial A(f)}{\partial a[k]} A^*(f) \right] I(f) df \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} [A(f) \exp(j2\pi f k) + A^*(f) \exp(-j2\pi f k)] I(f) df. \end{aligned}$$

Since $A(-f) = A^*(f)$ and $I(-f) = I(f)$, we can rewrite this as

$$\frac{\partial J(\mathbf{a})}{\partial a[k]} = 2 \int_{-\frac{1}{2}}^{\frac{1}{2}} A(f) I(f) \exp(j2\pi f k) df.$$

Setting it equal to zero produces

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \left(1 + \sum_{l=1}^p a[l] \exp(-j2\pi f l) \right) I(f) \exp(j2\pi f k) df = 0 \quad k = 1, 2, \dots, p \quad (7.68)$$

or

$$\sum_{l=1}^p a[l] \int_{-\frac{1}{2}}^{\frac{1}{2}} I(f) \exp[j2\pi f(k-l)] df = - \int_{-\frac{1}{2}}^{\frac{1}{2}} I(f) \exp[j2\pi f k] df.$$

But $\int_{-\frac{1}{2}}^{\frac{1}{2}} I(f) \exp[j2\pi f k] df$ is just the inverse Fourier transform of the periodogram evaluated at k . This can be shown to be the estimated ACF (see Problem 7.25):

$$\hat{r}_{xx}[k] = \begin{cases} \frac{1}{N} \sum_{n=0}^{N-1-|k|} x[n]x[n+|k|] & |k| \leq N-1 \\ 0 & |k| \geq N. \end{cases}$$

The set of equations to be solved for the approximate MLE of the AR filter parameters $\mathbf{a}$ becomes

$$\sum_{l=1}^p \hat{a}[l] \hat{r}_{xx}[k-l] = -\hat{r}_{xx}[k] \quad k = 1, 2, \dots, p$$

or in matrix form

$$\begin{bmatrix} \hat{r}_{xx}[0] & \hat{r}_{xx}[1] & \cdots & \hat{r}_{xx}[p-1] \\ \hat{r}_{xx}[1] & \hat{r}_{xx}[0] & \cdots & \hat{r}_{xx}[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ \hat{r}_{xx}[p-1] & \hat{r}_{xx}[p-2] & \cdots & \hat{r}_{xx}[0] \end{bmatrix} \begin{bmatrix} \hat{a}[1] \\ \hat{a}[2] \\ \vdots \\ \hat{a}[p] \end{bmatrix} = - \begin{bmatrix} \hat{r}_{xx}[1] \\ \hat{r}_{xx}[2] \\ \vdots \\ \hat{r}_{xx}[p] \end{bmatrix}. \quad (7.69)$$

These are the so-called *estimated Yule-Walker equations* and this is the *autocorrelation method of linear prediction*. Note the special form of the matrix and the right-hand vector, thereby allowing a recursive solution known as the Levinson recursion [Kay 1988]. To complete the discussion we determine an explicit form for the MLE of σ_u^2 . From (7.68) we note that

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \hat{A}(f) I(f) \exp(j2\pi f k) df = 0 \quad k = 1, 2, \dots, p.$$

Since

$$\begin{aligned} \hat{\sigma}_u^2 &= \int_{-\frac{1}{2}}^{\frac{1}{2}} |\hat{A}(f)|^2 I(f) df \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \hat{A}(f) I(f) \hat{A}^*(f) df \\ &= \sum_{k=0}^p \hat{a}[k] \int_{-\frac{1}{2}}^{\frac{1}{2}} \hat{A}(f) I(f) \exp(j2\pi f k) df \end{aligned}$$

we note that all terms in the sum are zero except for $k = 0$ due to (7.69). It follows that (we let $\hat{a}[0] = 1$)

$$\begin{aligned} \hat{\sigma}_u^2 &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \hat{A}(f) I(f) df \\ &= \sum_{k=0}^p \hat{a}[k] \int_{-\frac{1}{2}}^{\frac{1}{2}} I(f) \exp(-j2\pi f k) df \\ &= \sum_{k=0}^p \hat{a}[k] \hat{r}_{xx}[-k] \\ &= \sum_{k=0}^p \hat{a}[k] \hat{r}_{xx}[k]. \end{aligned}$$

Finally, the MLE is given by

$$\hat{\sigma}_u^2 = \hat{r}_{xx}[0] + \sum_{k=1}^p \hat{a}[k] \hat{r}_{xx}[k]. \quad (7.70)$$

References

- Bard, Y., *Nonlinear Parameter Estimation*, Academic Press, New York, 1974.
- Bendat, J.S., A.G. Piersol, *Random Data: Analysis and Measurement Procedures*, J. Wiley, New York, 1971.
- Bickel, P.J., K.A. Doksum, *Mathematical Statistics*, Holden-Day, San Francisco, 1977.
- Dempster, A.P., N.M. Laird, D.B. Rubin, "Maximum Likelihood From Incomplete Data via the EM Algorithm," *Ann. Roy. Statist. Soc.*, Vol. 39, pp. 1-38, Dec. 1977.
- Dudewicz, E.J., *Introduction to Probability and Statistics*, Holt, Rinehart, and Winston, New York, 1976.
- Feder, M., E. Weinstein, "Parameter Estimation of Superimposed Signals Using the EM Algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 36, pp. 477-489, April 1988.
- Hoel, P.G., S.C. Port, C.J. Stone, *Introduction to Statistical Theory*, Houghton Mifflin, Boston, 1971.
- Kay, S.M., *Modern Spectral Estimation: Theory and Application*, Prentice-Hall, Englewood Cliffs, N.J., 1988.
- Rao, C.R. *Linear Statistical Inference and Its Applications*, J. Wiley, New York, 1973.
- Schwartz, M., L. Shaw, *Signal Processing: Discrete Spectral Analysis, Detection, and Estimation*, McGraw-Hill, New York, 1975.

Problems

- 7.1 Show that the CRLB for $\hat{A}$ in Example 7.1 is given by (7.2).
- 7.2 Consider the sample mean estimator for the problem in Example 7.1. Find the variance and compare it to the CRLB. Does the sample mean estimator attain the CRLB for finite N ? How about if $N \rightarrow \infty$? Is the MLE or the sample mean a better estimator?
- 7.3 We observe N IID samples from the PDFs:

a. Gaussian

$$p(x; \mu) = \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2}(x - \mu)^2 \right].$$

b. Exponential

$$p(x; \lambda) = \begin{cases} \lambda \exp(-\lambda x) & x > 0 \\ 0 & x < 0. \end{cases}$$

In each case find the MLE of the unknown parameter and be sure to verify that it indeed maximizes the likelihood function. Do the estimators make sense?

- 7.4 Asymptotic results can sometimes be misleading and so should be carefully applied. As an example, for two unbiased estimators of θ the variances are given by

$$\begin{aligned} \text{var}(\hat{\theta}_1) &= \frac{2}{N} \\ \text{var}(\hat{\theta}_2) &= \frac{1}{N} + \frac{100}{N^2}. \end{aligned}$$

Plot the variances versus N to determine the better estimator.

- 7.5** A formal definition of the consistency of an estimator is given as follows. An estimator $\hat{\theta}$ is consistent if, given any $\epsilon > 0$,

$$\lim_{N \rightarrow \infty} \Pr\{|\hat{\theta} - \theta| > \epsilon\} = 0.$$

Prove that the sample mean is a consistent estimator for the problem of estimating a DC level A in white Gaussian noise of known variance. Hint: Use Chebyshev's inequality.

- 7.6** Another consistency result asserts that if $\alpha = g(\theta)$ for g , a continuous function, and $\hat{\theta}$ is consistent for θ , then $\hat{\alpha} = g(\hat{\theta})$ is consistent for α . Using the statistical linearization argument and the formal definition of consistency, show why this is true. Hint: Linearize g about the true value of θ .

- 7.7** Consider N IID observations from the exponential family of PDFs

$$p(x; \theta) = \exp[A(\theta)B(x) + C(x) + D(\theta)]$$

where A, B, C , and D are functions of their respective arguments. Find an equation to be solved for the MLE. Now apply your results to the PDFs in Problem 7.3.

- 7.8** If we observe N IID samples from a Bernoulli experiment (coin toss) with the probabilities

$$\begin{aligned}\Pr\{x[n] = 1\} &= p \\ \Pr\{x[n] = 0\} &= 1 - p\end{aligned}$$

find the MLE of p .

- 7.9** For N IID observations from a $\mathcal{U}[0, \theta]$ PDF find the MLE of θ .

- 7.10** If the data set

$$\mathbf{x}[n] = A\mathbf{s}[n] + \mathbf{w}[n] \quad n = 0, 1, \dots, N-1$$

is observed, where $\mathbf{s}[n]$ is known and $\mathbf{w}[n]$ is WGN with known variance σ^2 , find the MLE of A . Determine the PDF of the MLE and whether or not the asymptotic PDF holds.

- 7.11** Find an equation to be solved for the MLE of the correlation coefficient ρ in Problem 3.15. Can you find an approximate solution if $N \rightarrow \infty$? Hint: Let $\sum_{n=0}^{N-1} x_1^2[n]/N \rightarrow 1$ and $\sum_{n=0}^{N-1} x_2^2[n]/N \rightarrow 1$ in the equation where $\mathbf{x}[n] = [x_1[n] \ x_2[n]]^T$.

- 12 In this problem we prove that if an efficient estimator exists, the maximum likelihood method will produce it. Assuming a scalar parameter, if an efficient estimator exists, then we have

$$\frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} = I(\theta)(\hat{\theta} - \theta).$$

Use this to prove the theorem.

- 13 We observe N IID samples of a DC level in WGN or

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where A is to be estimated and $w[n]$ is WGN with known variance σ^2 . Using a Monte Carlo computer simulation, verify that the PDF of the MLE or sample mean is $\mathcal{N}(A, \sigma^2/N)$. Plot the theoretical and computer-generated PDFs for comparison. Use $A = 1$, $\sigma^2 = 0.1$, $N = 50$, and $M = 1000$ realizations. What happens if M is increased to 5000? Hint: Use the computer subroutines given in Appendix 7A.

- 14 Consider the IID data samples $x[n]$ for $n = 0, 1, \dots, N-1$, where $x[n] \sim \mathcal{N}(0, \sigma^2)$. Then, according to Slutsky's theorem [Bickel and Doksum 1977], if $\bar{x}$ denotes the sample mean and

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2$$

denotes the sample variance, then

$$\frac{\bar{x}}{\hat{\sigma}/\sqrt{N}} \stackrel{a}{\sim} \mathcal{N}(0, 1).$$

Although this may be proven analytically, it requires familiarity with convergence of random variables. Instead, implement a Monte Carlo computer simulation to find the PDF of $\bar{x}/(\hat{\sigma}/\sqrt{N})$ for $\sigma^2 = 1$ and $N = 10$, $N = 100$. Compare your results to the theoretical asymptotic PDF. Hint: Use the computer subroutines given in Appendix 7A.

- 15 In this problem we show that the MLE attains its asymptotic PDF when the estimation error is small. Consider Example 7.6 and let

$$\begin{aligned} \epsilon_s &= -\frac{2}{NA} \sum_{n=0}^{N-1} w[n] \sin 2\pi f_0 n \\ \epsilon_c &= \frac{2}{NA} \sum_{n=0}^{N-1} w[n] \cos 2\pi f_0 n \end{aligned}$$

in (7.20). Assume that f_0 is not near 0 or $1/2$, so that the same approximations can be made as in Example 3.4. Prove first that ϵ_s and ϵ_c are approximately

uncorrelated and hence independent Gaussian random variables. Next, determine the PDF of ϵ_s, ϵ_c . Then, assuming ϵ_s and ϵ_c to be small, use a truncated Taylor series expansion of $\hat{\phi}$ about the true value ϕ . This will yield

$$\begin{aligned}\hat{\phi} &= g(\epsilon_s, \epsilon_c) \\ &\approx g(0, 0) + \left. \frac{\partial g(\epsilon_s, \epsilon_c)}{\partial \epsilon_s} \right|_{\epsilon_s=0, \epsilon_c=0} \epsilon_s + \left. \frac{\partial g(\epsilon_s, \epsilon_c)}{\partial \epsilon_c} \right|_{\epsilon_s=0, \epsilon_c=0} \epsilon_c\end{aligned}$$

where the function g is given in (7.20). Use this expansion to find the PDF of $\hat{\phi}$. Compare the variance with the CRLB in Example 3.4. Note that for ϵ_s, ϵ_c to be small, either $N \rightarrow \infty$ and/or $A \rightarrow \infty$.

- 16** For Example 7.7 determine the $\text{var}(\hat{A})$ as well as the CRLB for the non-Gaussian PDF (termed a *Laplacian* PDF)

$$p(w[0]) = \frac{1}{2} \exp(-|w[0]|).$$

Does the MLE attain the CRLB as $N \rightarrow \infty$?

- 17** For N IID observations from a $\mathcal{N}(0, 1/\theta)$ PDF, where $\theta > 0$, find the MLE of θ and its asymptotic PDF.
- 18** Plot the function

$$g(x) = \exp\left[-\frac{1}{2}x^2\right] + 0.1 \exp\left[-\frac{1}{2}(x-10)^2\right]$$

over the domain $-3 \leq x \leq 13$. Find the maximum of the function from the graph. Then, use a Newton-Raphson iteration to find the maximum. In doing so use the initial guesses of $x_0 = 0.5$ and $x_0 = 9.5$. What can you say about the importance of the initial guess?

- 19** If

$$x[n] = \cos 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with known variance σ^2 , show that the MLE of frequency is obtained approximately by maximizing (f_0 not near 0 or $1/2$)

$$\sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n$$

over the interval $0 < f_0 < 1/2$. Next, use a Monte Carlo computer simulation (see Appendix 7A) for $N = 10, f_0 = 0.25, \sigma^2 = 0.01$ and plot the function to be maximized. Apply the Newton-Raphson method for determining the maximum and compare the results to those from a grid search.

20 Consider the data set

$$x[n] = s[n] + w[n] \quad n = 0, 1, \dots, N-1$$

where $s[n]$ is unknown for $n = 0, 1, \dots, N-1$ and $w[n]$ is WGN with known variance σ^2 . Determine the MLE of $s[n]$ and also its PDF. Do the asymptotic MLE properties hold? That is, is the MLE unbiased, efficient, Gaussian, and consistent?

21 For N IID observation from the PDF $\mathcal{N}(A, \sigma^2)$, where A and σ^2 are both unknown, find the MLE of the SNR $\alpha = A^2/\sigma^2$.

22 Prove that if

$$\mathcal{A}(z) = 1 + \sum_{k=1}^p a[k]z^{-k}$$

is a minimum-phase polynomial (all roots inside the unit circle of the z plane), then

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |A(f)|^2 df = 0.$$

To do so first show that

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \ln |A(f)|^2 df = 2 \operatorname{Re} \left\{ \frac{1}{2\pi j} \oint \ln \mathcal{A}(z) \frac{dz}{z} \right\}$$

where the contour is the unit circle in the z plane. Next, note that

$$\frac{1}{2\pi j} \oint \ln \mathcal{A}(z) \frac{dz}{z}$$

is the inverse z transform of $\ln \mathcal{A}(z)$ evaluated at $n = 0$. Finally, use the fact that a minimum-phase $\mathcal{A}(z)$ results in an inverse z transform of $\ln \mathcal{A}(z)$ that is causal.

23 Find the asymptotic MLE for the total power P_0 of the PSD

$$P_{xx}(f) = P_0 Q(f)$$

where

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} Q(f) df = 1.$$

If $Q(f) = 1$ for all f so that the process is WGN, simplify your results. Hint: Use the results from Problem 7.25 for the second part.

24 The peak of the periodogram was shown in Example 7.16 to be the MLE for frequency under the condition that f_0 is not near 0 or $1/2$. Plot the periodogram for $N = 10$ for the frequencies $f_0 = 0.25$ and $f_0 = 0.05$. Use the noiseless data

$$x[n] = \cos 2\pi f_0 n \quad n = 0, 1, \dots, N-1.$$

What happens if this approximation is not valid? Repeat the problem using the exact function

$$\mathbf{x}^T \mathbf{H} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$$

where $\mathbf{H}$ is defined in Example 7.16.

7.25 Prove that the inverse Fourier transform of the periodogram is

$$\hat{r}_{xx}[k] = \begin{cases} \frac{1}{N} \sum_{n=0}^{N-1-|k|} x[n]x[n+|k|] & |k| \leq N-1 \\ 0 & |k| \geq N. \end{cases}$$

Hint: Note that the periodogram can be written as

$$I(f) = \frac{1}{N} X'(f) X'(f)^*$$

where $X'(f)$ is the Fourier transform of the sequence

$$x'[n] = \begin{cases} x[n] & 0 \leq n \leq N-1 \\ 0 & \text{otherwise.} \end{cases}$$

7.26 An AR(1) PSD given as

$$P_{xx}(f) = \frac{\sigma_u^2}{|1 + a[1] \exp(-j2\pi f)|^2}.$$

If $a[1]$ and σ^2 are unknown, find the MLE of $P_{xx}(f)$ for $f = f_0$ by using (7.69) and (7.70). What is the asymptotic variance of the MLE? Hint: Use (3.30) and (3.44).

Appendix 7A

Monte Carlo Methods

We now describe the computer methods employed to generate realizations of the random variable (see Example 7.1)

$$\hat{A} = -\frac{1}{2} + \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{4}}$$

where $x[n] = A + w[n]$ for $w[n]$ WGN with variance $A > 0$. Additionally, how the statistical properties of $\hat{A}$ such as the mean, variance, and PDF are determined is discussed. The description applies more generally in that the methods are valid for any estimator whose performance is to be determined.

The steps are summarized as follows.

Data Generation:

1. Generate N independent $\mathcal{U}[0, 1]$ random variates.
2. Convert these variates to Gaussian random variates using the Box-Mueller transformation to produce $w[n]$.
3. Add A to $w[n]$ to yield $x[n]$ and then compute $\hat{A}$.
4. Repeat the procedure M times to yield the M realizations of $\hat{A}$.

Statistical Properties:

1. Determine the mean using (7.9).
2. Determine the variance using (7.10).
3. Determine the PDF using a histogram.

This approach is implemented in the accompanying Fortran computer program, MONTECARLO, which we now describe.




Figure 7A.1 Histogram of computer generated data

In the data generation we first use a standard pseudorandom noise generator which produces independent $\mathcal{U}[0, 1]$ random variates. As an example, on a VAX 11/780 the intrinsic function RAN can be used. In MONTECARLO the subroutine RANDOM calls the RAN function N times (for N even). Next, we convert these independent uniformly distributed random variates into independent Gaussian random variates with mean 0 and variance 1 by using the Box-Mueller transformation

$$\begin{aligned} w_1 &= \sqrt{-2 \ln u_1} \cos 2\pi u_2 \\ w_2 &= \sqrt{-2 \ln u_1} \sin 2\pi u_2 \end{aligned}$$

where u_1, u_2 are independent $\mathcal{U}[0, 1]$ random variables and w_1, w_2 are independent $\mathcal{N}(0, 1)$ random variables. Since the transformation operates on two random variables at a time, N needs to be even (if odd, just increment by 1 and discard the extra random variate). To convert to $\mathcal{N}(0, \sigma^2)$ random variates simply multiply by σ . This entire procedure is implemented in subroutine WGN. Next, A is added to $w[n]$ to generate one time series realization $x[n]$ for $n = 0, 1, \dots, N-1$. For this realization of $x[n]$, $\hat{A}$ is computed. We repeat the procedure M times to yield the M realizations of $\hat{A}$.

The mean and variance are determined using (7.9) and (7.10), respectively. The alternative form of the variance estimate

$$\widehat{\text{var}}(\hat{A}) = \frac{1}{M} \sum_{i=1}^M \hat{A}_i^2 - \widehat{E}(\hat{A})^2$$

is used by the subroutine STATS. The number of realizations, M , needed to provide an accurate estimate of the mean and variance can be determined by finding the *variance of the estimators* since (7.9) and (7.10) are nothing more than estimators themselves. A simpler procedure is to keep increasing M until the numbers given by (7.9) and (7.10) converge.

Finally, to determine the PDF of $\hat{A}$ we use an estimated PDF called the *histogram*. The histogram estimates the PDF by determining the number of times $\hat{A}$ falls within a specified interval. Then, a division by the total number of realizations to yield the probability, followed by a division by the interval length, produces the PDF estimate. A typical histogram is shown in Figure 7A.1 where the PDF is estimated over the interval

$(x_{\min}, x_{\max})$. Each subinterval $(x_i - \Delta x/2, x_i + \Delta x/2)$ is called a *cell*. The value for the i th cell is found as

$$\widehat{p}(x_i) = \frac{L_i}{M \Delta x}$$

where L_i is the number of realizations of x_i that lie within the i th cell. Hence, L_i/M estimates the probability that x_i will fall in the i th cell, and by dividing by Δx , we obtain the estimated PDF. If we associate $\widehat{p}(x_i)$ with the PDF at the cell center, then interpolating these points as shown in Figure 7A.1 produces an estimate of the continuous PDF. It is apparent from this discussion that for good estimates of the PDF we would like the cell width to be small. This is because we are estimating not the PDF $p(x)$ but

$$\frac{1}{\Delta x} \int_{x_i - \frac{\Delta x}{2}}^{x_i + \frac{\Delta x}{2}} p(x) dx$$

or the average PDF over the cell. Unfortunately, as the cell width becomes smaller (more cells), the probability of a realization falling into the cell becomes smaller. This yields highly variable PDF estimates. As before, a good strategy keeps increasing M until the estimated PDF appears to converge. The histogram approach is implemented in the subroutine HISTOG. A further discussion of PDF estimation can be found in [Bendat and Piersol 1971].

Fortran Program MONTECARLO

```

C MONTECARLO
C This program determines the asymptotic properties
C of the MLE for a N(A,A) PDF (see Examples 7.1-7.3, 7.5).
C
C The array dimensions are given as variable. Replace them
C with numerical values. To use this program you will need
C a plotting subroutine to replace PLOT and a random number
C generator to replace the intrinsic function RAN in the
C subroutine RANDOM.
C
    DIMENSION X(N),W(N),AHAT(M),PDF(NCELLS),HIST(NCELLS)
    * ,XX(NCELLS)
    PI=4.*ATAN(1.)
C Input the value of A, the number of data points N, and the
C number of realizations M.
    WRITE(6,10)
10  FORMAT(' INPUT A, N, AND M')
    READ(5,*)A,N,M
C Generate M realizations of the estimate of A.
    DO 40 K=1,M
C Generate the noise samples and add A to each one.

```

```

      CALL WGN(N,A,W)
      DO 20 I=1,N
20    X(I)=A+W(I)
C   Generate the estimate of A.
      XSQ=0.
      DO 30 I=1,N
30    XSQ=XSQ+X(I)*X(I)/N
      AHAT(K)=-0.5+SQRT(XSQ+0.25)
40    CONTINUE
C   Compute the mean and variance of the estimates of A.
      CALL STATS(AHAT,M,AMEAN,AVAR)
C   Normalize the variance by N.
      AVAR=N*AVAR
      WRITE(6,*)AMEAN,AVAR
C   Input the interval (XMIN,XMAX) and the number of cells
C   for the histogram.
      WRITE(6,50)
50    FORMAT(' INPUT XMIN,XMAX,NCELLS FOR HISTOGRAM')
      READ(5,*)XMIN,XMAX,NCELLS
C   Compute the histogram.
      CALL HISTOG(AHAT,M,XMIN,XMAX,NCELLS,XX,HIST)
C   Compute the asymptotic variance.
      SIG2=A*A/(N*(A+0.5))
C   Compute the asymptotic Gaussian PDF for comparison to
C   the histogram.
      DO 60 I=1,NCELLS
      ARG=(XX(I)-A)*(XX(I)-A)/SIG2
      IF(ARG.GT.50.)PDF(I)=0.
      IF(ARG.LT.50.)PDF(I)=(1./SQRT(2.*PI*SIG2))*EXP(-0.5*ARG)
60    CONTINUE
C   Compare graphically the asymptotic PDF to the histogram.
      CALL PLOT(XX,HIST,PDF)
      STOP
      END
      SUBROUTINE STATS(X,N,XMEAN,XVAR)
C   This program computes the sample mean and variance
C   of a data set.
C
C   Input parameters:
C
C   X      - Array of dimension Nx1 containing data
C   N      - Number of data set
C
C   Output parameters:
C

```



```

C      XMEAN - Sample mean
C      XVAR  - Sample variance
C
      DIMENSION X(1)
      XMEAN=0.
      XSQ=0.
      DO 10 I=1,N
      XMEAN=XMEAN+X(I)/N
10     XSQ=XSQ+X(I)*X(I)/N
      XVAR=XSQ-XMEAN*XMEAN
      RETURN
      END
      SUBROUTINE HISTOG(X,N,XMIN,XMAX,NCELLS,XX,HIST)
C This subroutine computes a histogram of a set of data.
C
C Input parameters:
C
C      X      - Input data array of dimension Nx1
C      N      - Number of data samples
C      XMIN   - Minimum value for histogram horizontal axis
C      XMAX   - Maximum value for histogram horizontal axis
C      NCELLS - Number of cells desired
C
C Output parameters:
C
C      XX      - Array of dimension NCELLSx1 containing
C               cell centers
C      HIST    - Histogram values for each cell where HIST(I)
C               is the value for XX(I), I=1,2,...,NCELLS
C
      DIMENSION X(1),XX(1),HIST(1)
      CELLWID=(XMAX-XMIN)/FLOAT(NCELLS)
      DO 10 I=1,NCELLS
10     XX(I)=XMIN+CELLWID/2.+(I-1.)*CELLWID
      DO 30 K=1,NCELLS
      HIST(K)=0.
      CELLMIN=XX(K)-CELLWID/2.
      CELLMAX=XX(K)+CELLWID/2.
      DO 20 I=1,N
20     IF(X(I).GT.CELLMIN.AND.X(I).LT.CELLMAX)HIST(K)=
      * HIST(K)+1./(N*CELLWID)
30     CONTINUE
      RETURN
      END
      SUBROUTINE PLOT(X,Y1,Y2)

```

```

      DIMENSION X(1),Y1(1),Y2(1)
C  Replace this subroutine with any standard Fortran-compatible
C  plotting routine.
      RETURN
      END
      SUBROUTINE WGN(N,VAR,W)
C  This subroutine generates samples of zero mean white
C  Gaussian noise.
C
C  Input parameters:
C
C    N    - Number of noise samples desired
C    VAR  - Variance of noise desired
C
C  Output parameters:
C
C    W    - Array of dimension Nx1 containing noise samples
C
      DIMENSION W(1)
      PI=4.*ATAN(1.)
C  Add 1 to desired number of samples if N is odd.
      N1=N
      IF(MOD(N,2).NE.0)N1=N+1
C  Generate N1 independent and uniformly distributed random
C  variates on [0,1].
      CALL RANDOM(N1,VAR,W)
      L=N1/2
C  Convert uniformly distributed random variates to Gaussian
C  ones using a Box-Mueller transformation.
      DO 10 I=1,L
        U1=W(2*I-1)
        U2=W(2*I)
        TEMP=SQRT(-2.*ALOG(U1))
        W(2*I-1)=TEMP*COS(2.*PI*U2)*SQRT(VAR)
10    W(2*I)=TEMP*SIN(2.*PI*U2)*SQRT(VAR)
      RETURN
      END
      SUBROUTINE RANDOM(N1,VAR,W)
      DIMENSION W(1)
      DO 10 I=1,N1
C  For machines other than DEC VAX 11/780 replace RAN(ISEED)
C  with a random number generator.
10    W(I)=RAN(11111)
      RETURN
      END

```

Appendix 7B

Asymptotic PDF of MLE for a Scalar Parameter

We now give an outline of the proof of Theorem 7.1. A rigorous proof of consistency can be found in [Dudewicz 1976] and of the asymptotic Gaussian property in [Rao 1973]. To simplify the discussion the observations are assumed to be IID. We further assume the following regularity conditions.

1. The first-order and second-order derivatives of the log-likelihood function are well defined.
- 2.

$$E \left[\frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right] = 0.$$

We first show that the MLE is consistent. To do so we will need the following inequality (see [Dudewicz 1976] for proof), which is related to the Kullback-Leibler information.

$$\int \ln \left[\frac{p(x[n]; \theta_1)}{p(x[n]; \theta_2)} \right] p(x[n]; \theta_1) dx[n] \geq 0 \quad (7B.1)$$

with equality if and only if $\theta_1 = \theta_2$. Now, in maximizing the log-likelihood function, we are equivalently maximizing

$$\begin{aligned} \frac{1}{N} \ln p(\mathbf{x}; \theta) &= \frac{1}{N} \ln \prod_{n=0}^{N-1} p(x[n]; \theta) \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \ln p(x[n]; \theta). \end{aligned}$$

But as $N \rightarrow \infty$, this converges to the expected value by the law of large numbers. Hence, if θ_0 denotes the true value of θ , we have

$$\frac{1}{N} \sum_{n=0}^{N-1} \ln p(x[n]; \theta) \rightarrow \int \ln p(x[n]; \theta) p(x[n]; \theta_0) dx[n]. \quad (7B.2)$$

However, from (7B.1)

$$\int \ln [p(x[n]; \theta_1)] p(x[n]; \theta_1) dx[n] \geq \int \ln [p(x[n]; \theta_2)] p(x[n]; \theta_1) dx[n].$$

The right-hand side of (7B.2) is therefore maximized for $\theta = \theta_0$. By a suitable continuity argument, the left-hand side of (7B.2) or the normalized log-likelihood function must also be maximized for $\theta = \theta_0$ or as $N \rightarrow \infty$, the MLE is $\hat{\theta} = \theta_0$. Thus, the MLE is consistent.

To derive the asymptotic PDF of the MLE we first use a Taylor expansion about the true value of θ , which is θ_0 . Then, by the mean value theorem

$$\left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\hat{\theta}} = \left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\theta_0} + \left. \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}} (\hat{\theta} - \theta_0)$$

where $\theta_0 < \bar{\theta} < \hat{\theta}$. But

$$\left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\hat{\theta}} = 0$$

by the definition of the MLE, so that

$$0 = \left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\theta_0} + \left. \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}} (\hat{\theta} - \theta_0). \quad (7B.3)$$

Now consider $\sqrt{N}(\hat{\theta} - \theta_0)$, so that (7B.3) becomes

$$\sqrt{N}(\hat{\theta} - \theta_0) = \frac{\frac{1}{\sqrt{N}} \left. \frac{\partial \ln p(\mathbf{x}; \theta)}{\partial \theta} \right|_{\theta=\theta_0}}{-\frac{1}{N} \left. \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}}}. \quad (7B.4)$$

Due to the IID assumption

$$\left. \frac{1}{N} \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}} = \frac{1}{N} \sum_{n=0}^{N-1} \left. \frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}}$$

Since $\theta_0 < \bar{\theta} < \hat{\theta}$, we must also have $\bar{\theta} \rightarrow \theta_0$ due to the consistency of the MLE. Hence,

$$\begin{aligned} \left. \frac{1}{N} \frac{\partial^2 \ln p(\mathbf{x}; \theta)}{\partial \theta^2} \right|_{\theta=\bar{\theta}} &\rightarrow \left. \frac{1}{N} \sum_{n=0}^{N-1} \frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right|_{\theta=\theta_0} \\ &\rightarrow E \left[\left. \frac{\partial^2 \ln p(x[n]; \theta)}{\partial \theta^2} \right|_{\theta=\theta_0} \right] \\ &= -i(\theta_0) \end{aligned}$$

where the last convergence is due to the law of large numbers. Also, the numerator term is

$$\frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \left. \frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right|_{\theta=\theta_0}.$$

Now

$$\xi_n = \left. \frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right|_{\theta=\theta_0}$$

is a random variable, being a function of $x[n]$. Additionally, since the $x[n]$'s are IID, so are the ξ_n 's. By the central limit theorem the numerator term in (7B.4) has a PDF that converges to a Gaussian with mean

$$E \left[\left. \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right|_{\theta=\theta_0} \right] = 0$$

and variance

$$\begin{aligned} E \left[\left(\left. \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right|_{\theta=\theta_0} \right)^2 \right] &= \frac{1}{N} \sum_{n=0}^{N-1} E \left[\left(\left. \frac{\partial \ln p(x[n]; \theta)}{\partial \theta} \right)^2 \right|_{\theta=\theta_0} \right] \\ &= i(\theta_0) \end{aligned}$$

due to the independence of the random variables. We now resort to Slutsky's theorem [Bickel and Doksum 1977], which says that if the sequence of random variables x_n has the asymptotic PDF of the random variable x and the sequence of random variables y_n converges to a constant c , then x_n/y_n has the same asymptotic PDF as the random variable x/c . In our case,

$$\begin{aligned} x &\sim \mathcal{N}(0, i(\theta_0)) \\ y_n &\rightarrow c = i(\theta_0) \end{aligned}$$

so that (7B.4) becomes

$$\sqrt{N}(\hat{\theta} - \theta_0) \stackrel{a}{\sim} \mathcal{N}(0, i^{-1}(\theta_0))$$

or, equivalently,

$$\hat{\theta} \stackrel{a}{\sim} \mathcal{N}\left(\theta_0, \frac{1}{Ni(\theta_0)}\right)$$

or finally

$$\hat{\theta} \stackrel{a}{\sim} \mathcal{N}(\theta_0, I^{-1}(\theta_0)).$$

Appendix 7C

Derivation of Conditional Log-Likelihood for EM Algorithm Example

From (7.51) we have, upon noting the independent data set assumption,

$$\begin{aligned}
 \ln p_{\mathbf{y}}(\mathbf{y}; \boldsymbol{\theta}) &= \sum_{i=1}^p \ln p(\mathbf{y}_i; \theta_i) \\
 &= \sum_{i=1}^p \ln \left\{ \frac{1}{(2\pi\sigma_i^2)^{\frac{N}{2}}} \exp \left[-\frac{1}{2\sigma_i^2} \sum_{n=0}^{N-1} (y_i[n] - \cos 2\pi f_i n)^2 \right] \right\} \\
 &= c - \sum_{i=1}^p \frac{1}{2\sigma_i^2} \sum_{n=0}^{N-1} (y_i[n] - \cos 2\pi f_i n)^2 \\
 &= g(\mathbf{y}) + \sum_{i=1}^p \frac{1}{\sigma_i^2} \sum_{n=0}^{N-1} (y_i[n] \cos 2\pi f_i n - \frac{1}{2} \cos^2 2\pi f_i n)
 \end{aligned}$$

where c is a constant and $g(\mathbf{y})$ does not depend on the frequencies. Making the approximation that $\sum_{n=0}^{N-1} \cos^2 2\pi f_i n \approx N/2$ for f_i not near 0 or $1/2$, we have

$$\ln p_{\mathbf{y}}(\mathbf{y}; \boldsymbol{\theta}) = h(\mathbf{y}) + \sum_{i=1}^p \frac{1}{\sigma_i^2} \sum_{n=0}^{N-1} y_i[n] \cos 2\pi f_i n$$

or letting $\mathbf{c}_i = [1 \cos 2\pi f_i \dots \cos 2\pi f_i (N-1)]^T$, we have

$$\ln p_{\mathbf{y}}(\mathbf{y}; \boldsymbol{\theta}) = h(\mathbf{y}) + \sum_{i=1}^p \frac{1}{\sigma_i^2} \mathbf{c}_i^T \mathbf{y}_i. \tag{7C.1}$$

Concatenating the $\mathbf{c}_i$'s and $\mathbf{y}_i$'s as

$$\mathbf{c} = \begin{bmatrix} \frac{1}{\sigma_1^2} \mathbf{c}_1 \\ \frac{1}{\sigma_2^2} \mathbf{c}_2 \\ \vdots \\ \frac{1}{\sigma_p^2} \mathbf{c}_p \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_p \end{bmatrix}$$

we have

$$\ln p_y(\mathbf{y}; \boldsymbol{\theta}) = h(\mathbf{y}) + \mathbf{c}^T \mathbf{y}.$$

From (7.56) we write the conditional expectation as

$$\begin{aligned} U(\boldsymbol{\theta}, \boldsymbol{\theta}_k) &= E[\ln p_y(\mathbf{y}; \boldsymbol{\theta}) | \mathbf{x}; \boldsymbol{\theta}_k] \\ &= E(h(\mathbf{y}) | \mathbf{x}; \boldsymbol{\theta}_k) + \mathbf{c}^T E(\mathbf{y} | \mathbf{x}; \boldsymbol{\theta}_k). \end{aligned} \quad (7C.2)$$

Since we wish to maximize $U(\boldsymbol{\theta}, \boldsymbol{\theta}_k)$ with respect to $\boldsymbol{\theta}$, we can omit the expected value of $h(\mathbf{y})$ since it does not depend on $\boldsymbol{\theta}$. Now, $\mathbf{y}$ and $\mathbf{x}$ are jointly Gaussian since, from (7.52),

$$\mathbf{x} = \sum_{i=1}^p \mathbf{y}_i = [\mathbf{I} \ \mathbf{I} \ \dots \ \mathbf{I}] \mathbf{y}$$

where $\mathbf{I}$ is the $N \times N$ identity matrix and the transformation matrix is composed of p identity matrices. Using the standard result for conditional expectations of jointly Gaussian random vectors (see Appendix 10A), we have

$$E(\mathbf{y} | \mathbf{x}; \boldsymbol{\theta}_k) = E(\mathbf{y}) + \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})).$$

The means are given by

$$\begin{aligned} E(\mathbf{y}) &= \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \vdots \\ \mathbf{c}_p \end{bmatrix} \\ E(\mathbf{x}) &= \sum_{i=1}^p \mathbf{c}_i \end{aligned}$$

while the covariance matrices are

$$\begin{aligned} \mathbf{C}_{xx} &= \sigma^2 \mathbf{I} \\ \mathbf{C}_{yx} &= E \left(\begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \vdots \\ \mathbf{w}_p \end{bmatrix} \mathbf{w}^T \right) \end{aligned}$$

$$\begin{aligned}
&= E \left(\begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \vdots \\ \mathbf{w}_p \end{bmatrix} \left\{ [\mathbf{I} \ \mathbf{I} \ \dots \ \mathbf{I}] \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \vdots \\ \mathbf{w}_p \end{bmatrix} \right\}^T \right) \\
&= \begin{bmatrix} \sigma_1^2 \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \sigma_2^2 \mathbf{I} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \sigma_p^2 \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \mathbf{I} \\ \vdots \\ \mathbf{I} \end{bmatrix} \\
&= \begin{bmatrix} \sigma_1^2 \mathbf{I} \\ \sigma_2^2 \mathbf{I} \\ \vdots \\ \sigma_p^2 \mathbf{I} \end{bmatrix},
\end{aligned}$$

so that

$$\begin{aligned}
E(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_k) &= \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \vdots \\ \mathbf{c}_p \end{bmatrix} + \frac{1}{\sigma^2} \begin{bmatrix} \sigma_1^2 \mathbf{I} \\ \sigma_2^2 \mathbf{I} \\ \vdots \\ \sigma_p^2 \mathbf{I} \end{bmatrix} \left(\mathbf{x} - \sum_{i=1}^p \mathbf{c}_i \right) \\
&= \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \vdots \\ \mathbf{c}_p \end{bmatrix} + \begin{bmatrix} \frac{\sigma_1^2}{\sigma^2} (\mathbf{x} - \sum_{i=1}^p \mathbf{c}_i) \\ \frac{\sigma_2^2}{\sigma^2} (\mathbf{x} - \sum_{i=1}^p \mathbf{c}_i) \\ \vdots \\ \frac{\sigma_p^2}{\sigma^2} (\mathbf{x} - \sum_{i=1}^p \mathbf{c}_i) \end{bmatrix}
\end{aligned}$$

or

$$E(\mathbf{y}_i|\mathbf{x}; \boldsymbol{\theta}_k) = \mathbf{c}_i + \frac{\sigma_i^2}{\sigma^2} (\mathbf{x} - \sum_{i=1}^p \mathbf{c}_i) \quad i = 1, 2, \dots, p$$

where $\mathbf{c}_i$ is computed using $\boldsymbol{\theta}_k$. Note that $E(\mathbf{y}_i|\mathbf{x}; \boldsymbol{\theta}_k)$ can be thought of as an estimate of the $y_i[n]$ data set since, letting $\hat{\mathbf{y}}_i = E(\mathbf{y}_i|\mathbf{x}; \boldsymbol{\theta}_k)$,

$$\hat{y}_i[n] = \cos 2\pi f_{i_k} n + \frac{\sigma_i^2}{\sigma^2} \left(x[n] - \sum_{i=1}^p \cos 2\pi f_{i_k} n \right). \quad (7C.3)$$

Using (7C.2) and dropping the term that is independent of $\boldsymbol{\theta}$, we have

$$U'(\boldsymbol{\theta}, \boldsymbol{\theta}_k) = \sum_{i=1}^p \mathbf{c}_i^T \hat{\mathbf{y}}_i$$

which is maximized over θ by maximizing each term in the sum separately or

$$f_{i_{k+1}} = \arg \max_{f_i} \mathbf{c}_i^T \hat{\mathbf{y}}_i. \quad (7C.4)$$

Finally, since the σ_i^2 's are not unique, they can be chosen arbitrarily as long as (see (7.53))

$$\sum_{i=1}^p \sigma_i^2 = \sigma^2$$

or, equivalently,

$$\sum_{i=1}^p \beta_i = \sum_{i=1}^p \frac{\sigma_i^2}{\sigma^2} = 1.$$

Chapter 8

Least Squares

8.1 Introduction

In previous chapters we have attempted to find an optimal or nearly optimal (for large data records) estimator by considering the class of unbiased estimators and determining the one exhibiting minimum variance, the so-called MVU estimator. We now depart from this philosophy to investigate a class of estimators that in general have no optimality properties associated with them but make good sense for many problems of interest. This is the method of least squares and dates back to 1795 when Gauss used the method to study planetary motions. A salient feature of the method is that no probabilistic assumptions are made about the data, only a signal model is assumed. The advantage, then, is its broader range of possible applications. On the negative side, no claims about optimality can be made, and furthermore, the statistical performance cannot be assessed without some specific assumptions about the probabilistic structure of the data. Nonetheless, the least squares estimator is widely used in practice due to its ease of implementation, amounting to the minimization of a least squares error criterion.

8.2 Summary

The least squares approach to parameter estimation chooses θ to minimize (8.1), where the signal depends on θ . Linear versus nonlinear least squares problems are described in Section 8.3. The general linear least squares problem which minimizes (8.9) leads to the least squares estimator of (8.10) and the minimum least squares error of (8.11)–(8.13). A weighted least squares error criterion of (8.14) results in the estimator of (8.16) and the minimum least squares error of (8.17). A geometrical interpretation of least squares is described in Section 8.5 and leads to the important orthogonality principle. When the dimensionality of the vector parameter is not known, an order-recursive least squares approach can be useful. It computes the least squares estimator recursively as the number of unknown parameters increases. It is summarized by (8.28)–(8.31). If it is desired to update in time the least squares estimator as more data become




Figure 8.1 Least squares approach

available, then a sequential approach can be employed. It determines the least squares estimator based on the estimate at the previous time and the new data. The equations (8.46)–(8.48) summarize the calculations required. At times the parameter vector is constrained, as in (8.50). In such a case the constrained least squares estimator is given by (8.52). Nonlinear least squares is discussed in Section 8.9. Some methods for converting the problem to a linear one are described, followed by iterative minimization approaches if this is not possible. The two methods that are generally used are the Newton-Raphson iteration of (8.61) and the Gauss-Newton iteration of (8.62).

8.3 The Least Squares Approach

Our focus in determining a good estimator has been to find one that was unbiased and had minimum variance. In choosing the variance as our measure of goodness we implicitly sought to minimize the discrepancy (on the average) between our estimate and the true parameter value. In the least squares (LS) approach we attempt to minimize the squared difference between the given data $x[n]$ and the assumed signal or noiseless data. This is illustrated in Figure 8.1. The signal is generated by some model which in turn depends upon our unknown parameter θ . The signal $s[n]$ is purely deterministic. Due to observation noise or model inaccuracies we observe a perturbed version of $s[n]$, which we denote by $x[n]$. The least squares estimator (LSE) of θ chooses the value that

makes $s[n]$ closest to the observed data $x[n]$. Closeness is measured by the LS error criterion

$$J(\theta) = \sum_{n=0}^{N-1} (x[n] - s[n])^2 \quad (8.1)$$

where the observation interval is assumed to be $n = 0, 1, \dots, N-1$, and the dependence of J on θ is via $s[n]$. The value of θ that minimizes $J(\theta)$ is the LSE. *Note that no probabilistic assumptions have been made about the data $x[n]$.* The method is equally valid for Gaussian as well as non-Gaussian noise. Of course, the *performance* of the LSE will undoubtedly depend upon the properties of the corrupting noise as well as any modeling errors. LSEs are usually applied in situations where a precise statistical characterization of the data is unknown or where an optimal estimator cannot be found or may be too complicated to apply in practice.

Example 8.1 - DC Level Signal

Assume that the signal model in Figure 8.1 is $s[n] = A$ and we observe $x[n]$ for $n = 0, 1, \dots, N-1$. Then, according to the LS approach, we can estimate A by minimizing (8.1) or

$$J(A) = \sum_{n=0}^{N-1} (x[n] - A)^2.$$

Differentiating with respect to A and setting the result equal to zero produces

$$\begin{aligned} \hat{A} &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] \\ &= \bar{x} \end{aligned}$$

or the sample mean estimator. Our familiar estimator, however, cannot be claimed to be optimal in the MVU sense but only in that it minimizes the LS error. We know, however, from our previous discussions that if $x[n] = A + w[n]$, where $w[n]$ is zero mean WGN, then the LSE will also be the MVU estimator, but otherwise not. To underscore the potential difficulties, consider what would happen if the noise were not zero mean. Then, the sample mean estimator would actually be an estimator of $A + E(w[n])$ since $w[n]$ could be written as

$$w[n] = E(w[n]) + w'[n]$$

where $w'[n]$ is zero mean noise. The data are more appropriately described by

$$x[n] = A + E(w[n]) + w'[n].$$

It should be clear that in using this approach, it must be assumed that the observed data are composed of a deterministic signal and *zero mean* noise. If this is the case, the error $\epsilon[n] = x[n] - s[n]$ will tend to be zero on the average for the correct choice of the signal parameters. The minimization of (8.1) is then a reasonable approach. The reader might also consider what would happen if the assumed DC level signal model

were incorrect, as for instance if $x[n] = A + Bn + w[n]$ described the data. This modeling error would also cause the LSE to be biased. $\diamond$

Example 8.2 - Sinusoidal Frequency Estimation

Consider the signal model

$$s[n] = \cos 2\pi f_0 n$$

in which the frequency f_0 is to be estimated. The LSE is found by minimizing

$$J(f_0) = \sum_{n=0}^{N-1} (x[n] - \cos 2\pi f_0 n)^2.$$

In contrast to the DC level signal for which the minimum is easily found, here the LS error is highly nonlinear in f_0 . The minimization cannot be done in closed form. Since the error criterion is a quadratic function of the signal, a signal that is *linear* in the unknown parameter yields a quadratic function for J , as in the previous example. The minimization is then easily carried out. A signal model that is *linear in the unknown parameter* is said to generate a *linear least squares* problem. Otherwise, as in this example, the problem is a *nonlinear least squares* problem. Nonlinear LS problems are solved via grid searches or iterative minimization methods as described in Section 8.9. It should be noted that the signal itself need not be linear but only in the unknown parameter, as the next example illustrates. $\diamond$

Example 8.3 - Sinusoidal Amplitude Estimation

If the signal is $s[n] = A \cos 2\pi f_0 n$, where f_0 is *known* and A is to be estimated, then the LSE minimizes

$$J(A) = \sum_{n=0}^{N-1} (x[n] - A \cos 2\pi f_0 n)^2$$

over A . This is easily accomplished by differentiation since $J(A)$ is quadratic in A . This linear LS problem is clearly very desirable from a practical viewpoint. If, however, A were known and the frequency were to be estimated, the problem would be equivalent to that in Example 8.2, that is to say, it would be a nonlinear LS problem. A final possibility, in the vector parameter case, is that both A and f_0 might need to be estimated. Then, the error criterion

$$J(A, f_0) = \sum_{n=0}^{N-1} (x[n] - A \cos 2\pi f_0 n)^2$$

is quadratic in A but nonquadratic in f_0 . The net result is that J can be minimized in closed form with respect to A for a given f_0 , reducing the minimization of J to one

over f_0 only. This type of problem, in which the signal is linear in some parameters but nonlinear in others, is termed a *separable least squares* problem. In Section 8.9 we will discuss this further. $\diamond$

8.4 Linear Least Squares

In applying the linear LS approach for a scalar parameter we must assume that

$$s[n] = \theta h[n] \quad (8.2)$$

where $h[n]$ is a known sequence. (The reader may want to refer back to Chapter 6 on the BLUE for a comparison to similar signal models.) The LS error criterion becomes

$$J(\theta) = \sum_{n=0}^{N-1} (x[n] - \theta h[n])^2. \quad (8.3)$$

A minimization is readily shown to produce the LSE

$$\hat{\theta} = \frac{\sum_{n=0}^{N-1} x[n]h[n]}{\sum_{n=0}^{N-1} h^2[n]}. \quad (8.4)$$

The minimum LS error, obtained by substituting (8.4) into (8.3), is

$$\begin{aligned} J_{\min} = J(\hat{\theta}) &= \sum_{n=0}^{N-1} (x[n] - \hat{\theta}h[n])(x[n] - \hat{\theta}h[n]) \\ &= \sum_{n=0}^{N-1} x[n](x[n] - \hat{\theta}h[n]) - \underbrace{\hat{\theta} \sum_{n=0}^{N-1} h[n](x[n] - \hat{\theta}h[n])}_S \\ &= \sum_{n=0}^{N-1} x^2[n] - \hat{\theta} \sum_{n=0}^{N-1} x[n]h[n]. \end{aligned} \quad (8.5)$$

The last step follows because the sum S is zero (substitute $\hat{\theta}$ to verify). Alternatively, by using (8.4) we can rewrite $J_{\min}$ as

$$J_{\min} = \sum_{n=0}^{N-1} x^2[n] - \frac{\left(\sum_{n=0}^{N-1} x[n]h[n] \right)^2}{\sum_{n=0}^{N-1} h^2[n]}. \quad (8.6)$$

The minimum LS error is the original energy of the data or $\sum_{n=0}^{N-1} x^2[n]$ less that due to the signal fitting. For Example 8.1 in which $\theta = A$ we have $h[n] = 1$, so that from (8.4) $\hat{A} = \bar{x}$ and from (8.5)

$$J_{\min} = \sum_{n=0}^{N-1} x^2[n] - N\bar{x}^2.$$

If the data were noiseless so that $x[n] = A$, then $J_{\min} = 0$ or we would have a perfect LS fit to the data. On the other hand, if $x[n] = A + w[n]$, where $E(w^2[n]) \gg A^2$, then $\sum_{n=0}^{N-1} x^2[n]/N \gg \bar{x}^2$. The minimum LS error would then be

$$J_{\min} \approx \sum_{n=0}^{N-1} x^2[n]$$

or not much different than the original error. It can be shown (see Problem 8.2) that the minimum LS error is always between these two extremes or

$$0 \leq J_{\min} \leq \sum_{n=0}^{N-1} x^2[n]. \quad (8.7)$$

The extension of these results to a vector parameter θ of dimension $p \times 1$ is straightforward and of great practical utility. For the signal $s = [s[0] s[1] \dots s[N-1]]^T$ to be linear in the unknown parameters, we assume, using matrix notation,

$$s = \mathbf{H}\theta \quad (8.8)$$

where $\mathbf{H}$ is a known $N \times p$ matrix ($N > p$) of full rank p . The matrix $\mathbf{H}$ is referred to as the *observation matrix*. This is, of course, the linear model, albeit without the usual noise PDF assumption. Many examples of signals satisfying this model can be found in Chapter 4. The LSE is found by minimizing

$$\begin{aligned} J(\theta) &= \sum_{n=0}^{N-1} (x[n] - s[n])^2 \\ &= (\mathbf{x} - \mathbf{H}\theta)^T (\mathbf{x} - \mathbf{H}\theta). \end{aligned} \quad (8.9)$$

This is easily accomplished (since J is a quadratic function of θ) by using (4.3). Since

$$\begin{aligned} J(\theta) &= \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \mathbf{H}\theta - \theta^T \mathbf{H}^T \mathbf{x} + \theta^T \mathbf{H}^T \mathbf{H}\theta \\ &= \mathbf{x}^T \mathbf{x} - 2\mathbf{x}^T \mathbf{H}\theta + \theta^T \mathbf{H}^T \mathbf{H}\theta \end{aligned}$$

(note that $\mathbf{x}^T \mathbf{H}\theta$ is a scalar), the gradient is

$$\frac{\partial J(\theta)}{\partial \theta} = -2\mathbf{H}^T \mathbf{x} + 2\mathbf{H}^T \mathbf{H}\theta.$$

Setting the gradient equal to zero yields the LSE

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}. \quad (8.10)$$

The equations $\mathbf{H}^T \mathbf{H} \boldsymbol{\theta} = \mathbf{H}^T \mathbf{x}$ to be solved for $\hat{\boldsymbol{\theta}}$ are termed the *normal equations*. The assumed full rank of $\mathbf{H}$ guarantees the invertibility of $\mathbf{H}^T \mathbf{H}$. See also Problem 8.4 for another derivation. Somewhat surprisingly, we obtain an estimator that has the *identical functional form* as the efficient estimator for the linear model as well as the BLUE. That $\hat{\boldsymbol{\theta}}$ as given by (8.10) is not the *identical estimator* stems from the assumptions made about the data. For it to be the BLUE would require $E(\mathbf{x}) = \mathbf{H}\boldsymbol{\theta}$ and $\mathbf{C}_x = \sigma^2 \mathbf{I}$ (see Chapter 6), and to be efficient would in addition to these properties require $\mathbf{x}$ to be Gaussian (see Chapter 4). As a side issue, if these assumptions hold, we can easily determine the statistical properties of the LSE (see Problem 8.6), having been given in Chapters 4 and 6. Otherwise, this may be quite difficult. The minimum LS error is found from (8.9) and (8.10) as

$$\begin{aligned} J_{\min} &= J(\hat{\boldsymbol{\theta}}) \\ &= (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}) \\ &= (\mathbf{x} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x})^T (\mathbf{x} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}) \\ &= \mathbf{x}^T (\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T) (\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T) \mathbf{x} \\ &= \mathbf{x}^T (\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T) \mathbf{x}. \end{aligned} \quad (8.11)$$

The last step results from the fact that $\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T$ is an idempotent matrix or it has the property $\mathbf{A}^2 = \mathbf{A}$. Other forms for $J_{\min}$ are

$$J_{\min} = \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \quad (8.12)$$

$$= \mathbf{x}^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}). \quad (8.13)$$

An extension of the linear LS problem is to *weighted* LS. Instead of minimizing (8.9), we include an $N \times N$ positive definite (and by definition therefore symmetric) weighting matrix $\mathbf{W}$, so that

$$J(\boldsymbol{\theta}) = (\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{W} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}). \quad (8.14)$$

If, for instance, $\mathbf{W}$ is diagonal with diagonal elements $[\mathbf{W}]_{ii} = w_i > 0$, then the LS error for Example 8.1 will be

$$J(A) = \sum_{n=0}^{N-1} w_n (x[n] - A)^2.$$

The rationale for introducing weighting factors into the error criterion is to emphasize the contributions of those data samples that are deemed to be more reliable. Again, considering Example 8.1, if $x[n] = A + w[n]$, where $w[n]$ is zero mean uncorrelated noise with variance σ_n^2 , then it is reasonable to choose $w_n = 1/\sigma_n^2$. This choice will result in

the estimator (see Problem 8.8)

$$\hat{A} = \frac{\sum_{n=0}^{N-1} \frac{x[n]}{\sigma_n^2}}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2}}. \quad (8.15)$$

This familiar estimator is of course the BLUE since the $w[n]$'s are uncorrelated so that $\mathbf{W} = \mathbf{C}^{-1}$ (see Example 6.2).

The general form of the weighted LSE is readily shown to be

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{W} \mathbf{x} \quad (8.16)$$

and its minimum LS error is

$$J_{\min} = \mathbf{x}^T (\mathbf{W} - \mathbf{W} \mathbf{H} (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{W}) \mathbf{x} \quad (8.17)$$

(see Problem 8.9).

8.5 Geometrical Interpretations

We now reexamine the linear LS approach from a geometrical perspective. This has the advantage of more clearly revealing the essence of the approach and leads to additional useful properties and insights into the estimator. Recall the general signal model $\mathbf{s} = \mathbf{H}\boldsymbol{\theta}$. If we denote the columns of $\mathbf{H}$ by $\mathbf{h}_i$, we have

$$\begin{aligned} \mathbf{s} &= [\mathbf{h}_1 \quad \mathbf{h}_2 \quad \dots \quad \mathbf{h}_p] \begin{bmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_p \end{bmatrix} \\ &= \sum_{i=1}^p \theta_i \mathbf{h}_i \end{aligned}$$

so that the signal model is seen to be a linear combination of the “signal” vectors $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_p\}$.

Example 8.4 - Fourier Analysis

Referring to Example 4.2 (with $M = 1$), we suppose the signal model to be

$$s[n] = a \cos 2\pi f_0 n + b \sin 2\pi f_0 n \quad n = 0, 1, \dots, N-1$$

where f_0 is a known frequency and $\theta = [a \ b]^T$ is to be estimated. Then, in vector form we have

$$\begin{bmatrix} s[0] \\ s[1] \\ \vdots \\ s[N-1] \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \cos 2\pi f_0 & \sin 2\pi f_0 \\ \vdots & \vdots \\ \cos 2\pi f_0(N-1) & \sin 2\pi f_0(N-1) \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix}. \quad (8.18)$$

It is seen that the columns of $\mathbf{H}$ are composed of the samples of the cosinusoidal and sinusoidal sequences. Alternatively, since

$$\begin{aligned} \mathbf{h}_1 &= [1 \ \cos 2\pi f_0 \ \dots \ \cos 2\pi f_0(N-1)]^T \\ \mathbf{h}_2 &= [0 \ \sin 2\pi f_0 \ \dots \ \sin 2\pi f_0(N-1)]^T, \end{aligned}$$

we have

$$\mathbf{s} = a\mathbf{h}_1 + b\mathbf{h}_2.$$

◇

The LS error was defined to be

$$J(\theta) = (\mathbf{x} - \mathbf{H}\theta)^T(\mathbf{x} - \mathbf{H}\theta).$$

If we further define the Euclidean length of an $N \times 1$ vector $\xi = [\xi_1 \ \xi_2 \ \dots \ \xi_N]^T$ as

$$\|\xi\| = \sqrt{\sum_{i=1}^N \xi_i^2} = \sqrt{\xi^T \xi},$$

then the LS error can be also written as

$$\begin{aligned} J(\theta) &= \|\mathbf{x} - \mathbf{H}\theta\|^2 \\ &= \|\mathbf{x} - \sum_{i=1}^p \theta_i \mathbf{h}_i\|^2. \end{aligned} \quad (8.19)$$

We now see that the linear LS approach attempts to minimize the square of the distance from the data vector $\mathbf{x}$ to a signal vector $\sum_{i=1}^p \theta_i \mathbf{h}_i$, which must be a linear combination of the columns of $\mathbf{H}$. The data vector can lie anywhere in an N -dimensional space, termed R^N , while all possible signal vectors, being linear combinations of $p < N$ vectors, must lie in a p -dimensional subspace of R^N , termed S^p . (The full rank of $\mathbf{H}$ assumption assures us that the columns are linearly independent and hence the subspace spanned is truly p -dimensional.) For $N = 3$ and $p = 2$ we illustrate this in Figure 8.2. Note that all possible choices of θ_1, θ_2 (where we assume $-\infty < \theta_1 < \infty$ and $-\infty < \theta_2 < \infty$) produce signal vectors constrained to lie in the subspace S^2 and that in general $\mathbf{x}$ does not lie in the subspace. It should be intuitively clear that the vector $\hat{\mathbf{s}}$ that lies in S^2 and




Figure 8.2 Geometrical viewpoint of linear least squares in $\mathbb{R}^3$

that is closest to $\mathbf{x}$ in the Euclidean sense is the component of $\mathbf{x}$ in S^2 . Alternatively, $\hat{\mathbf{s}}$ is the *orthogonal projection* of $\mathbf{x}$ onto S^2 . This means that the error vector $\mathbf{x} - \hat{\mathbf{s}}$ must be orthogonal to all vectors in S^2 . Two vectors in $\mathbb{R}^N$ are defined to be orthogonal if $\mathbf{x}^T \mathbf{y} = 0$. To actually determine $\hat{\mathbf{s}}$ for this example we use the orthogonality condition. This says that the error vector is orthogonal to the signal subspace or

$$(\mathbf{x} - \hat{\mathbf{s}}) \perp S^2$$

where $\perp$ denotes orthogonal (or perpendicular). For this to be true we must have

$$(\mathbf{x} - \hat{\mathbf{s}}) \perp \mathbf{h}_1$$

$$(\mathbf{x} - \hat{\mathbf{s}}) \perp \mathbf{h}_2$$

since then the error vector will be orthogonal to any linear combination of $\mathbf{h}_1$ and $\mathbf{h}_2$. Using the definition of orthogonality, we have

$$(\mathbf{x} - \hat{\mathbf{s}})^T \mathbf{h}_1 = 0$$

$$(\mathbf{x} - \hat{\mathbf{s}})^T \mathbf{h}_2 = 0.$$

Letting $\hat{\mathbf{s}} = \theta_1 \mathbf{h}_1 + \theta_2 \mathbf{h}_2$, we have

$$(\mathbf{x} - \theta_1 \mathbf{h}_1 - \theta_2 \mathbf{h}_2)^T \mathbf{h}_1 = 0$$

$$(\mathbf{x} - \theta_1 \mathbf{h}_1 - \theta_2 \mathbf{h}_2)^T \mathbf{h}_2 = 0.$$

In matrix form this is

$$(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{h}_1 = 0$$

$$(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{h}_2 = 0.$$

Combining the two equations yields

$$(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T [\mathbf{h}_1 \quad \mathbf{h}_2] = \mathbf{0}^T$$

or

$$(\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{H} = \mathbf{0}^T. \quad (8.20)$$

Finally, we have as our LSE

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

Note that if $\boldsymbol{\epsilon} = \mathbf{x} - \mathbf{H}\boldsymbol{\theta}$ denotes the error vector, then the LSE is found from (8.20) by invoking the condition

$$\boldsymbol{\epsilon}^T \mathbf{H} = \mathbf{0}^T. \quad (8.21)$$

The error vector must be orthogonal to the columns of $\mathbf{H}$. This is the well-known *orthogonality principle*. In effect, the error represents the part of $\mathbf{x}$ that cannot be described by the signal model. A similar orthogonality principle will arise in Chapter 12 in our study of estimation of random parameters.

Again referring to Figure 8.2b, the minimum LS error is $\|\mathbf{x} - \hat{\mathbf{s}}\|^2$ or

$$\begin{aligned} \|\mathbf{x} - \hat{\mathbf{s}}\|^2 &= \|\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}\|^2 \\ &= (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}). \end{aligned}$$

In evaluating this error we can make use of (8.21) (we have already done so for the scalar case in arriving at (8.5)). This produces

$$\begin{aligned} J_{\min} &= (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}) \\ &= \mathbf{x}^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}) - \hat{\boldsymbol{\theta}}^T \mathbf{H}^T \boldsymbol{\epsilon} \\ &= \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \mathbf{H}\hat{\boldsymbol{\theta}} \\ &= \mathbf{x}^T (\mathbf{I} - \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T) \mathbf{x}. \end{aligned} \quad (8.22)$$

In summary, the LS approach can be interpreted as the problem of fitting or approximating a data vector $\mathbf{x}$ in R^N by another vector $\hat{\mathbf{s}}$, which is a linear combination of vectors $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_p\}$ that lie in a p -dimensional subspace of R^N . The problem is solved by choosing $\hat{\mathbf{s}}$ in the subspace to be the orthogonal projection of $\mathbf{x}$. Many of our intuitive notions about vector geometry may be used to our advantage once this connection is made. We now discuss some of these consequences.

Referring to Figure 8.3a, if it had happened that $\mathbf{h}_1$ and $\mathbf{h}_2$ were orthogonal, then $\hat{\mathbf{s}}$ could have easily been found. This is because the component of $\hat{\mathbf{s}}$ along $\mathbf{h}_1$ or $\hat{\mathbf{s}}_1$ does not contain a component of $\hat{\mathbf{s}}$ along $\mathbf{h}_2$. If it did, then we would have the situation in Figure 8.3b. Making the orthogonality assumption and also assuming that $\|\mathbf{h}_1\| = \|\mathbf{h}_2\| = 1$ (orthonormal vectors), we have

$$\begin{aligned} \hat{\mathbf{s}} &= \hat{\mathbf{s}}_1 + \hat{\mathbf{s}}_2 \\ &= (\mathbf{h}_1^T \mathbf{x}) \mathbf{h}_1 + (\mathbf{h}_2^T \mathbf{x}) \mathbf{h}_2 \end{aligned}$$




Figure 8.3 Effect of nonorthogonal columns of observation matrix

where $\mathbf{h}_i^T \mathbf{x}$ is the length of the vector $\mathbf{x}$ along $\mathbf{h}_i$. In matrix notation this is

$$\begin{aligned}
 \hat{\mathbf{s}} &= \begin{bmatrix} \mathbf{h}_1 & \mathbf{h}_2 \end{bmatrix} \begin{bmatrix} \mathbf{h}_1^T \mathbf{x} \\ \mathbf{h}_2^T \mathbf{x} \end{bmatrix} \\
 &= \begin{bmatrix} \mathbf{h}_1 & \mathbf{h}_2 \end{bmatrix} \begin{bmatrix} \mathbf{h}_1^T \\ \mathbf{h}_2^T \end{bmatrix} \mathbf{x} \\
 &= \mathbf{H} \mathbf{H}^T \mathbf{x}
 \end{aligned}$$

so that

$$\hat{\boldsymbol{\theta}} = \mathbf{H}^T \mathbf{x}.$$

This result is due to the orthonormal columns of $\mathbf{H}$. As a result, we have

$$(\mathbf{H}^T \mathbf{H})^{-1} = (\mathbf{I})^{-1} = \mathbf{I}$$

and therefore

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} = \mathbf{H}^T \mathbf{x}.$$

No inversion is necessary. An example follows.

Example 8.5 - Fourier Analysis (continued)

Continuing Example 8.4, if $f_0 = k/N$, where k is an integer taking on any of the values $k = 1, 2, \dots, N/2 - 1$, it is easily shown (see (4.13)) that

$$\mathbf{h}_1^T \mathbf{h}_2 = \sum_{n=0}^{N-1} \cos\left(2\pi \frac{k}{N} n\right) \sin\left(2\pi \frac{k}{N} n\right) = 0$$

and also

$$\begin{aligned}\mathbf{h}_1^T \mathbf{h}_1 &= \frac{N}{2} \\ \mathbf{h}_2^T \mathbf{h}_2 &= \frac{N}{2}\end{aligned}$$

so that $\mathbf{h}_1$ and $\mathbf{h}_2$ are orthogonal but not orthonormal. Combining these results produces $\mathbf{H}^T \mathbf{H} = (N/2)\mathbf{I}$, and therefore,

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \begin{bmatrix} \hat{a} \\ \hat{b} \end{bmatrix} \\ &= (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \\ &= \frac{2}{N} \mathbf{H}^T \mathbf{x} \\ &= \begin{bmatrix} \frac{2}{N} \sum_{n=0}^{N-1} x[n] \cos\left(2\pi \frac{k}{N} n\right) \\ \frac{2}{N} \sum_{n=0}^{N-1} x[n] \sin\left(2\pi \frac{k}{N} n\right) \end{bmatrix}.\end{aligned}$$

If we had instead defined the signal as

$$s[n] = a' \sqrt{\frac{2}{N}} \cos\left(2\pi \frac{k}{N} n\right) + b' \sqrt{\frac{2}{N}} \sin\left(2\pi \frac{k}{N} n\right)$$

then the columns of $\mathbf{H}$ would have been orthonormal. ◇

In general, the columns of $\mathbf{H}$ will not be orthogonal, so that the signal vector estimate is obtained as

$$\hat{\mathbf{s}} = \mathbf{H} \hat{\boldsymbol{\theta}} = \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

The signal estimate is the orthogonal projection of $\mathbf{x}$ onto the p -dimensional subspace. The $N \times N$ matrix $\mathbf{P} = \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T$ is known as the *orthogonal projection matrix* or just the *projection matrix*. It has the properties

1. $\mathbf{P}^T = \mathbf{P}$, symmetric

2. $\mathbf{P}^2 = \mathbf{P}$, idempotent.

That the projection matrix must be symmetric is shown in Problem 8.11, that it must be idempotent follows from the observation that if $\mathbf{P}$ is applied to $\mathbf{P}\mathbf{x}$, then the same vector must result since $\mathbf{P}\mathbf{x}$ is already in the subspace. Additionally, the projection matrix must be singular (for independent columns of $\mathbf{H}$ it has rank p , as shown in Problem 8.12). If it were not, then $\mathbf{x}$ could be recovered from $\hat{\mathbf{s}}$, which is clearly impossible since many $\mathbf{x}$'s have the same projection, as shown in Figure 8.4.




Figure 8.4 Vectors with same projection onto subspace

Likewise, the error vector $\epsilon = \mathbf{x} - \hat{\mathbf{s}} = (\mathbf{I} - \mathbf{P})\mathbf{x}$ is the projection of $\mathbf{x}$ onto the *complement* subspace or the subspace orthogonal to the signal subspace. The matrix $\mathbf{P}^\perp = \mathbf{I} - \mathbf{P}$ is also a projection matrix, as can be easily verified from the properties given above. As a result, the minimum LS error is, from (8.22),

$$\begin{aligned} J_{\min} &= \mathbf{x}^T (\mathbf{I} - \mathbf{P}) \mathbf{x} \\ &= \mathbf{x}^T \mathbf{P}^\perp \mathbf{x} \\ &= \mathbf{x}^T \mathbf{P}^{\perp T} \mathbf{P}^\perp \mathbf{x} \\ &= \|\mathbf{P}^\perp \mathbf{x}\|^2 \end{aligned}$$

which is just $\|\epsilon\|^2$.

In the next section we further utilize the geometrical theory to derive an order-recursive LS solution.

8.6 Order-Recursive Least Squares

In many cases the signal model is unknown and must be assumed. For example, consider the experimental data shown in Figure 8.5a. The following models might be assumed

$$\begin{aligned} s_1(t) &= A \\ s_2(t) &= A + Bt \end{aligned}$$

for $0 \leq t \leq T$. If data $x[n]$ are obtained by sampling $x(t)$ at times $t = n\Delta$, where $\Delta = 1$ and $n = 0, 1, \dots, N-1$, the corresponding discrete-time signal models become

$$\begin{aligned} s_1[n] &= A \\ s_2[n] &= A + Bn. \end{aligned}$$

Using a LSE with

$$\mathbf{H}_1 = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} \quad \mathbf{H}_2 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ \vdots & \vdots \\ 1 & N-1 \end{bmatrix}$$






Figure 8.5 Experimental data fitting by least squares




Figure 8.6 Effect of chosen number of parameters on minimum least squares error

would produce the estimates for the intercept and slope as

$$\hat{A}_1 = \bar{x} \quad (8.23)$$

and (see Problem 8.13)

$$\begin{aligned} \hat{A}_2 &= \frac{2(2N-1)}{N(N+1)} \sum_{n=0}^{N-1} x[n] - \frac{6}{N(N+1)} \sum_{n=0}^{N-1} nx[n] \\ \hat{B}_2 &= -\frac{6}{N(N+1)} \sum_{n=0}^{N-1} x[n] + \frac{12}{N(N^2-1)} \sum_{n=0}^{N-1} nx[n]. \end{aligned} \quad (8.24)$$

In Figures 8.5b and 8.5c we have plotted $\hat{s}_1(t) = \hat{A}_1$ and $\hat{s}_2(t) = \hat{A}_2 + \hat{B}_2 t$, where $T = 100$. The fit using two parameters is better, as expected. We will later show that the minimum LS error must decrease as we add more parameters. The question might be asked whether or not we should add a quadratic term to the signal model. If we did, the fit would become better yet. Realizing that the data are subject to error, we may very well be *fitting the noise*. Such a situation is undesirable but in the absence of a known signal model may be unavoidable to some extent. In practice, we choose the simplest signal model that adequately describes the data. One such scheme might be to increase the order of the polynomial until the minimum LS error decreases only slightly as the order is increased. For the data in Figure 8.5 the minimum LS error versus the number of parameters is shown in Figure 8.6. Models $s_1(t)$ and $s_2(t)$ correspond to $k = 1$ and $k = 2$, respectively. It is seen that a large drop occurs at $k = 2$. For larger orders there is only a slight decrease, indicating a modeling of the noise. In this example the signal was actually

$$s(t) = 1 + 0.03t$$

and the noise $w[n]$ was WGN with variance $\sigma^2 = 0.1$. Note that if A, B had been estimated perfectly, the minimum LS error would have been

$$\begin{aligned} J_{\min} &= J(\hat{A}, \hat{B}) \\ &= J(A, B) \\ &= \sum_{n=0}^{N-1} (x[n] - s[n])^2 \\ &= \sum_{n=0}^{N-1} w^2[n] \\ &\approx N\sigma^2 \end{aligned}$$

so that when the true order is reached, $J_{\min} \approx 10$. This is verified in Figure 8.6 and also increases our confidence in the chosen model.

In the preceding example we saw that it was important to be able to determine the LSE for several signal models. A straightforward approach would compute the LSE for each model using (8.10). Alternatively, the computation may be reduced by using an *order-recursive* LS approach. In this method we *update* the LSE in order. Specifically, we are able to compute the LSE based on an $\mathbf{H}$ of dimension $N \times (k+1)$ from the solution based on an $\mathbf{H}$ of dimension $N \times k$. For the previous example this update in order would have been exceedingly simple had the columns of $\mathbf{H}_2$ been orthogonal. To see why assume the signal models to be

$$\begin{aligned} s_1[n] &= A_1 \\ s_2[n] &= A_2 + B_2 n \end{aligned}$$

for $-M \leq n \leq M$. We have now altered the observation interval to be the symmetric interval $[-M, M]$ as opposed to the original $[0, N-1]$ interval. The effect of this assumption is to orthogonalize the columns of $\mathbf{H}_2$ since now

$$\mathbf{H}_2 = \begin{bmatrix} 1 & -M \\ 1 & -(M-1) \\ \vdots & \vdots \\ 1 & M \end{bmatrix}.$$

The LSE is readily found since

$$\mathbf{H}_2^T \mathbf{H}_2 = \begin{bmatrix} 2M+1 & 0 \\ 0 & \sum_{n=-M}^M n^2 \end{bmatrix}$$

is a diagonal matrix. The solutions are

$$\hat{A}_1 = \frac{1}{2M+1} \sum_{n=-M}^M x[n]$$




Figure 8.7 Gram-Schmidt orthogonalization

and

$$\hat{A}_2 = \frac{1}{2M+1} \sum_{n=-M}^M x[n]$$

$$\hat{B}_2 = \frac{\sum_{n=-M}^M nx[n]}{\sum_{n=-M}^M n^2}.$$

In this case the LSE of A does not change as we add a parameter to the model and follows algebraically from the diagonal nature of $\mathbf{H}_2^T \mathbf{H}_2$. In geometric terms this result is readily apparent from Figure 8.3a. The orthogonality of the $\mathbf{h}_i$'s allows us to project $\mathbf{x}$ along $\mathbf{h}_1$ and $\mathbf{h}_2$ separately and then add the results. Each projection is independent of the other. In general, the column vectors will not be orthogonal but can be replaced by another set of p vectors that are orthogonal. The procedure for doing so is shown in Figure 8.7 and is called a *Gram-Schmidt orthogonalization*. The new column $\mathbf{h}_2$ is projected onto the subspace orthogonal to $\mathbf{h}_1$. Then, since $\mathbf{h}_1$ and $\mathbf{h}_2'$ are orthogonal, the LS signal estimate becomes

$$\hat{\mathbf{s}} = \mathbf{h}_1 \hat{\theta}_1 + \mathbf{h}_2' \hat{\theta}_2'$$

where $\mathbf{h}_1 \hat{\theta}_1$ is also the LSE for the signal based on $\mathbf{H} = \mathbf{h}_1$ only. This procedure recursively adds terms to the signal model; it updates the order. In Problem 8.14 this geometrical viewpoint is used together with a Gram-Schmidt orthogonalization to derive the order-update equations. A purely algebraic derivation is given in Appendix 8A. It is now summarized.

Denote the $N \times k$ observation matrix as $\mathbf{H}_k$, and the LSE based on $\mathbf{H}_k$ as $\hat{\theta}_k$ or

$$\hat{\theta}_k = (\mathbf{H}_k^T \mathbf{H}_k)^{-1} \mathbf{H}_k^T \mathbf{x}. \quad (8.25)$$

The minimum LS error based on $\mathbf{H}_k$ is

$$J_{\min k} = (\mathbf{x} - \mathbf{H}_k \hat{\theta}_k)^T (\mathbf{x} - \mathbf{H}_k \hat{\theta}_k). \quad (8.26)$$

By increasing the order to $k + 1$ we add a column to the observation matrix. This generates a new observation matrix which in partitioned form is

$$\mathbf{H}_{k+1} = \begin{bmatrix} \mathbf{H}_k & \mathbf{h}_{k+1} \end{bmatrix} = \begin{bmatrix} N \times k & N \times 1 \end{bmatrix}. \quad (8.27)$$

To update $\hat{\theta}_k$ and $J_{\min k}$ we use

$$\hat{\theta}_{k+1} = \begin{bmatrix} \hat{\theta}_k - \frac{(\mathbf{H}_k^T \mathbf{H}_k)^{-1} \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ \frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} = \begin{bmatrix} k \times 1 \\ 1 \times 1 \end{bmatrix} \quad (8.28)$$

where

$$\mathbf{P}_k^\perp = \mathbf{I} - \mathbf{H}_k (\mathbf{H}_k^T \mathbf{H}_k)^{-1} \mathbf{H}_k^T$$

is the projection matrix onto the subspace orthogonal to that spanned by the columns of $\mathbf{H}_k$. To avoid inverting $\mathbf{H}_k^T \mathbf{H}_k$ we let

$$\mathbf{D}_k = (\mathbf{H}_k^T \mathbf{H}_k)^{-1} \quad (8.29)$$

and use the recursive formula

$$\begin{aligned} \mathbf{D}_{k+1} &= \begin{bmatrix} \mathbf{D}_k + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & -\frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ -\frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & \frac{1}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} \\ &= \begin{bmatrix} k \times k & k \times 1 \\ 1 \times k & 1 \times 1 \end{bmatrix} \end{aligned} \quad (8.30)$$

where $\mathbf{P}_k^\perp = \mathbf{I} - \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T$. The minimum LS error is updated by using

$$J_{\min k+1} = J_{\min k} - \frac{(\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x})^2}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}. \quad (8.31)$$

The entire algorithm requires no matrix inversions. The recursion begins by determining $\hat{\theta}_1$, $J_{\min 1}$, and $\mathbf{D}_1$ using (8.25), (8.26), and (8.29), respectively. We term (8.28), (8.30), and (8.31) the *order-recursive least squares method*. To illustrate the computations involved we apply the method to the previous line fitting example.

Example 8.6 - Line Fitting

Since $s_1[n] = A_1$ and $s_2[n] = A_2 + B_2 n$ for $n = 0, 1, \dots, N-1$, we have

$$\begin{aligned} \mathbf{H}_1 &= \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \mathbf{1} \\ \mathbf{H}_2 &= \begin{bmatrix} \mathbf{H}_1 & \mathbf{h}_2 \end{bmatrix} \end{aligned}$$

where

$$\mathbf{h}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ N-1 \end{bmatrix}.$$

Using (8.25) and (8.26) to begin the recursion

$$\begin{aligned} \hat{A}_1 &= \hat{\theta}_1 \\ &= (\mathbf{H}_1^T \mathbf{H}_1)^{-1} \mathbf{H}_1^T \mathbf{x} \\ &= \bar{x} \end{aligned}$$

and

$$\begin{aligned} J_{\min_1} &= (\mathbf{x} - \mathbf{H}_1 \hat{\theta}_1)^T (\mathbf{x} - \mathbf{H}_1 \hat{\theta}_1) \\ &= \sum_{n=0}^{N-1} (x[n] - \bar{x})^2. \end{aligned}$$

Next, we use (8.28) to find $\hat{\theta}_2 = [\hat{A}_2 \hat{B}_2]^T$ or

$$\hat{\theta}_2 = \begin{bmatrix} \hat{\theta}_1 - \frac{(\mathbf{H}_1^T \mathbf{H}_1)^{-1} \mathbf{H}_1^T \mathbf{h}_2 \mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{x}}{\mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{h}_2} \\ \frac{\mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{x}}{\mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{h}_2} \end{bmatrix}.$$

The necessary terms are

$$\begin{aligned} (\mathbf{H}_1^T \mathbf{H}_1)^{-1} &= \frac{1}{N} \\ \mathbf{P}_1^\perp &= \mathbf{I} - \mathbf{H}_1 (\mathbf{H}_1^T \mathbf{H}_1)^{-1} \mathbf{H}_1^T \\ &= \mathbf{I} - \frac{1}{N} \mathbf{1} \mathbf{1}^T \\ \mathbf{P}_1^\perp \mathbf{x} &= \mathbf{x} - \frac{1}{N} \mathbf{1} \mathbf{1}^T \mathbf{x} \\ &= \mathbf{x} - \bar{x} \mathbf{1} \\ \mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{x} &= \mathbf{h}_2^T \mathbf{x} - \bar{x} \mathbf{h}_2^T \mathbf{1} \\ &= \sum_{n=0}^{N-1} nx[n] - \bar{x} \sum_{n=0}^{N-1} n \\ \mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{h}_2 &= \mathbf{h}_2^T \mathbf{h}_2 - \frac{1}{N} (\mathbf{h}_2^T \mathbf{1})^2 \\ &= \sum_{n=0}^{N-1} n^2 - \frac{1}{N} \left(\sum_{n=0}^{N-1} n \right)^2. \end{aligned}$$

Upon substitution we have

$$\begin{aligned}\hat{\theta}_2 &= \begin{bmatrix} \bar{x} - \frac{\frac{1}{N} \sum_{n=0}^{N-1} n \left[\sum_{n=0}^{N-1} nx[n] - \bar{x} \sum_{n=0}^{N-1} n \right]}{\sum_{n=0}^{N-1} n^2 - \frac{1}{N} \left(\sum_{n=0}^{N-1} n \right)^2} \\ \frac{\sum_{n=0}^{N-1} nx[n] - \bar{x} \sum_{n=0}^{N-1} n}{\sum_{n=0}^{N-1} n^2 - \frac{1}{N} \left(\sum_{n=0}^{N-1} n \right)^2} \end{bmatrix} \\ &= \begin{bmatrix} \bar{x} - \frac{1}{N} \sum_{n=0}^{N-1} n \hat{B}_2 \\ \hat{B}_2 \end{bmatrix}.\end{aligned}$$

Simplifying produces

$$\begin{aligned}\hat{B}_2 &= \frac{\sum_{n=0}^{N-1} nx[n] - \bar{x} \sum_{n=0}^{N-1} n}{\sum_{n=0}^{N-1} n^2 - \frac{1}{N} \left(\sum_{n=0}^{N-1} n \right)^2} \\ &= \frac{\sum_{n=0}^{N-1} nx[n] - \frac{N(N-1)}{2} \bar{x}}{\frac{N(N^2-1)}{12}} \\ &= -\frac{6}{N(N+1)} \sum_{n=0}^{N-1} x[n] + \frac{12}{N(N^2-1)} \sum_{n=0}^{N-1} nx[n]\end{aligned}$$

and

$$\begin{aligned}\hat{A}_2 &= \bar{x} - \frac{1}{N} \sum_{n=0}^{N-1} n \hat{B}_2 \\ &= \bar{x} - \frac{N-1}{2} \hat{B}_2 \\ &= \frac{2(2N-1)}{N(N+1)} \sum_{n=0}^{N-1} x[n] - \frac{6}{N(N+1)} \sum_{n=0}^{N-1} nx[n].\end{aligned}$$

These are seen to agree with (8.24). The minimum LS error is, from (8.31),

$$\begin{aligned}
 J_{\min_2} &= J_{\min_1} - \frac{(\mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{x})^2}{\mathbf{h}_2^T \mathbf{P}_1^\perp \mathbf{h}_2} \\
 &= J_{\min_1} - \frac{\left(\sum_{n=0}^{N-1} nx[n] - \bar{x} \sum_{n=0}^{N-1} n \right)^2}{\sum_{n=0}^{N-1} n^2 - \frac{1}{N} \left(\sum_{n=0}^{N-1} n \right)^2} \\
 &= J_{\min_1} - \frac{\left(\sum_{n=0}^{N-1} nx[n] - \frac{N}{2}(N-1)\bar{x} \right)^2}{N(N^2-1)/12}.
 \end{aligned}$$

◇

In solving for the LSE of the parameters of the second-order model we first solved for the LSE of the first-order model. In general, the recursive procedure solves successive LS problems until the desired order is attained. Hence, the recursive order-update solution not only determines the LSE for the desired model but also *determines the LSE for all lower-order models as well*. As discussed previously, this property is useful when the model order is not known a priori.

Several interesting observations can be made based on (8.28).

1. If the new column $\mathbf{h}_{k+1}$ is orthogonal to all the previous ones, then $\mathbf{H}_k^T \mathbf{h}_{k+1} = \mathbf{0}$ and from (8.28)

$$\hat{\boldsymbol{\theta}}_{k+1} = \begin{bmatrix} \hat{\boldsymbol{\theta}}_k \\ \frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix}$$

or the LSE for the first k components of $\hat{\boldsymbol{\theta}}_{k+1}$ remains the same. This generalizes the geometrical illustration in Figure 8.3a.

2. The term $\mathbf{P}_k^\perp \mathbf{x} = (\mathbf{I} - \mathbf{H}_k(\mathbf{H}_k^T \mathbf{H}_k)^{-1} \mathbf{H}_k^T) \mathbf{x} = \mathbf{x} - \mathbf{H}_k \hat{\boldsymbol{\theta}}_k = \boldsymbol{\epsilon}_k$ is the LS error vector or the *data residual* that cannot be modeled by the columns of $\mathbf{H}_k$. It represents the component of $\mathbf{x}$ orthogonal to the space spanned by the columns of $\mathbf{H}_k$ (see Problem 8.17). We can view $\mathbf{P}_k^\perp \mathbf{x}$, the residual, as the part of $\mathbf{x}$ yet to be modeled.
3. If $\mathbf{h}_{k+1}$ is nearly in the space spanned by the columns of $\mathbf{H}_k$, then, as shown in Figure 8.8, $\mathbf{P}_k^\perp \mathbf{h}_{k+1}$ will be small. Consequently,

$$\begin{aligned}
 \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1} &= \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{P}_k^\perp \mathbf{h}_{k+1} \\
 &= \|\mathbf{P}_k^\perp \mathbf{h}_{k+1}\|^2 \\
 &\approx 0
 \end{aligned}$$




Figure 8.8 Near collinearity of columns of observation matrix

and the recursive procedure will “blow up” (see (8.28)). In essence, since h_{k+1} nearly lies in the subspace spanned by the columns of H_k , the new observation matrix H_{k+1} will be nearly of rank k and hence $H_{k+1}^T H_{k+1}$ will be nearly singular. In practice, we could monitor the term $h_{k+1}^T P_k^\perp h_{k+1}$ and exclude from the recursion those column vectors that produce small values of this term.

4. The minimum LS error can also be written in a more suggestive fashion, indicating the contribution of the new parameter in reducing the error. From (8.22) and (8.31) we have

$$\begin{aligned}
 J_{\min_{k+1}} &= \mathbf{x}^T P_k^\perp \mathbf{x} - \frac{(\mathbf{h}_{k+1}^T P_k^\perp \mathbf{x})^2}{\mathbf{h}_{k+1}^T P_k^\perp \mathbf{h}_{k+1}} \\
 &= \mathbf{x}^T P_k^\perp \mathbf{x} \left[1 - \frac{(\mathbf{h}_{k+1}^T P_k^\perp \mathbf{x})^2}{\mathbf{x}^T P_k^\perp \mathbf{x} \mathbf{h}_{k+1}^T P_k^\perp \mathbf{h}_{k+1}} \right] \\
 &= J_{\min_k} (1 - r_{k+1}^2)
 \end{aligned} \tag{8.32}$$

where

$$r_{k+1}^2 = \frac{[(P_k^\perp \mathbf{h}_{k+1})^T (P_k^\perp \mathbf{x})]^2}{\|P_k^\perp \mathbf{h}_{k+1}\|^2 \|P_k^\perp \mathbf{x}\|^2}. \tag{8.33}$$

Letting $(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y}$ denote the inner product in R^N , we have

$$r_{k+1}^2 = \frac{(P_k^\perp \mathbf{h}_{k+1}, P_k^\perp \mathbf{x})^2}{\|P_k^\perp \mathbf{h}_{k+1}\|^2 \|P_k^\perp \mathbf{x}\|^2}$$

where r_{k+1}^2 is seen to be the square of a correlation coefficient and as such has the property

$$0 \leq r_{k+1}^2 \leq 1.$$

Intuitively, $\mathbf{P}_k^\perp \mathbf{x}$ represents the residual or error in modeling $\mathbf{x}$ based on k parameters, while $\mathbf{P}_k^\perp \mathbf{h}_{k+1}$ represents the new model information contributed by the $(k+1)$ st parameter (see Figure 8.8). For instance, if $\mathbf{P}_k^\perp \mathbf{x}$ and $\mathbf{P}_k^\perp \mathbf{h}_{k+1}$ are collinear, then $r_{k+1}^2 = 1$ and $J_{\min k+1} = 0$. This says that the part of $\mathbf{x}$ that could not be modeled by the columns of $\mathbf{H}_k$ can be perfectly modeled by $\mathbf{h}_{k+1}$. Assuming that $r_{k+1} \neq 0$, the expression of (8.32) also shows that the minimum LS error monotonically decreases with order as in Figure 8.6.

5. Recall that the LS signal estimate is

$$\hat{\mathbf{s}} = \mathbf{H}\hat{\boldsymbol{\theta}} = \mathbf{P}\mathbf{x}$$

where $\mathbf{P}$ is the projection matrix. It is shown in Appendix 8B that $\mathbf{P}$ may be recursively updated by using (8.28). The result is

$$\mathbf{P}_{k+1} = \mathbf{P}_k + \frac{(\mathbf{I} - \mathbf{P}_k)\mathbf{h}_{k+1}\mathbf{h}_{k+1}^T(\mathbf{I} - \mathbf{P}_k)}{\mathbf{h}_{k+1}^T(\mathbf{I} - \mathbf{P}_k)\mathbf{h}_{k+1}}. \quad (8.34)$$

This is termed the *recursive orthogonal projection matrix*. See Problem 8.18 for its utility in determining the recursive formula for the minimum LS error. If we define the unit length vector

$$\begin{aligned} \mathbf{u}_{k+1} &= \frac{(\mathbf{I} - \mathbf{P}_k)\mathbf{h}_{k+1}}{\|(\mathbf{I} - \mathbf{P}_k)\mathbf{h}_{k+1}\|} \\ &= \frac{\mathbf{P}_k^\perp \mathbf{h}_{k+1}}{\|\mathbf{P}_k^\perp \mathbf{h}_{k+1}\|} \end{aligned}$$

the recursive projection operator becomes

$$\mathbf{P}_{k+1} = \mathbf{P}_k + \mathbf{u}_{k+1}\mathbf{u}_{k+1}^T$$

where $\mathbf{u}_{k+1}$ points in the direction of the new information (see Figure 8.8), or by letting $\hat{\mathbf{s}}_k = \mathbf{P}_k \mathbf{x}$, we have

$$\begin{aligned} \hat{\mathbf{s}}_{k+1} &= \mathbf{P}_{k+1} \mathbf{x} \\ &= \mathbf{P}_k \mathbf{x} + (\mathbf{u}_{k+1}^T \mathbf{x}) \mathbf{u}_{k+1} \\ &= \hat{\mathbf{s}}_k + (\mathbf{u}_{k+1}^T \mathbf{x}) \mathbf{u}_{k+1}. \end{aligned}$$

8.7 Sequential Least Squares

In many signal processing applications the received data are obtained by sampling a continuous-time waveform. Such data are on-going in that as time progresses, more data become available. We have the option of either waiting for all the available data or, as we now describe, of processing the data sequentially in time. This yields a sequence of LSEs in time. Specifically, assume we have determined the LSE $\hat{\boldsymbol{\theta}}$ based on

$\{x[0], x[1], \dots, x[N-1]\}$. If we now observe $x[N]$, can we update $\hat{\theta}$ (in time) without having to resolve the linear equations of (8.10)? The answer is yes, and the procedure is termed *sequential least squares* to distinguish it from our original approach in which we processed all the data at once, termed the *batch* approach.

Consider Example 8.1 in which the DC signal level is to be estimated. We saw that the LSE is

$$\hat{A}[N-1] = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

where the argument of $\hat{A}$ denotes the index of the most recent data point observed. If we now observe the new data sample $x[N]$, then the LSE becomes

$$\hat{A}[N] = \frac{1}{N+1} \sum_{n=0}^N x[n].$$

In computing this new estimator we do not have to recompute the sum of the observations since

$$\begin{aligned} \hat{A}[N] &= \frac{1}{N+1} \left(\sum_{n=0}^{N-1} x[n] + x[N] \right) \\ &= \frac{N}{N+1} \hat{A}[N-1] + \frac{1}{N+1} x[N]. \end{aligned} \quad (8.35)$$

The new LSE is found by using the previous one and the new observation. The sequential approach also lends itself to an interesting interpretation. Rearranging (8.35), we have

$$\hat{A}[N] = \hat{A}[N-1] + \frac{1}{N+1} (x[N] - \hat{A}[N-1]). \quad (8.36)$$

The new estimate is equal to the old one plus a *correction* term. The correction term decreases with N , reflecting the fact that the estimate $\hat{A}[N-1]$ is based on many more data samples and therefore should be given more weight. Also, $x[N] - \hat{A}[N-1]$ can be thought of as the error in predicting $x[N]$ by the previous samples, which are summarized by $\hat{A}[N-1]$. If this error is zero, then no correction takes place for that update. Otherwise, the new estimate differs from the old one.

The minimum LS error may also be computed recursively. Based on data samples up to time $N-1$, the error is

$$J_{\min}[N-1] = \sum_{n=0}^{N-1} (x[n] - \hat{A}[N-1])^2$$

and thus using (8.36)

$$J_{\min}[N] = \sum_{n=0}^N (x[n] - \hat{A}[N])^2$$

$$\begin{aligned}
&= \sum_{n=0}^{N-1} \left[x[n] - \hat{A}[N-1] - \frac{1}{N+1} (x[N] - \hat{A}[N-1]) \right]^2 + (x[N] - \hat{A}[N])^2 \\
&= J_{\min}[N-1] - \frac{2}{N+1} \sum_{n=0}^{N-1} (x[n] - \hat{A}[N-1])(x[N] - \hat{A}[N-1]) \\
&\quad + \frac{N}{(N+1)^2} (x[N] - \hat{A}[N-1])^2 + (x[N] - \hat{A}[N])^2.
\end{aligned}$$

Noting that the middle term on the right-hand side is zero, we have after some simplification

$$J_{\min}[N] = J_{\min}[N-1] + \frac{N}{N+1} (x[N] - \hat{A}[N-1])^2.$$

The apparent paradoxical behavior of an increase in the minimum LS error is readily explained if we note that with each new data point, the number of squared error terms increases. Thus, we need to fit more points with the same number of parameters.

A more interesting example of a sequential LS approach arises in the *weighted* LS problem. For the present example, if the weighting matrix $\mathbf{W}$ is diagonal, with $[\mathbf{W}]_{ii} = 1/\sigma_i^2$, then the weighted LSE is, from (8.15),

$$\hat{A}[N-1] = \frac{\sum_{n=0}^{N-1} \frac{x[n]}{\sigma_n^2}}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2}}.$$

To find the sequential version

$$\begin{aligned}
\hat{A}[N] &= \frac{\sum_{n=0}^N \frac{x[n]}{\sigma_n^2}}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} \\
&= \frac{\sum_{n=0}^{N-1} \frac{x[n]}{\sigma_n^2} + \frac{x[N]}{\sigma_N^2}}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} \\
&= \frac{\left(\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2} \right) \hat{A}[N-1] + \frac{x[N]}{\sigma_N^2}}{\sum_{n=0}^N \frac{1}{\sigma_n^2}}
\end{aligned}$$

$$= \hat{A}[N-1] - \frac{\frac{1}{\sigma_N^2} \hat{A}[N-1]}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} + \frac{\frac{x[N]}{\sigma_N^2}}{\sum_{n=0}^N \frac{1}{\sigma_n^2}}$$

or finally

$$\hat{A}[N] = \hat{A}[N-1] + \frac{1}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} (\sigma_N^2 (x[N] - \hat{A}[N-1])). \quad (8.37)$$

As expected, if $\sigma_n^2 = \sigma^2$ for all n , we have our previous result. The gain factor that multiplies the correction term now depends on our confidence in the new data sample. If the new sample is noisy or $\sigma_N^2 \rightarrow \infty$, we do not correct the previous LSE. On the other hand, if the new sample is noise-free or $\sigma_N^2 \rightarrow 0$, then $\hat{A}[N] \rightarrow x[N]$. We discard all the previous samples. The gain factor then represents our confidence in the new data sample relative to the previous ones. We can make a further interpretation of our results if indeed $x[n] = A + w[n]$, where $w[n]$ is zero mean uncorrelated noise with variance σ_n^2 . Then, we know from Chapter 6 (see Example 6.2) that the LSE is actually the BLUE, and therefore,

$$\text{var}(\hat{A}[N-1]) = \frac{1}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2}}$$

and the gain factor for the N th correction is (see (8.37))

$$\begin{aligned} K[N] &= \frac{\frac{1}{\sigma_N^2}}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} \\ &= \frac{\frac{1}{\sigma_N^2}}{\frac{1}{\sigma_N^2} + \frac{1}{\text{var}(\hat{A}[N-1])}} \\ &= \frac{\text{var}(\hat{A}[N-1])}{\text{var}(\hat{A}[N-1]) + \sigma_N^2}. \end{aligned} \quad (8.38)$$

Since $0 \leq K[N] \leq 1$, the correction is large if $K[N]$ is large or $\text{var}(\hat{A}[N-1])$ is large. Likewise, if the variance of the previous estimator is small, then so is the correction. Further expressions may be developed for determining the gain recursively, since $K[N]$

depends on $\text{var}(\hat{A}[N-1])$. We can express the latter as

$$\begin{aligned}
 \text{var}(\hat{A}[N]) &= \frac{1}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} \\
 &= \frac{1}{\sum_{n=0}^{N-1} \frac{1}{\sigma_n^2} + \frac{1}{\sigma_N^2}} \\
 &= \frac{1}{\frac{1}{\text{var}(\hat{A}[N-1])} + \frac{1}{\sigma_N^2}} \\
 &= \frac{\text{var}(\hat{A}[N-1])\sigma_N^2}{\text{var}(\hat{A}[N-1]) + \sigma_N^2} \\
 &= \left(1 - \frac{\text{var}(\hat{A}[N-1])}{\text{var}(\hat{A}[N-1]) + \sigma_N^2}\right) \text{var}(\hat{A}[N-1])
 \end{aligned}$$

or finally

$$\text{var}(\hat{A}[N]) = (1 - K[N])\text{var}(\hat{A}[N-1]). \quad (8.39)$$

To find $K[N]$ recursively we can use (8.38) and (8.39) as summarized below. In the process of recursively finding the gain, the variance of the LSE is also found recursively. Summarizing our results, we have

Estimator Update:

$$\hat{A}[N] = \hat{A}[N-1] + K[N](x[N] - \hat{A}[N-1]) \quad (8.40)$$

where

$$K[N] = \frac{\text{var}(\hat{A}[N-1])}{\text{var}(\hat{A}[N-1]) + \sigma_N^2}. \quad (8.41)$$

Variance Update:

$$\text{var}(\hat{A}[N]) = (1 - K[N])\text{var}(\hat{A}[N-1]). \quad (8.42)$$

To start the recursion we use

$$\begin{aligned}
 \hat{A}[0] &= x[0] \\
 \text{var}(\hat{A}[0]) &= \sigma_0^2.
 \end{aligned}$$

Then, $K[1]$ is found from (8.41), and $\hat{A}[1]$ from (8.40). Next, $\text{var}(\hat{A}[1])$ is determined from (8.42). Continuing in this manner, we find $K[2]$, $\hat{A}[2]$, $\text{var}(\hat{A}[2])$, etc. An example of the sequential LS approach is shown in Figure 8.9 for $A = 10$, $\sigma_n^2 = 1$ using a Monte Carlo computer simulation. The variance and gain sequences have been computed




Figure 8.9 Sequential least squares for DC level in white Gaussian noise

recursively from (8.42) and (8.41), respectively. Note that they decrease to zero since

$$\text{var}(\hat{A}[N]) = \frac{1}{\sum_{n=0}^N \frac{1}{\sigma_n^2}} = \frac{1}{N+1}$$

and from (8.41)

$$K[N] = \frac{\frac{1}{N}}{\frac{1}{N} + 1} = \frac{1}{N+1}.$$

Also, as seen in Figure 8.9c, the estimate appears to be converging to the true value of $A = 10$. This is in agreement with the variance approaching zero. Finally, it can be shown (see Problem 8.19) that the minimum LS error can be computed recursively as

$$J_{\min}[N] = J_{\min}[N-1] + \frac{(x[N] - \hat{A}[N-1])^2}{\text{var}(\hat{A}[N-1]) + \sigma_N^2}. \quad (8.43)$$

We now generalize our results to obtain the sequential LSE for a vector parameter. The derivation is given in Appendix 8C. Consider the minimization of the weighted LS error criterion J with $\mathbf{W} = \mathbf{C}^{-1}$, where $\mathbf{C}$ denotes the covariance matrix of the zero mean noise or

$$J = (\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}).$$

This is implicitly the same assumptions as in the BLUE. We know from (8.16) that the weighted LSE is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$$

and since $\mathbf{C}$ is the covariance matrix of the noise, we have from (6.17)

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}$$

where $\mathbf{C}_{\hat{\boldsymbol{\theta}}}$ is the covariance matrix of $\hat{\boldsymbol{\theta}}$. If $\mathbf{C}$ is diagonal or the noise is uncorrelated, then $\hat{\boldsymbol{\theta}}$ may be computed sequentially in time, but otherwise not. Assuming this condition to hold, let

$$\begin{aligned} \mathbf{C}[n] &= \text{diag}(\sigma_0^2, \sigma_1^2, \dots, \sigma_n^2) \\ \mathbf{H}[n] &= \begin{bmatrix} \mathbf{H}[n-1] \\ \mathbf{h}^T[n] \end{bmatrix} = \begin{bmatrix} n \times p \\ 1 \times p \end{bmatrix} \\ \mathbf{x}[n] &= [x[0] \ x[1] \ \dots \ x[n]]^T \end{aligned}$$

and denote the weighted LSE of $\boldsymbol{\theta}$ based on $\mathbf{x}[n]$ or the $(n+1)$ data samples as $\hat{\boldsymbol{\theta}}[n]$. Then, the batch estimator is

$$\hat{\boldsymbol{\theta}}[n] = (\mathbf{H}^T[n] \mathbf{C}^{-1}[n] \mathbf{H}[n])^{-1} \mathbf{H}^T[n] \mathbf{C}^{-1}[n] \mathbf{x}[n] \quad (8.44)$$




Figure 8.10 Sequential least squares estimator

with covariance matrix

$$\mathbf{C}_{\hat{\theta}} = \Sigma[n] = (\mathbf{H}^T[n] \mathbf{C}^{-1}[n] \mathbf{H}[n])^{-1}. \quad (8.45)$$

It should be kept in mind that $\mathbf{C}[n]$ is the covariance matrix of the noise, while $\Sigma[n]$ is the covariance matrix of the LSE. The sequential LSE becomes

Estimator Update:

$$\hat{\theta}[n] = \hat{\theta}[n-1] + \mathbf{K}[n] (x[n] - \mathbf{h}^T[n] \hat{\theta}[n-1]) \quad (8.46)$$

where

$$\mathbf{K}[n] = \frac{\Sigma[n-1] \mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n]}. \quad (8.47)$$

Covariance Update:

$$\Sigma[n] = (\mathbf{I} - \mathbf{K}[n] \mathbf{h}^T[n]) \Sigma[n-1]. \quad (8.48)$$

The gain factor $\mathbf{K}[n]$ is a $p \times 1$ vector, and the covariance matrix $\Sigma[n]$ has dimension $p \times p$. It is of great interest that no matrix inversions are required. The estimator update is summarized in Figure 8.10, where the thick arrows indicate vector processing. To start the recursion we need to specify initial values for $\hat{\theta}[n-1]$ and $\Sigma[n-1]$, so that $\mathbf{K}[n]$ can be determined from (8.47) and then $\hat{\theta}[n]$ from (8.46). In deriving (8.46)–(8.48) it was assumed that $\hat{\theta}[n-1]$ and $\Sigma[n-1]$ were available or that $\mathbf{H}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{H}[n-1]$ was invertible as per (8.44) and (8.45). For this to be invertible $\mathbf{H}[n-1]$ must have rank greater than or equal to p . Since $\mathbf{H}[n-1]$ is an $n \times p$ matrix, we must have $n \geq p$ (assuming all its columns are linearly independent). Hence, the sequential LS procedure normally determines $\hat{\theta}[p-1]$ and $\Sigma[p-1]$ using the batch estimator (8.44) and (8.45), and then employs the sequential equations (8.46)–(8.48) for $n \geq p$. A second method of initializing the recursion is to assign values for $\hat{\theta}[-1]$ and $\Sigma[-1]$. Then, the sequential

LS estimator is computed for $n \geq 0$. This has the effect of biasing the estimator toward $\hat{\theta}[-1]$. Typically, to minimize the biasing effect we choose $\Sigma[-1]$ to be large (little confidence in $\hat{\theta}[-1]$) or $\Sigma[-1] = \alpha \mathbf{I}$, where α is large, and also $\hat{\theta}[-1] = \mathbf{0}$. The LSE for $n \geq p$ will be the same as when the batch estimator is used for initialization if $\alpha \rightarrow \infty$ (see Problem 8.23). In the next example we show how to set up the equations. In Example 8.13 we apply this procedure to a signal processing problem.

Example 8.7 - Fourier Analysis

We now continue Example 8.4 in which the signal model is

$$s[n] = a \cos 2\pi f_0 n + b \sin 2\pi f_0 n \quad n \geq 0$$

and $\theta = [a \ b]^T$ is to be estimated by a sequential LSE. We furthermore assume that the noise is uncorrelated ($\mathbf{C}[n]$ must be diagonal for sequential LS to apply) and has equal variance σ^2 for each data sample. Since there are two parameters to be estimated, we need at least two observations or $x[0], x[1]$ to initialize the procedure using a batch initialization approach. We compute our first LSE, using (8.44), as

$$\begin{aligned} \hat{\theta}[1] &= \left(\mathbf{H}^T[1] \left(\frac{1}{\sigma^2} \mathbf{I} \right) \mathbf{H}[1] \right)^{-1} \mathbf{H}^T[1] \left(\frac{1}{\sigma^2} \mathbf{I} \right) \mathbf{x}[1] \\ &= (\mathbf{H}^T[1] \mathbf{H}[1])^{-1} \mathbf{H}^T[1] \mathbf{x}[1] \end{aligned}$$

where

$$\begin{aligned} \mathbf{H}[1] &= \begin{bmatrix} 1 & 0 \\ \cos 2\pi f_0 & \sin 2\pi f_0 \end{bmatrix} \\ \mathbf{x}[1] &= \begin{bmatrix} x[0] \\ x[1] \end{bmatrix}. \end{aligned}$$

($\mathbf{H}[1]$ is the 2×2 partition of $\mathbf{H}$ given in (8.18).) The initial covariance matrix is, from (8.45),

$$\begin{aligned} \Sigma[1] &= \left[\mathbf{H}^T[1] \left(\frac{1}{\sigma^2} \mathbf{I} \right) \mathbf{H}[1] \right]^{-1} \\ &= \sigma^2 (\mathbf{H}^T[1] \mathbf{H}[1])^{-1}. \end{aligned}$$

Next we determine $\hat{\theta}[2]$. To do so we first compute the 2×1 gain vector from (8.47) as

$$\mathbf{K}[2] = \frac{\Sigma[1] \mathbf{h}[2]}{\sigma^2 + \mathbf{h}^T[2] \Sigma[1] \mathbf{h}[2]}$$

where $\mathbf{h}^T[2]$ is the new row of the $\mathbf{H}[2]$ matrix or

$$\mathbf{h}^T[2] = [\cos 4\pi f_0 \quad \sin 4\pi f_0].$$

Once the gain vector has been found, $\hat{\theta}[2]$ is determined from (8.46) as

$$\hat{\theta}[2] = \hat{\theta}[1] + \mathbf{K}[2](x[2] - \mathbf{h}^T[2]\hat{\theta}[1]).$$

Finally, the 2×2 LSE covariance matrix is updated as per (8.48) for use in the next computation of the gain vector or

$$\Sigma[2] = (\mathbf{I} - \mathbf{K}[2]\mathbf{h}^T[2])\Sigma[1].$$

It should be clear that, in general, computer evaluation of $\hat{\theta}[n]$ is necessary. Also, except for the initialization procedure, no matrix inversion is required. Alternatively, we could have avoided even the matrix inversion of the initialization by assuming $\hat{\theta}[-1] = \mathbf{0}$ and $\Sigma[-1] = \alpha\mathbf{I}$ with α large. The recursion then would have begun at $n = 0$, and for $n \geq 2$ we would have the same result as before for large enough α . $\diamond$

Finally, we remark that if the minimum LS error is desired, it too can be found sequentially as (see Appendix 8C)

$$J_{\min}[n] = J_{\min}[n-1] + \frac{(x[n] - \mathbf{h}^T[n]\hat{\theta}[n-1])^2}{\sigma_n^2 + \mathbf{h}^T[n]\Sigma[n-1]\mathbf{h}[n]}. \quad (8.49)$$

8.8 Constrained Least Squares

At times we are confronted with LS problems whose unknown parameters must be constrained. Such would be the case if we wished to estimate the amplitudes of several signals but knew a priori that some of the amplitudes were equal. Then, the total number of parameters to be estimated should be reduced to take advantage of the a priori knowledge. For this example the parameters are linearly related, leading to a linear least squares problem with *linear constraints*. This problem is easily solved, as we now show.

Assume that the parameter θ is subject to $r < p$ linear constraints. The constraints must be independent, ruling out the possibility of redundant constraints such as $\theta_1 + \theta_2 = 0$, $2\theta_1 + 2\theta_2 = 0$, where θ_i is the i th element of θ . We summarize the constraints as

$$\mathbf{A}\theta = \mathbf{b} \quad (8.50)$$

where $\mathbf{A}$ is a known $r \times p$ matrix and $\mathbf{b}$ is a known $r \times 1$ vector. If, for instance, $p = 2$ and one parameter is known to be the negative of the other, then the constraint would be $\theta_1 + \theta_2 = 0$. We would then have $\mathbf{A} = [1 \ 1]$ and $\mathbf{b} = 0$. It is always assumed that the matrix $\mathbf{A}$ is full rank (equal to r), which is necessary for the constraints to be independent. It should be realized that in the constrained LS problem there are actually only $(p - r)$ independent parameters.

To find the LSE subject to the linear constraints we use the technique of Lagrangian multipliers. We determine $\hat{\theta}_c$ (c denotes the constrained LSE) by minimizing the Lagrangian

$$J_c = (\mathbf{x} - \mathbf{H}\theta)^T(\mathbf{x} - \mathbf{H}\theta) + \lambda^T(\mathbf{A}\theta - \mathbf{b})$$

where λ is a $r \times 1$ vector of Lagrangian multipliers. Expanding this expression, we have

$$J_c = \mathbf{x}^T \mathbf{x} - 2\boldsymbol{\theta}^T \mathbf{H}^T \mathbf{x} + \boldsymbol{\theta}^T \mathbf{H}^T \mathbf{H} \boldsymbol{\theta} + \boldsymbol{\lambda}^T \mathbf{A} \boldsymbol{\theta} - \boldsymbol{\lambda}^T \mathbf{b}.$$

Taking the gradient with respect to $\boldsymbol{\theta}$ and using (4.3) yields

$$\frac{\partial J_c}{\partial \boldsymbol{\theta}} = -2\mathbf{H}^T \mathbf{x} + 2\mathbf{H}^T \mathbf{H} \boldsymbol{\theta} + \mathbf{A}^T \boldsymbol{\lambda}$$

and setting it equal to zero produces

$$\begin{aligned} \hat{\boldsymbol{\theta}}_c &= (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} - \frac{1}{2} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T \boldsymbol{\lambda} \\ &= \hat{\boldsymbol{\theta}} - (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T \frac{\boldsymbol{\lambda}}{2} \end{aligned} \quad (8.51)$$

where $\hat{\boldsymbol{\theta}}$ is the *unconstrained* LSE and $\boldsymbol{\lambda}$ is yet to be determined. To find $\boldsymbol{\lambda}$ we impose the constraint of (8.50), so that

$$\mathbf{A} \boldsymbol{\theta}_c = \mathbf{A} \hat{\boldsymbol{\theta}} - \mathbf{A} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T \frac{\boldsymbol{\lambda}}{2} = \mathbf{b}$$

and hence

$$\frac{\boldsymbol{\lambda}}{2} = [\mathbf{A} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T]^{-1} (\mathbf{A} \hat{\boldsymbol{\theta}} - \mathbf{b}).$$

Substituting into (8.51) produces the solution

$$\hat{\boldsymbol{\theta}}_c = \hat{\boldsymbol{\theta}} - (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T [\mathbf{A} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{A}^T]^{-1} (\mathbf{A} \hat{\boldsymbol{\theta}} - \mathbf{b}) \quad (8.52)$$

where $\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$. The constrained LSE is a corrected version of the unconstrained LSE. If it happens that the constraint is fortuitously satisfied by $\hat{\boldsymbol{\theta}}$ or $\mathbf{A} \hat{\boldsymbol{\theta}} = \mathbf{b}$, then according to (8.52) the estimators are identical. Such is usually not the case, however. We consider a simple example.

Example 8.8 - Constrained Signal

If the signal model is

$$s[n] = \begin{cases} \theta_1 & n = 0 \\ \theta_2 & n = 1 \\ 0 & n = 2 \end{cases}$$

and we observe $\{x[0], x[1], x[2]\}$, then the observation matrix is

$$\mathbf{H} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}.$$




Figure 8.11 Comparison of least squares and constrained least squares

Observe that the signal vector

$$\mathbf{s} = \mathbf{H}\boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \theta_2 \\ 0 \end{bmatrix}$$

must lie in the plane shown in Figure 8.11a. The unconstrained LSE is

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} = \begin{bmatrix} x[0] \\ x[1] \end{bmatrix}$$

and the signal estimate is

$$\hat{\mathbf{s}} = \mathbf{H}\hat{\boldsymbol{\theta}} = \begin{bmatrix} x[0] \\ x[1] \\ 0 \end{bmatrix}.$$

As shown in Figure 8.11a, this is intuitively reasonable. Now assume that we know a priori that $\theta_1 = \theta_2$. In terms of (8.50) we have

$$\begin{bmatrix} 1 & -1 \end{bmatrix} \boldsymbol{\theta} = 0$$

so that $\mathbf{A} = [1 \ -1]$ and $\mathbf{b} = 0$. Noting that $\mathbf{H}^T \mathbf{H} = \mathbf{I}$, we have from (8.52)

$$\begin{aligned} \hat{\boldsymbol{\theta}}_c &= \hat{\boldsymbol{\theta}} - \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1} \mathbf{A} \hat{\boldsymbol{\theta}} \\ &= [\mathbf{I} - \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1} \mathbf{A}] \hat{\boldsymbol{\theta}}. \end{aligned}$$

After some simple algebra this becomes

$$\hat{\boldsymbol{\theta}}_c = \begin{bmatrix} \frac{1}{2}(x[0] + x[1]) \\ \frac{1}{2}(x[0] + x[1]) \end{bmatrix}$$

and the constrained signal estimate becomes

$$\hat{\mathbf{s}}_c = \mathbf{H}\hat{\boldsymbol{\theta}}_c = \begin{bmatrix} \frac{1}{2}(x[0] + x[1]) \\ \frac{1}{2}(x[0] + x[1]) \\ 0 \end{bmatrix}$$

as illustrated in Figure 8.11b. Since $\theta_1 = \theta_2$, we just average the two observations, which is again intuitively reasonable. In this simple problem we could have just as easily incorporated our parameter constraints into the signal model directly to yield

$$s[n] = \begin{cases} \theta & n = 0 \\ \theta & n = 1 \\ 0 & n = 2 \end{cases}$$

and estimated θ . This would have produced the same result, as it must. This new model is sometimes referred to as the *reduced* model [Graybill 1976] for obvious reasons. It is worthwhile to view the constrained LS problem geometrically. Again referring to Figure 8.11b, if $\theta_1 = \theta_2$, then the signal model is

$$\mathbf{s} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \theta$$

and $\mathbf{s}$ must lie in the constraint subspace shown. The constrained signal estimate $\hat{\mathbf{s}}_c$ may be viewed as the projection of the unconstrained signal estimate $\hat{\mathbf{s}}$ onto the constrained subspace. This accounts for the correction term of (8.52). In fact, (8.52) may be obtained geometrically using projection theory. See also Problem 8.24. $\diamond$

8.9 Nonlinear Least Squares

In Section 8.3 we introduced the nonlinear LS problem. We now investigate this in more detail. Recall that the LS procedure estimates model parameters $\boldsymbol{\theta}$ by minimizing the LS error criterion

$$J = (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta}))^T (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta}))$$

where $\mathbf{s}(\boldsymbol{\theta})$ is the signal model for $\mathbf{x}$, with its dependence on $\boldsymbol{\theta}$ explicitly shown. (Note that if $\mathbf{x} - \mathbf{s}(\boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, the LSE is also the MLE.) In the linear LS problem the signal takes on the special form $\mathbf{s}(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\theta}$, which leads to the simple linear LSE. In general, $\mathbf{s}(\boldsymbol{\theta})$ cannot be expressed in this manner but is an N -dimensional nonlinear function of $\boldsymbol{\theta}$. In such a case the minimization of J becomes much more difficult, if not impossible. This type of nonlinear LS problem is often termed a *nonlinear regression problem* in statistics [Bard 1974, Seber and Wild 1989], and much theoretical work on it can be found. Practically, the determination of the nonlinear LSE must be based on

iterative approaches and so suffers from the same limitations discussed in Chapter 7 for MLE determination by numerical methods. The preferable method of a grid search is practical only if the dimensionality of θ is small, perhaps $p \leq 5$.

Before discussing general methods for determining nonlinear LSEs we first describe two methods that can reduce the complexity of the problem. They are

1. transformation of parameters
2. separability of parameters.

In the first case we seek a one-to-one transformation of θ that produces a *linear* signal model in the new space. To do so we let

$$\alpha = \mathbf{g}(\theta)$$

where $\mathbf{g}$ is a p -dimensional function of θ whose inverse exists. If a $\mathbf{g}$ can be found so that

$$\mathbf{s}(\theta(\alpha)) = \mathbf{s}(\mathbf{g}^{-1}(\alpha)) = \mathbf{H}\alpha$$

then the signal will be linear in α . We can then easily find the *linear* LSE of α and thus the *nonlinear* LSE of θ by

$$\hat{\theta} = \mathbf{g}^{-1}(\hat{\alpha})$$

where

$$\hat{\alpha} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

This approach relies on the property that the minimization can be carried out in any transformed space that is obtained by a one-to-one mapping and then converted back to the original space (see Problem 8.26). The determination of the transformation $\mathbf{g}$, if it exists, is usually quite difficult. Suffice it to say, only a few nonlinear LS problems may be solved in this manner.

Example 8.9 - Sinusoidal Parameter Estimation

For a sinusoidal signal model

$$s[n] = A \cos(2\pi f_0 n + \phi) \quad n = 0, 1, \dots, N-1$$

it is desired to estimate the amplitude A , where $A > 0$, and phase ϕ . The frequency f_0 is assumed known. The LSE is obtained by minimizing

$$J = \sum_{n=0}^{N-1} (x[n] - A \cos(2\pi f_0 n + \phi))^2$$

over A and ϕ , a nonlinear LS problem. However, because

$$A \cos(2\pi f_0 n + \phi) = A \cos \phi \cos 2\pi f_0 n - A \sin \phi \sin 2\pi f_0 n$$

if we let

$$\begin{aligned}\alpha_1 &= A \cos \phi \\ \alpha_2 &= -A \sin \phi,\end{aligned}$$

then the signal model becomes

$$s[n] = \alpha_1 \cos 2\pi f_0 n + \alpha_2 \sin 2\pi f_0 n.$$

In matrix form this is

$$\mathbf{s} = \mathbf{H}\boldsymbol{\alpha}$$

where

$$\mathbf{H} = \begin{bmatrix} 1 & 0 \\ \cos 2\pi f_0 & \sin 2\pi f_0 \\ \vdots & \vdots \\ \cos 2\pi f_0(N-1) & \sin 2\pi f_0(N-1) \end{bmatrix},$$

which is now linear in the new parameters. The LSE of $\boldsymbol{\alpha}$ is

$$\hat{\boldsymbol{\alpha}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$$

and to find $\hat{\boldsymbol{\theta}}$ we must find the inverse transformation $\mathbf{g}^{-1}(\boldsymbol{\alpha})$. This is

$$\begin{aligned}A &= \sqrt{\alpha_1^2 + \alpha_2^2} \\ \phi &= \arctan\left(\frac{-\alpha_2}{\alpha_1}\right)\end{aligned}$$

so that the nonlinear LSE for this problem is given by

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \begin{bmatrix} \hat{A} \\ \hat{\phi} \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{\hat{\alpha}_1^2 + \hat{\alpha}_2^2} \\ \arctan\left(\frac{-\hat{\alpha}_2}{\hat{\alpha}_1}\right) \end{bmatrix}\end{aligned}$$

where $\hat{\boldsymbol{\alpha}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$. The reader may wish to refer back to Example 7.16 in which this approach was used to find the MLE of frequency. $\diamond$

A second type of nonlinear LS problem that is less complex than the general one exhibits the separability property. Although the signal model is nonlinear, it may be linear in some of the parameters, as illustrated in Example 8.3. In general, a separable signal model has the form

$$\mathbf{s} = \mathbf{H}(\boldsymbol{\alpha})\boldsymbol{\beta}$$

where

$$\theta = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} (p-q) \times 1 \\ q \times 1 \end{bmatrix}$$

and $\mathbf{H}(\alpha)$ is an $N \times q$ matrix dependent on α . This model is linear in β but nonlinear in α . As a result, the LS error may be minimized with respect to β and thus reduced to a function of α only. Since

$$J(\alpha, \beta) = (\mathbf{x} - \mathbf{H}(\alpha)\beta)^T (\mathbf{x} - \mathbf{H}(\alpha)\beta)$$

the β that minimizes J for a given α is

$$\hat{\beta} = (\mathbf{H}^T(\alpha)\mathbf{H}(\alpha))^{-1} \mathbf{H}^T(\alpha)\mathbf{x} \quad (8.53)$$

and the resulting LS error is, from (8.22),

$$J(\alpha, \hat{\beta}) = \mathbf{x}^T \left[\mathbf{I} - \mathbf{H}(\alpha) (\mathbf{H}^T(\alpha)\mathbf{H}(\alpha))^{-1} \mathbf{H}^T(\alpha) \right] \mathbf{x}.$$

The problem now reduces to a *maximization* of

$$\mathbf{x}^T \mathbf{H}(\alpha) (\mathbf{H}^T(\alpha)\mathbf{H}(\alpha))^{-1} \mathbf{H}^T(\alpha)\mathbf{x} \quad (8.54)$$

over α . If, for instance, $q = p - 1$, so that α is a scalar, then a grid search can possibly be used. This should be contrasted with the original minimization of a p -dimensional function. (See also Example 7.16.)

Example 8.10 - Damped Exponentials

Assume we have a signal model of the form

$$s[n] = A_1 r^n + A_2 r^{2n} + A_3 r^{3n}$$

where the unknown parameters are $\{A_1, A_2, A_3, r\}$. It is known that $0 < r < 1$. Then, the model is linear in the amplitudes $\beta = [A_1 \ A_2 \ A_3]^T$, and nonlinear in the damping factor $\alpha = r$. Using (8.54), the nonlinear LSE is obtained by maximizing

$$\mathbf{x}^T \mathbf{H}(r) (\mathbf{H}^T(r)\mathbf{H}(r))^{-1} \mathbf{H}^T(r)\mathbf{x}$$

over $0 < r < 1$, where

$$\mathbf{H}(r) = \begin{bmatrix} 1 & 1 & 1 \\ r & r^2 & r^3 \\ \vdots & \vdots & \vdots \\ r^{N-1} & r^{2(N-1)} & r^{3(N-1)} \end{bmatrix}.$$

Once $\hat{r}$ is found we have the LSE for the amplitudes

$$\hat{\beta} = (\mathbf{H}^T(\hat{r})\mathbf{H}(\hat{r}))^{-1} \mathbf{H}^T(\hat{r})\mathbf{x}.$$

This maximization is easily carried out on a digital computer. $\diamond$

It may also be possible to combine the approaches to reduce a nonseparable problem to a separable one by an appropriate transformation. When these approaches fail, we are forced to tackle the original nonlinear LS problem or to minimize

$$J = (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta}))^T (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta})).$$

Necessary conditions can be obtained by differentiation of J . They are

$$\frac{\partial J}{\partial \theta_j} = -2 \sum_{i=0}^{N-1} (x[i] - s[i]) \frac{\partial s[i]}{\partial \theta_j} = 0$$

for $j = 1, 2, \dots, p$. If we define an $N \times p$ Jacobian matrix as

$$\left[\frac{\partial \mathbf{s}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right]_{ij} = \frac{\partial s[i]}{\partial \theta_j} \quad \begin{array}{l} i = 0, 1, \dots, N-1 \\ j = 1, 2, \dots, p, \end{array}$$

then the necessary conditions become

$$\sum_{i=0}^{N-1} (x[i] - s[i]) \left[\frac{\partial \mathbf{s}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right]_{ij} = 0 \quad j = 1, 2, \dots, p$$

or in matrix form

$$\frac{\partial \mathbf{s}(\boldsymbol{\theta})^T}{\partial \boldsymbol{\theta}} (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta})) = 0 \quad (8.55)$$

which is a set of p simultaneous nonlinear equations. (If the signal is linear in the unknown parameters or $\mathbf{s}(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\theta}$ so that $\partial \mathbf{s}(\boldsymbol{\theta})/\partial \boldsymbol{\theta} = \mathbf{H}$, we have our usual LS equations.) They may be solved using the Newton-Raphson iteration described in Chapter 7. Letting

$$\mathbf{g}(\boldsymbol{\theta}) = \frac{\partial \mathbf{s}(\boldsymbol{\theta})^T}{\partial \boldsymbol{\theta}} (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta})) \quad (8.56)$$

the iteration becomes

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \left(\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^{-1} \mathbf{g}(\boldsymbol{\theta}) \bigg|_{\boldsymbol{\theta}=\boldsymbol{\theta}_k}. \quad (8.57)$$

To find the Jacobian of $\mathbf{g}$, which is actually the scaled Hessian of J , we need

$$\frac{\partial [\mathbf{g}(\boldsymbol{\theta})]_i}{\partial \theta_j} = \frac{\partial}{\partial \theta_j} \left[\sum_{n=0}^{N-1} (x[n] - s[n]) \frac{\partial s[n]}{\partial \theta_i} \right]$$

which follows from (8.56). Evaluating this, we have

$$\frac{\partial [\mathbf{g}(\boldsymbol{\theta})]_i}{\partial \theta_j} = \sum_{n=0}^{N-1} \left[(x[n] - s[n]) \frac{\partial^2 s[n]}{\partial \theta_i \partial \theta_j} - \frac{\partial s[n]}{\partial \theta_j} \frac{\partial s[n]}{\partial \theta_i} \right].$$

To put this in more succinct form let

$$[\mathbf{H}(\boldsymbol{\theta})]_{ij} = \left[\frac{\partial s(\boldsymbol{\theta})}{\partial \theta} \right]_{ij} = \frac{\partial s[i]}{\partial \theta_j} \quad (8.58)$$

for $i = 0, 1, \dots, N-1; j = 1, 2, \dots, p$. (Again if $s(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\theta}$, then $\mathbf{H}(\boldsymbol{\theta}) = \mathbf{H}$.) Also let

$$[\mathbf{G}_n(\boldsymbol{\theta})]_{ij} = \frac{\partial^2 s[n]}{\partial \theta_i \partial \theta_j} \quad (8.59)$$

for $i = 1, 2, \dots, p; j = 1, 2, \dots, p$. Then,

$$\begin{aligned} \frac{\partial [\mathbf{g}(\boldsymbol{\theta})]_i}{\partial \theta_j} &= \sum_{n=0}^{N-1} (x[n] - s[n]) [\mathbf{G}_n(\boldsymbol{\theta})]_{ij} - [\mathbf{H}(\boldsymbol{\theta})]_{nj} [\mathbf{H}(\boldsymbol{\theta})]_{ni} \\ &= \sum_{n=0}^{N-1} [\mathbf{G}_n(\boldsymbol{\theta})]_{ij} (x[n] - s[n]) - [\mathbf{H}^T(\boldsymbol{\theta})]_{in} [\mathbf{H}(\boldsymbol{\theta})]_{nj} \end{aligned}$$

and

$$\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \sum_{n=0}^{N-1} \mathbf{G}_n(\boldsymbol{\theta}) (x[n] - s[n]) - \mathbf{H}^T(\boldsymbol{\theta}) \mathbf{H}(\boldsymbol{\theta}). \quad (8.60)$$

In summary, we have upon using (8.57) with (8.56) and (8.60) the Newton-Raphson iteration

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \boldsymbol{\theta}_k + \left(\mathbf{H}^T(\boldsymbol{\theta}_k) \mathbf{H}(\boldsymbol{\theta}_k) - \sum_{n=0}^{N-1} \mathbf{G}_n(\boldsymbol{\theta}_k) (x[n] - s(\boldsymbol{\theta}_k)_n) \right)^{-1} \\ &\quad \cdot \mathbf{H}^T(\boldsymbol{\theta}_k) (\mathbf{x} - s(\boldsymbol{\theta}_k)) \end{aligned} \quad (8.61)$$

where $\mathbf{H}(\boldsymbol{\theta})$ is given by (8.58) and $\mathbf{G}_n(\boldsymbol{\theta})$ is given by (8.59). These are the first-order and second-order partials of the signal with respect to the unknown parameters. It is interesting to note that if $s(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\theta}$, then $\mathbf{G}_n(\boldsymbol{\theta}) = \mathbf{0}$ and $\mathbf{H}(\boldsymbol{\theta}) = \mathbf{H}$ and we have

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \boldsymbol{\theta}_k + (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}_k) \\ &= (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \end{aligned}$$

or convergence is attained in one step. This is, of course, due to the exact quadratic nature of the LS error criterion, which results in a linear $\mathbf{g}(\boldsymbol{\theta})$. It would be expected then that for signal models that are approximately linear, convergence will occur rapidly.

A second method for solving a nonlinear LS problem is to linearize the signal model about some nominal $\boldsymbol{\theta}$ and then apply the linear LS procedure. This differs from the Newton-Raphson method in which the *derivative* of J is linearized about the current iterate. To appreciate the difference consider a scalar parameter and let θ_0 be the nominal value of θ . Then, linearizing about θ_0 we have

$$s[n; \theta] \approx s[n; \theta_0] + \left. \frac{\partial s[n; \theta]}{\partial \theta} \right|_{\theta=\theta_0} (\theta - \theta_0)$$

where the dependence of $s[n]$ on θ is now shown. The LS error becomes

$$\begin{aligned} J &= \sum_{n=0}^{N-1} (x[n] - s[n; \theta])^2 \\ &\approx \sum_{n=0}^{N-1} \left(x[n] - s[n; \theta_0] + \frac{\partial s[n; \theta]}{\partial \theta} \bigg|_{\theta=\theta_0} \theta_0 - \frac{\partial s[n; \theta]}{\partial \theta} \bigg|_{\theta=\theta_0} \theta \right)^2 \\ &= (\mathbf{x} - \mathbf{s}(\theta_0) + \mathbf{H}(\theta_0)\theta_0 - \mathbf{H}(\theta_0)\theta)^T (\mathbf{x} - \mathbf{s}(\theta_0) + \mathbf{H}(\theta_0)\theta_0 - \mathbf{H}(\theta_0)\theta). \end{aligned}$$

Since $\mathbf{x} - \mathbf{s}(\theta_0) + \mathbf{H}(\theta_0)\theta_0$ is known, we have as the LSE

$$\begin{aligned} \hat{\theta} &= (\mathbf{H}^T(\theta_0)\mathbf{H}(\theta_0))^{-1} \mathbf{H}^T(\theta_0) (\mathbf{x} - \mathbf{s}(\theta_0) + \mathbf{H}(\theta_0)\theta_0) \\ &= \theta_0 + (\mathbf{H}^T(\theta_0)\mathbf{H}(\theta_0))^{-1} \mathbf{H}^T(\theta_0) (\mathbf{x} - \mathbf{s}(\theta_0)). \end{aligned}$$

If we now iterate the solution, it becomes

$$\theta_{k+1} = \theta_k + (\mathbf{H}^T(\theta_k)\mathbf{H}(\theta_k))^{-1} \mathbf{H}^T(\theta_k) (\mathbf{x} - \mathbf{s}(\theta_k))$$

which is identical to the Newton-Raphson iteration except for the omission of the second derivatives or for the presence of $\mathbf{G}_n$. The linearization method is termed the *Gauss-Newton method* and can easily be generalized to the vector parameter case as

$$\theta_{k+1} = \theta_k + (\mathbf{H}^T(\theta_k)\mathbf{H}(\theta_k))^{-1} \mathbf{H}^T(\theta_k) (\mathbf{x} - \mathbf{s}(\theta_k)) \quad (8.62)$$

where

$$[\mathbf{H}(\theta)]_{ij} = \frac{\partial s[i]}{\partial \theta_j}.$$

In Example 8.14 the Gauss method is illustrated. Both the Newton-Raphson and the Gauss methods can have convergence problems. It has been argued that neither method is reliable enough to use without safeguards. The interested reader should consult [Seber and Wild 1989] for additional implementation details.

8.10 Signal Processing Examples

We now describe some typical signal processing problems for which a LSE is used. In these applications the optimal MVU estimator is unavailable. The statistical characterization of the noise may not be known or, even if it is, the optimal MVU estimator cannot be found. For known noise statistics the asymptotically optimal MLE is generally too complicated to implement. Faced with these practical difficulties we resort to least squares.

Example 8.11 - Digital Filter Design

A common problem in digital signal processing is to design a digital filter whose frequency response closely matches a given frequency response specification [Oppenheim and Schaffer 1975]. Alternatively, we could attempt to match the impulse response, which is the approach examined in this example. A general infinite impulse response (IIR) filter has the system function

$$\begin{aligned}\mathcal{H}(z) &= \frac{\mathcal{B}(z)}{\mathcal{A}(z)} \\ &= \frac{b[0] + b[1]z^{-1} + \cdots + b[q]z^{-q}}{1 + a[1]z^{-1} + \cdots + a[p]z^{-p}}.\end{aligned}$$

If the desired frequency response is $H_d(f) = \mathcal{H}_d(\exp[j2\pi f])$, then the inverse Fourier transform is the desired impulse response or

$$h_d[n] = \mathcal{F}^{-1}\{H_d(f)\}.$$

The digital filter impulse response is given as a recursive difference equation, obtained by taking the inverse z transform of $\mathcal{H}(z)$ to produce

$$h[n] = \begin{cases} -\sum_{k=1}^p a[k]h[n-k] + \sum_{k=0}^q b[k]\delta[n-k] & n \geq 0 \\ 0 & n < 0. \end{cases}$$

A straightforward LS solution would seek to choose $\{a[k], b[k]\}$ to minimize

$$J = \sum_{n=0}^{N-1} (h_d[n] - h[n])^2$$

where N is some sufficiently large integer for which $h_d[n]$ is essentially zero. Unfortunately, this approach produces a nonlinear LS problem (see Problem 8.28 for the exact solution). (The reader should note that $h_d[n]$ plays the role of the “data,” and $h[n]$ that of the “signal” model. Also, the “noise” in the data is attributable to modeling error whose statistical characterization is unknown.) As an example, if

$$\mathcal{H}(z) = \frac{b[0]}{1 + a[1]z^{-1}},$$

then

$$h[n] = \begin{cases} b[0](-a[1])^n & n \geq 0 \\ 0 & n < 0 \end{cases}$$

and

$$J = \sum_{n=0}^{N-1} (h_d[n] - b[0](-a[1])^n)^2$$




Figure 8.12 Conversion of nonlinear least squares filter design to linear least squares

which is clearly very nonlinear in $a[1]$. In fact, it is the presence of $\mathcal{A}(z)$ that causes the LS error to be nonquadratic. To alleviate this problem we can filter both $h_d[n]$ and $h[n]$ by $\mathcal{A}(z)$, as shown in Figure 8.12. Then, we minimize the *filtered* LS error

$$J_f = \sum_{n=0}^{N-1} (h_{d_f}[n] - b[n])^2$$

where $h_{d_f}[n]$ is given as

$$h_{d_f}[n] = \sum_{k=0}^p a[k] h_d[n-k]$$

and $a[0] = 1$. The filtered LS error then becomes

$$J_f = \sum_{n=0}^{N-1} \left(\sum_{k=0}^p a[k] h_d[n-k] - b[n] \right)^2$$

which is now a quadratic function of the $a[k]$'s and $b[k]$'s. Alternatively,

$$J_f = \sum_{n=0}^{N-1} \left[h_d[n] - \left(- \sum_{k=1}^p a[k] h_d[n-k] + b[n] \right) \right]^2.$$

In minimizing this over the filter coefficients note that the $b[n]$'s appear only in the first $(q+1)$ terms since $b[n] = 0$ for $n > q$. As a result, we have upon letting $\mathbf{a} = [a[1] \ a[2] \ \dots \ a[p]]^T$ and $\mathbf{b} = [b[0] \ b[1] \ \dots \ b[q]]^T$,

$$\begin{aligned} J_f(\mathbf{a}, \mathbf{b}) &= \sum_{n=0}^q \left[h_d[n] - \left(- \sum_{k=1}^p a[k] h_d[n-k] + b[n] \right) \right]^2 \\ &\quad + \sum_{n=q+1}^{N-1} \left[h_d[n] - \left(- \sum_{k=1}^p a[k] h_d[n-k] \right) \right]^2. \end{aligned}$$

The first sum may be minimized (actually made to be zero) by setting

$$h_d[n] - \left(- \sum_{k=1}^p a[k] h_d[n-k] + b[n] \right) = 0 \quad n = 0, 1, \dots, q$$

or

$$b[n] = h_d[n] + \sum_{k=1}^p a[k] h_d[n-k].$$

In matrix form the LSE of the numerator coefficients is

$$\hat{\mathbf{b}} = \mathbf{h} + \mathbf{H}_0 \hat{\mathbf{a}}$$

where

$$\begin{aligned} \mathbf{h} &= [h_d[0] \ h_d[1] \ \dots \ h_d[q]]^T \\ \mathbf{H}_0 &= \begin{bmatrix} 0 & 0 & \dots & 0 \\ h_d[0] & 0 & \dots & 0 \\ h_d[1] & h_d[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h_d[q-1] & h_d[q-2] & \dots & h_d[q-p] \end{bmatrix}. \end{aligned}$$

The vector $\mathbf{h}$ has dimension $(q+1) \times 1$, while the matrix $\mathbf{H}_0$ has dimension $(q+1) \times p$. To find the LSE of the denominator coefficients or $\hat{\mathbf{a}}$ we must minimize

$$\begin{aligned} J_f(\mathbf{a}, \hat{\mathbf{b}}) &= \sum_{n=q+1}^{N-1} \left[h_d[n] - \left(- \sum_{k=1}^p a[k] h_d[n-k] \right) \right]^2 \\ &= (\mathbf{x} - \mathbf{H}\mathbf{a})^T (\mathbf{x} - \mathbf{H}\mathbf{a}) \end{aligned}$$

where $\theta = \mathbf{a}$ and

$$\mathbf{x} = [h_d[q+1] \quad h_d[q+2] \quad \dots \quad h_d[N-1]]^T \quad (N-1-q \times 1)$$

$$\mathbf{H} = - \begin{bmatrix} h_d[q] & h_d[q-1] & \dots & h_d[q-p+1] \\ h_d[q+1] & h_d[q] & \dots & h_d[q-p+2] \\ \vdots & \vdots & \ddots & \vdots \\ h_d[N-2] & h_d[N-3] & \dots & h_d[N-1-p] \end{bmatrix} \quad (N-1-q \times p).$$

The LSE of the denominator coefficients is

$$\hat{\mathbf{a}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

This method for designing digital filters is termed the *least squares Prony method* [Parks and Burrus 1987]. As an example, consider the design of a low-pass filter for which the desired frequency response is

$$H'_d(f) = \begin{cases} 1 & |f| < f_c \\ 0 & |f| > f_c \end{cases}$$


where f_c is the cutoff frequency. The corresponding impulse response is

$$h'_d[n] = \frac{\sin 2\pi f_c n}{\pi n} \quad -\infty < n < \infty$$

and as expected is not causal (due to the zero phase frequency response assumption). To ensure causality we delay the impulse response by n_0 samples and then set the resultant impulse response equal to zero for $n < 0$. Next, to approximate the desired impulse response using the LS Prony method we assume that N samples are available or

$$h_d[n] = \frac{\sin 2\pi f_c (n - n_0)}{\pi (n - n_0)} \quad n = 0, 1, \dots, N-1.$$

For a cutoff frequency of $f_c = 0.1$ and a delay of $n_0 = 25$ samples, the desired impulse response is shown in Figure 8.13a for $N = 50$. The corresponding desired magnitude frequency response is shown in Figure 8.13b. The effects of truncating the impulse response are manifested by the sidelobe structure of the response. Using the LS Prony method with $p = q = 10$, a digital filter was designed to match the desired low-pass filter frequency response. The result is shown in Figure 8.13c, where the magnitude of the frequency response or $|\mathcal{H}(\exp(j2\pi f))|$ has been plotted in dB. The agreement is generally good, with the Prony digital filter exhibiting a peaky structure in the pass-band (frequencies below the cutoff) and a smooth rolloff in the stopband (frequencies above the cutoff) in contrast to the desired response. Larger values for p and q would presumably improve the match. $\diamond$

**Figure 8.13** Low pass filter design by least squares Prony method

Example 8.12 - AR Parameter Estimation for the ARMA Model

We now describe a method for estimation of the AR parameters of an autoregressive moving average (ARMA) model. An ARMA model for a WSS random process assumes a PSD of

$$P_{xx}(f) = \frac{\sigma_u^2 |B(f)|^2}{|A(f)|^2}$$

where

$$B(f) = 1 + \sum_{k=1}^q b[k] \exp(-j2\pi f k)$$

$$A(f) = 1 + \sum_{k=1}^p a[k] \exp(-j2\pi f k).$$

The $b[k]$'s are termed the MA filter parameters, the $a[k]$'s the AR filter parameters, and σ_u^2 the driving white noise variance. The process is obtained by exciting a *causal* filter whose frequency response is $B(f)/A(f)$ with white noise of variance σ_u^2 . If the $b[k]$'s are zero, then we have the AR process introduced in Example 3.16. The estimation of the AR parameters of an AR model using an asymptotic MLE was discussed in Example 7.18. For the ARMA process even the asymptotic MLE proves to be intractable. An alternative approach relies on *equation error modeling* of the ACF. It estimates the AR filter parameters, leaving the MA filter parameters and white noise variance to be found by other means.

To determine the ACF we can take the inverse z transform of the PSD extended to the z plane or

$$\mathcal{P}_{xx}(z) = \frac{\sigma^2 \mathcal{B}(z) \mathcal{B}(z^{-1})}{\mathcal{A}(z) \mathcal{A}(z^{-1})}$$

where $B(f) = \mathcal{B}(\exp(j2\pi f))$, $A(f) = \mathcal{A}(\exp(j2\pi f))$. The usual PSD is obtained by evaluating $\mathcal{P}_{xx}(z)$ on the unit circle of the z plane or for $z = \exp(j2\pi f)$. It is convenient to develop a difference equation for the ACF. This difference equation will serve as the basis for the equation error modeling approach. Taking the inverse z transform of $\mathcal{A}(z) \mathcal{P}_{xx}(z)$, we have

$$\mathcal{Z}^{-1}\{\mathcal{A}(z) \mathcal{P}_{xx}(z)\} = \mathcal{Z}^{-1}\left\{\sigma^2 \mathcal{B}(z) \frac{\mathcal{B}(z^{-1})}{\mathcal{A}(z^{-1})}\right\}.$$

Since the filter impulse response is causal, it follows that

$$h[n] = \mathcal{Z}^{-1}\left\{\frac{\mathcal{B}(z)}{\mathcal{A}(z)}\right\}$$

$$= 0$$

for $n < 0$. Consequently, we have that

$$h[-n] = \mathcal{Z}^{-1}\left\{\frac{\mathcal{B}(z^{-1})}{\mathcal{A}(z^{-1})}\right\} = 0$$

for $n > 0$ and hence is anticausal. Thus,

$$\begin{aligned} \mathcal{Z}^{-1} \left\{ \sigma^2 \mathcal{B}(z) \frac{\mathcal{B}(z^{-1})}{\mathcal{A}(z^{-1})} \right\} &= \sigma^2 b[n] \star h[-n] \\ &= 0 \quad \text{for } n > q. \end{aligned}$$

Continuing, we have

$$\begin{aligned} \mathcal{Z}^{-1} \{ \mathcal{A}(z) \mathcal{P}_{xx}(z) \} &= \mathcal{Z}^{-1} \left\{ \sigma^2 \mathcal{B}(z) \frac{\mathcal{B}(z^{-1})}{\mathcal{A}(z^{-1})} \right\} \\ &= 0 \quad \text{for } n > q. \end{aligned}$$

Finally, a difference equation for the ACF can be written as

$$\sum_{k=0}^p a[k] r_{xx}[n-k] = 0 \quad \text{for } n > q \quad (8.63)$$

where $a[0] = 1$. These equations are called the *modified Yule-Walker equations*. Recall from Example 7.18 that the Yule-Walker equations for an AR process were identical except that they held for $n > 0$ (since $q = 0$ for an AR process). Only the AR filter parameters appear in these equations.

In practice, we must estimate the ACF lags as

$$\hat{r}_{xx}[k] = \frac{1}{N} \sum_{n=0}^{N-1-|k|} x[n]x[n+|k|],$$

assuming $x[n]$ is available for $n = 0, 1, \dots, N-1$. Substituting these into (8.63) yields

$$\sum_{k=0}^p a[k] \hat{r}_{xx}[n-k] = \epsilon[n] \quad n > q$$

where $\epsilon[n]$ denotes the error due to the effect of errors in the ACF function estimate. The model becomes

$$\hat{r}_{xx}[n] = - \sum_{k=1}^p a[k] \hat{r}_{xx}[n-k] + \epsilon[n] \quad n > q$$

which can be seen to be linear in the unknown AR filter parameters. If the ACF is estimated for lags $n = 0, 1, \dots, M$ (where we must have $M \leq N-1$), then a LSE of $a[k]$ will minimize

$$\begin{aligned} J &= \sum_{n=q+1}^M \left[\hat{r}_{xx}[n] - \left(- \sum_{k=1}^p a[k] \hat{r}_{xx}[n-k] \right) \right]^2 \\ &= (\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}) \end{aligned} \quad (8.64)$$

where

$$\mathbf{x} = \begin{bmatrix} \hat{r}_{xx}[q+1] \\ \hat{r}_{xx}[q+2] \\ \vdots \\ \hat{r}_{xx}[M] \end{bmatrix}$$

$$\boldsymbol{\theta} = \begin{bmatrix} a[1] \\ a[2] \\ \vdots \\ a[p] \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} \hat{r}_{xx}[q] & \hat{r}_{xx}[q-1] & \dots & \hat{r}_{xx}[q-p+1] \\ \hat{r}_{xx}[q+1] & \hat{r}_{xx}[q] & \dots & \hat{r}_{xx}[q-p+2] \\ \vdots & \vdots & \ddots & \vdots \\ \hat{r}_{xx}[M-1] & \hat{r}_{xx}[M-2] & \dots & \hat{r}_{xx}[M-p] \end{bmatrix}.$$

The LSE of $\boldsymbol{\theta}$ is $(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$ and is termed the *least squares modified Yule-Walker equations*. It is interesting to observe that in this problem we use a LSE for the estimated ACF data, not the original data. Additionally, the observation matrix $\mathbf{H}$, which is usually a known deterministic matrix, is now a random matrix. As expected, the statistics of the LSE are difficult to determine. In practice, M should not be chosen too large since the ACF estimate is less reliable at higher lags due to the $(N-k)$ lag products averaged in the estimate of $r_{xx}[k]$. Some researchers advocate a weighted LSE to reflect the tendency for the errors of (8.64) to increase as n increases. As an example, the choice $w_n = 1 - (n/(M+1))$ might be made and the weighted LSE of (8.16) implemented. Further discussions of this problem can be found in [Kay 1988].

◇

Example 8.13 - Adaptive Noise Canceler

A common problem in signal processing is to reduce unwanted noise. It arises in such diverse problems as the suppression of a 60 Hz interference in an electronic circuit and the suppression of a mother's heartbeat masking the fetal heartbeat in an EKG trace [Widrow and Stearns 1985]. A particularly effective means of accomplishing this is the adaptive noise canceler (ANC). It assumes the availability of a reference noise that after suitable filtering can be subtracted from the noise that is to be canceled. As an example, to cancel a 60 Hz interference we need to know the amplitude and phase of the interference, both of which are generally unknown. A suitable adaptive filter as shown in Figure 8.14 can modify the amplitude and phase so as to make it nearly identical to the interference to be canceled. A discrete-time ANC is shown in Figure 8.15. The primary channel contains noise $x[n]$ to be canceled. To effect cancellation a reference channel containing a known sequence $x_R[n]$ that is similar but not identical to the noise is then filtered to improve the match. The filter $\mathcal{H}_n(z)$ has weights that are determined




Figure 8.14 Adaptive noise canceler for 60 Hz interference




Figure 8.15 Generic adaptive noise canceler

at each time n to make $\hat{x}[k] \approx x[k]$ for $k = 0, 1, \dots, n$. Since we want $\epsilon[n] \approx 0$, it makes sense to choose the weights at time n to minimize

$$\begin{aligned} J[n] &= \sum_{k=0}^n \epsilon^2[k] \\ &= \sum_{k=0}^n (x[k] - \hat{x}[k])^2 \\ &= \sum_{k=0}^n \left(x[k] - \sum_{l=0}^{p-1} h_n[l]x_R[k-l] \right)^2 \end{aligned}$$




Figure 8.16 Stationarity of interference model

where $h_n[l]$ are the filter weights at time n . It is immediately recognized that the weights can be determined as the solution of a sequential LS problem. Before giving the solution we remark that the LSE assumes that the interference is fixed over time or that it is stationary. For instance, the primary channel containing the sinusoid is assumed to appear as shown in Figure 8.16a. If, however, the interference appears as in Figure 8.16b, then the adaptive filter must quickly change its coefficients to respond to the changing interference for $t > t_0$. How fast it can change depends upon how many error terms are included in $J[n]$ before and after the transition at $t = t_0$. The chosen weights for $t > t_0$ will be a compromise between the old weights and the desired new ones. If the transition occurs in discrete time at $n = n_0$, then we should probably expect the wrong weights until $n \gg n_0$. To allow the filter to adapt more quickly we can downweight previous errors in $J[n]$ by incorporating a weighting or “forgetting factor” λ , where $0 < \lambda < 1$, as follows:

$$J[n] = \sum_{k=0}^n \lambda^{n-k} \left(x[k] - \sum_{l=0}^{p-1} h_n[l] x_R[k-l] \right)^2.$$

This modification downweights previous errors exponentially, allowing the filter to react more quickly to interference changes. The penalty paid is that the estimates of the filter weights are noisier, relying on fewer effective LS errors. Note that the solution will not change if we minimize instead

$$J'[n] = \sum_{k=0}^n \frac{1}{\lambda^k} \left(x[k] - \sum_{l=0}^{p-1} h_n[l] x_R[k-l] \right)^2 \quad (8.65)$$

for each n . Now we have our standard sequential *weighted* LS problem described in Section 8.7. Referring to (8.46), we identify the sequential LSE of the filter weights as

$$\hat{\theta}[n] = \begin{bmatrix} \hat{h}_n[0] & \hat{h}_n[1] & \dots & \hat{h}_n[p-1] \end{bmatrix}^T.$$

The data vector $\mathbf{h}[n]$ is given as (not to be confused with the impulse response $h[n]$)

$$\mathbf{h}[n] = [x_R[n] \quad x_R[n-1] \quad \dots \quad x_R[n-p+1]]^T$$

and the weights σ_n^2 are given by λ^n . Also, note that

$$\begin{aligned} e[n] &= x[n] - \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n-1] \\ &= x[n] - \sum_{l=0}^{p-1} \hat{h}_{n-1}[l] x_R[n-l] \end{aligned}$$

is the error at time n based on the filter weights at time $n-1$. (This is not the same as $\epsilon[n] = x[n] - \hat{x}[n] = x[n] - \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n]$.) The algorithm is summarized from (8.46)–(8.48) as

$$\hat{\boldsymbol{\theta}}[n] = \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n]e[n]$$

where

$$\begin{aligned} e[n] &= x[n] - \sum_{l=0}^{p-1} \hat{h}_{n-1}[l] x_R[n-l] \\ \mathbf{K}[n] &= \frac{\boldsymbol{\Sigma}[n-1] \mathbf{h}[n]}{\lambda^n + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]} \\ \mathbf{h}[n] &= [x_R[n] \quad x_R[n-1] \quad \dots \quad x_R[n-p+1]]^T \\ \boldsymbol{\Sigma}[n] &= (\mathbf{I} - \mathbf{K}[n] \mathbf{h}^T[n]) \boldsymbol{\Sigma}[n-1]. \end{aligned}$$

The forgetting factor is chosen as $0 < \lambda < 1$, with λ near 1 or $0.9 < \lambda < 1$ being typical. As an example, for a sinusoidal interference

$$x[n] = 10 \cos(2\pi(0.1)n + \pi/4)$$

and a reference signal


$$x_R[n] = \cos(2\pi(0.1)n)$$

we implement the ANC using two filter coefficients ($p = 2$) since the reference signal need only be modified in amplitude and phase to match the interference. To initialize the sequential LS estimator we choose

$$\begin{aligned} \hat{\boldsymbol{\theta}}[-1] &= \mathbf{0} \\ \boldsymbol{\Sigma}[-1] &= 10^5 \mathbf{I} \end{aligned}$$

and a forgetting factor of $\lambda = 0.99$ is assumed. The interference $x[n]$ and output $\epsilon[n]$ of the ANC are shown in Figures 8.17a and 8.17b, respectively. As expected, the interference is canceled. The LSE of the filter coefficients is shown in Figure 8.17c and is seen to rapidly converge to the steady-state value. To find the steady state values of the weights we need to solve

$$\mathcal{H}[\exp(2\pi(0.1))] = 10 \exp(j\pi/4)$$

**Figure 8.17** Interference cancellation example

which says that the adaptive filter must increase the gain of the reference signal by 10 and the phase by $\pi/4$ to match the interference. Solving, we have

$$h[0] + h[1] \exp(-j2\pi(0.1)) = 10 \exp(j\pi/4)$$

which results in $h[0] = 16.8$ and $h[1] = -12.0$. $\diamond$

Example 8.14 - Phase-Locked Loop

We now consider the problem of carrier recovery in a communication system, which is necessary for coherent demodulation. It is assumed that the carrier is received but is embedded in noise. The received noise-free carrier is

$$s[n] = \cos(2\pi f_0 n + \phi) \quad n = -M, \dots, 0, \dots, M$$

where the frequency f_0 and phase ϕ are to be estimated. A symmetric observation interval is chosen to simplify the algebra. A LSE is to be used. Due to the nonlinearity we employ the linearization technique of (8.62). Determining first $\mathbf{H}(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = [f_0 \ \phi]^T$, we have

$$\begin{aligned} \frac{\partial s[n]}{\partial f_0} &= -n2\pi \sin(2\pi f_0 n + \phi) \\ \frac{\partial s[n]}{\partial \phi} &= -\sin(2\pi f_0 n + \phi) \end{aligned}$$

so that

$$\mathbf{H}(\boldsymbol{\theta}) = - \begin{bmatrix} -M2\pi \sin(-2\pi f_0 M + \phi) & \sin(-2\pi f_0 M + \phi) \\ -(M-1)2\pi \sin(-2\pi f_0 (M-1) + \phi) & \sin(-2\pi f_0 (M-1) + \phi) \\ \vdots & \vdots \\ M2\pi \sin(2\pi f_0 M + \phi) & \sin(2\pi f_0 M + \phi) \end{bmatrix}$$

and

$$\mathbf{H}^T(\boldsymbol{\theta})\mathbf{H}(\boldsymbol{\theta}) = \begin{bmatrix} 4\pi^2 \sum_{n=-M}^M n^2 \sin^2(2\pi f_0 n + \phi) & 2\pi \sum_{n=-M}^M n \sin^2(2\pi f_0 n + \phi) \\ 2\pi \sum_{n=-M}^M n \sin^2(2\pi f_0 n + \phi) & \sum_{n=-M}^M \sin^2(2\pi f_0 n + \phi) \end{bmatrix}.$$

But

$$\sum_{n=-M}^M n^2 \sin^2(2\pi f_0 n + \phi) = \sum_{n=-M}^M \left[\frac{n^2}{2} - \frac{n^2}{2} \cos(4\pi f_0 n + 2\phi) \right]$$

$$\begin{aligned}\sum_{n=-M}^M n \sin^2(2\pi f_0 n + \phi) &= \sum_{n=-M}^M \left[\frac{n}{2} - \frac{n}{2} \cos(4\pi f_0 n + 2\phi) \right] \\ \sum_{n=-M}^M \sin^2(2\pi f_0 n + \phi) &= \sum_{n=-M}^M \left[\frac{1}{2} - \frac{1}{2} \cos(4\pi f_0 n + 2\phi) \right]\end{aligned}$$

and since [Stoica et al. 1989]

$$\frac{1}{(2M+1)^{i+1}} \sum_{n=-M}^M n^i \cos(4\pi f_0 n + 2\phi) \approx 0 \quad i = 0, 1, 2$$

we have from (3.22)

$$\begin{aligned}\mathbf{H}^T(\boldsymbol{\theta})\mathbf{H}(\boldsymbol{\theta}) &\approx \begin{bmatrix} \frac{4\pi^2}{3}M(M+1)(2M+1) & 0 \\ 0 & \frac{2M+1}{2} \end{bmatrix} \\ &\approx \begin{bmatrix} \frac{8\pi^2 M^3}{3} & 0 \\ 0 & M \end{bmatrix}\end{aligned}$$

for $M \gg 1$. We have then from (8.62)

$$\begin{aligned}f_{0_{k+1}} &= f_{0_k} - \frac{3}{4\pi M^3} \sum_{n=-M}^M n \sin(2\pi f_{0_k} n + \phi_k) (x[n] - \cos(2\pi f_{0_k} n + \phi_k)) \\ \phi_{k+1} &= \phi_k - \frac{1}{M} \sum_{n=-M}^M \sin(2\pi f_{0_k} n + \phi_k) (x[n] - \cos(2\pi f_{0_k} n + \phi_k))\end{aligned}$$

or since

$$\begin{aligned}\frac{1}{2M+1} \sum_{n=-M}^M \sin(2\pi f_{0_k} n + \phi_k) \cos(2\pi f_{0_k} n + \phi_k) \\ = \frac{1}{2(2M+1)} \sum_{n=-M}^M \sin(4\pi f_{0_k} n + 2\phi_k) \\ \approx 0\end{aligned}$$

it follows that

$$\begin{aligned}f_{0_{k+1}} &= f_{0_k} - \frac{3}{4\pi M^3} \sum_{n=-M}^M n x[n] \sin(2\pi f_{0_k} n + \phi_k) \\ \phi_{k+1} &= \phi_k - \frac{1}{M} \sum_{n=-M}^M x[n] \sin(2\pi f_{0_k} n + \phi_k).\end{aligned}$$

The LSE of frequency and phase can be implemented as a phase-locked loop [Proakis 1983]. In Problem 8.29 it is shown that at convergence for a high enough SNR the solution will be

$$\begin{aligned} f_{0k+1} &= f_{0k} \approx f_0 \\ \phi_{k+1} &= \phi_k \approx \phi. \end{aligned}$$

◇

References

- Bard, Y., *Nonlinear Parameter Estimation*, Academic Press, New York, 1974.
- Graybill, F.A., *Theory and Application of the Linear Model*, Duxbury Press, North Scituate, Mass., 1976.
- Kay, S.M., *Modern Spectral Estimation: Theory and Application*, Prentice-Hall, Englewood Cliffs, N.J., 1988.
- Oppenheim, A.V., R.W. Schaffer, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1975.
- Parks, T.W., C.S. Burrus, *Digital Filter Design*, J. Wiley, New York, 1987.
- Proakis, J.G., *Digital Communications*, McGraw-Hill, New York, 1983.
- Scharf, L.L., *Statistical Signal Processing*, Addison-Wesley, New York, 1991.
- Seber, G.A.F., C.J. Wild, *Nonlinear Regression*, J. Wiley, New York, 1989.
- Stoica, P., R.L. Moses, B. Friedlander, T. Soderstrom, "Maximum Likelihood Estimation of the Parameters of Multiple Sinusoids from Noisy Measurements," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 37, pp. 378-392, March 1989.
- Widrow, B., S.D. Stearns, *Adaptive Signal Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1985.

Problems

8.1 The LS error

$$J = \sum_{n=0}^{N-1} (x[n] - A \cos 2\pi f_0 n - B r^n)^2$$

is to be minimized to find the LSE of $\theta = [A f_0 B r]^T$, where $0 < r < 1$. Is this a linear or nonlinear LS problem? Is the LS error quadratic in any of the parameters, and if so, which ones? How could you solve this minimization problem using a digital computer?

8.2 Show that the inequality given in (8.7) holds.

8.3 For the signal model

$$s[n] = \begin{cases} A & 0 \leq n \leq M-1 \\ -A & M \leq n \leq N-1, \end{cases}$$

find the LSE of A and the minimum LS error. Assume that $x[n] = s[n] + w[n]$ for $n = 0, 1, \dots, N-1$ are observed. If now $w[n]$ is WGN with variance σ^2 , find the PDF of the LSE.

8.4 Derive the LSE for a vector parameter as given by (8.10) by verifying the identity

$$J = (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T(\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}}) + (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T \mathbf{H}^T \mathbf{H} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$$

where

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}.$$

To complete the proof show that J is minimized when $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, assuming that $\mathbf{H}$ is full rank and therefore that $\mathbf{H}^T \mathbf{H}$ is positive definite.

8.5 For the signal model

$$s[n] = \sum_{i=1}^p A_i \cos 2\pi f_i n$$

where the frequencies f_i are known and the amplitudes A_i are to be estimated, find the LSE *normal equations* (do not attempt to solve them). Then, if the frequencies are specifically known to be $f_i = i/N$, explicitly find the LSE and the minimum LS error. Finally, if $x[n] = s[n] + w[n]$, where $w[n]$ is WGN with variance σ^2 , determine the PDF of the LSE, assuming the given frequencies. Hint: The columns of $\mathbf{H}$ are orthogonal for the given frequencies.

8.6 For the general LSE of (8.10) find the PDF of the LSE if it is known that $\mathbf{x} \sim \mathcal{N}(\mathbf{H}\boldsymbol{\theta}, \sigma^2 \mathbf{I})$. Is the LSE unbiased?

8.7 In this problem we consider the estimation of the noise variance σ^2 in the model $\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$, where $\mathbf{w}$ is zero mean with covariance matrix $\sigma^2 \mathbf{I}$. The estimator

$$\hat{\sigma}^2 = \frac{1}{N} J_{\min} = \frac{1}{N} (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})$$

where $\hat{\boldsymbol{\theta}}$ is the LSE of (8.10), is proposed. Is this estimator unbiased, and if not, can you propose one that is? Explain your results. Hint: The identities $E(\mathbf{x}^T \mathbf{y}) = E(\text{tr}(\mathbf{y} \mathbf{x}^T)) = \text{tr}(E(\mathbf{y} \mathbf{x}^T))$ and $\text{tr}(\mathbf{A}\mathbf{B}) = \text{tr}(\mathbf{B}\mathbf{A})$ will be useful.

8.8 Verify the LSE for A given in (8.15). Find the mean and variance for $\hat{A}$ if $x[n] = A + w[n]$, where $w[n]$ is zero mean uncorrelated noise with variance σ_n^2 .

8.9 Verify (8.16) and (8.17) for the weighted LSE by noting that if $\mathbf{W}$ is positive definite, we can write it as $\mathbf{W} = \mathbf{D}^T \mathbf{D}$, where $\mathbf{D}$ is an invertible $N \times N$ matrix.

8.10 Referring to Figure 8.2b, prove that

$$\|\hat{\mathbf{s}}\|^2 + \|\mathbf{x} - \hat{\mathbf{s}}\|^2 = \|\mathbf{x}\|^2.$$

This can be thought of as the least squares Pythagorean theorem.

8.11 In this problem we prove that a projection matrix $\mathbf{P}$ must be symmetric. Let $\mathbf{x} = \boldsymbol{\xi} + \boldsymbol{\xi}^\perp$, where $\boldsymbol{\xi}$ lies in a subspace which is the range of the projection matrix or $\mathbf{P}\mathbf{x} = \boldsymbol{\xi}$ and $\boldsymbol{\xi}^\perp$ lies in the orthogonal subspace or $\mathbf{P}\boldsymbol{\xi}^\perp = \mathbf{0}$. For arbitrary vectors $\mathbf{x}_1, \mathbf{x}_2$ in R^N show that

$$\mathbf{x}_1^T \mathbf{P} \mathbf{x}_2 - \mathbf{x}_2^T \mathbf{P} \mathbf{x}_1 = 0$$

by decomposing $\mathbf{x}_1$ and $\mathbf{x}_2$ as discussed above. Finally, prove the desired result.

8.12 Prove the following properties of the projection matrix

$$\mathbf{P} = \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T.$$

- $\mathbf{P}$ is idempotent.
- $\mathbf{P}$ is positive semidefinite.
- The eigenvalues of $\mathbf{P}$ are either 1 or 0.
- The rank of $\mathbf{P}$ is p . Use the fact that the trace of a matrix is equal to the sum of its eigenvalues.

8.13 Verify the result given in (8.24).

8.14 In this problem we derive the order-update LS equations using a geometrical argument. Assume that $\hat{\boldsymbol{\theta}}_k$ is available and therefore

$$\hat{\mathbf{s}}_k = \mathbf{H}_k \hat{\boldsymbol{\theta}}_k$$

is known. Now, if $\mathbf{h}_{k+1}$ is used, the LS signal estimate becomes

$$\hat{\mathbf{s}}_{k+1} = \hat{\mathbf{s}}_k + \alpha \mathbf{h}'_{k+1}$$

where $\mathbf{h}'_{k+1}$ is the component of $\mathbf{h}_{k+1}$ orthogonal to the subspace spanned by $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_k\}$. First, find $\mathbf{h}'_{k+1}$ and then determine α by noting that $\alpha \mathbf{h}'_{k+1}$ is the projection of $\mathbf{x}$ onto $\mathbf{h}'_{k+1}$. Finally, since

$$\hat{\mathbf{s}}_{k+1} = \mathbf{H}_{k+1} \hat{\boldsymbol{\theta}}_{k+1}$$

determine $\hat{\boldsymbol{\theta}}_{k+1}$.

8.15 Use order-recursive LS to find the values of A and B that minimize

$$J = \sum_{n=0}^{N-1} (x[n] - A - Br^n)^2.$$

The parameter r is assumed to be known.

8.16 Consider the models

$$\begin{aligned}s_1[n] &= A & 0 \leq n \leq N-1 \\ s_2[n] &= \begin{cases} A & 0 \leq n \leq M-1 \\ B & M \leq n \leq N-1. \end{cases}\end{aligned}$$

The second signal is useful for modeling a jump in level at $n = M$. We can express the second model in the alternative form

$$s[n] = Au_1[n] + (B - A)u_2[n]$$

where

$$\begin{aligned}u_1[n] &= 1 & 0 \leq n \leq N-1 \\ u_2[n] &= \begin{cases} 0 & 0 \leq n \leq M-1 \\ 1 & M \leq n \leq N-1. \end{cases}\end{aligned}$$

Now let $\theta_1 = A$ and $\theta_2 = [A \ (B - A)]^T$ and find the LSE and minimum LS error for each model using order-recursive LS. Discuss how you might use it to detect a jump in level.

- 8.17** Prove that $\mathbf{P}_k^\perp \mathbf{x}$ is orthogonal to the space spanned by the columns of $\mathbf{H}_k$.
- 8.18** Using the orthogonal recursive projection matrix formula (8.34), derive the update formula for the minimum LS error (8.31).
- 8.19** Verify the sequential formula for the minimum LS error (8.43) by using the sequential LSE update (8.40).
- 8.20** Let $x[n] = Ar^n + w[n]$, where $w[n]$ is WGN with variance $\sigma^2 = 1$. Find the sequential LSE for A , assuming that r is known. Also, determine the variance of the LSE in sequential form and then solve explicitly for the variance as a function of n . Let $\hat{A}[0] = x[0]$ and $\text{var}(\hat{A}[0]) = \sigma^2 = 1$.
- 8.21** Using the sequential update formulas for the gain (8.41) and variance (8.42), solve for the gain and variance sequences if $\sigma_n^2 = r^n$. Use $\text{var}(\hat{A}[0]) = \text{var}(x[0]) = \sigma_0^2 = 1$ for initialization. Then, examine what happens to the gain and variance as $N \rightarrow \infty$ if $r = 1$, $0 < r < 1$, and $r > 1$. Hint: Solve for $1/\text{var}(\hat{A}[N])$.
- 8.22** By implementing a Monte Carlo computer simulation plot $\hat{A}[N]$ as given by (8.40). Assume that the data are given by

$$x[n] = A + w[n]$$

where $w[n]$ is zero mean WGN with $\sigma_n^2 = r^n$. Use $A = 10$ and $r = 1, 0.95, 1.05$. Initialize the estimator by using $\hat{A}[0] = x[0]$ and $\text{var}(\hat{A}[0]) = \text{var}(x[0]) = \sigma_0^2 = 1$. Also, plot the gain and variance sequences.

8.23 In this problem we examine the initialization of a sequential LSE. Assume that we choose $\hat{\theta}[-1]$ and $\Sigma[-1] = \alpha \mathbf{I}$ to initialize the LSE. We will show that, as $\alpha \rightarrow \infty$, the batch LSE

$$\begin{aligned}\hat{\theta}_B[n] &= (\mathbf{H}^T[n] \mathbf{C}^{-1}[n] \mathbf{H}[n])^{-1} \mathbf{H}^T[n] \mathbf{C}^{-1}[n] \mathbf{x}[n] \\ &= \left(\sum_{k=0}^n \frac{1}{\sigma_k^2} \mathbf{h}[k] \mathbf{h}^T[k] \right)^{-1} \left(\sum_{k=0}^n \frac{1}{\sigma_k^2} x[k] \mathbf{h}^T[k] \right)\end{aligned}$$

for $n \geq p$ is identical to the sequential LSE. First, assume that the initial observation vectors $\{\mathbf{h}[-p], \mathbf{h}[-(p-1)], \dots, \mathbf{h}[-1]\}$ and noise variances $\{\sigma_{-p}^2, \sigma_{-(p-1)}^2, \dots, \sigma_{-1}^2\}$ exist, so that for the chosen initial LSE and covariance we can use a batch estimator. Hence,

$$\hat{\theta}[-1] = \hat{\theta}_B[-1] = (\mathbf{H}^T[-1] \mathbf{C}^{-1}[-1] \mathbf{H}[-1])^{-1} \mathbf{H}^T[-1] \mathbf{C}^{-1}[-1] \mathbf{x}[-1]$$

and

$$\Sigma[-1] = (\mathbf{H}^T[-1] \mathbf{C}^{-1}[-1] \mathbf{H}[-1])^{-1}$$

where

$$\begin{aligned}\mathbf{H}[-1] &= \begin{bmatrix} \mathbf{h}^T[-p] \\ \mathbf{h}^T[-(p-1)] \\ \vdots \\ \mathbf{h}^T[-1] \end{bmatrix} \\ \mathbf{C}[-1] &= \text{diag}(\sigma_{-p}^2, \sigma_{-(p-1)}^2, \dots, \sigma_{-1}^2).\end{aligned}$$

Thus, we may view the initial estimate for the sequential LSE as the result of applying a batch estimator to the initial observation vectors. Since the sequential LSE initialized using a batch estimator is identical to the batch LSE using all the observation vectors, we have for the sequential LSE with the assumed initial conditions

$$\hat{\theta}_S[n] = \left(\sum_{k=-p}^n \frac{1}{\sigma_k^2} \mathbf{h}[k] \mathbf{h}^T[k] \right)^{-1} \left(\sum_{k=-p}^n \frac{1}{\sigma_k^2} x[k] \mathbf{h}^T[k] \right).$$

Prove that this can be rewritten as

$$\begin{aligned}\hat{\theta}_S[n] &= \\ &\left(\Sigma^{-1}[-1] + \sum_{k=0}^n \frac{1}{\sigma_k^2} \mathbf{h}[k] \mathbf{h}^T[k] \right)^{-1} \left(\Sigma^{-1}[-1] \hat{\theta}[-1] + \sum_{k=0}^n \frac{1}{\sigma_k^2} x[k] \mathbf{h}^T[k] \right).\end{aligned}$$

Then examine what happens as $\alpha \rightarrow \infty$ for $0 \leq n \leq p-1$ and $n \geq p$.

8.24 In Example 8.8 determine $\hat{s}_c$ by first projecting $\mathbf{x}$ onto the subspace spanned by $\mathbf{h}_1$ and $\mathbf{h}_2$ to produce $\hat{\mathbf{s}}$, and then projecting $\hat{\mathbf{s}}$ onto the constrained subspace.

8.25 If the signal model is

$$s[n] = A + B(-1)^n \quad n = 0, 1, \dots, N-1$$

and N is even, find the LSE of $\boldsymbol{\theta} = [A \ B]^T$. Now assume that $A = B$ and repeat the problem using a constrained LS approach. Compare your results.

8.26 Consider the minimization of $h(\theta)$ with respect to θ and assume that

$$\hat{\theta} = \arg \min_{\theta} h(\theta).$$

Let $\alpha = g(\theta)$, where g is a one-to-one mapping. Prove that if $\hat{\alpha}$ minimizes $h(g^{-1}(\alpha))$, then $\hat{\theta} = g^{-1}(\hat{\alpha})$.

8.27 Let the observed data be

$$x[n] = \exp(\theta) + w[n] \quad n = 0, 1, \dots, N-1.$$

Set up a Newton-Raphson iteration to find the LSE of θ . Can you avoid the nonlinear optimization and find the LSE analytically?

8.28 In Example 8.11 the true LSE must minimize

$$J = \sum_{n=0}^{N-1} (h_d[n] - h[n])^2.$$

In this problem we derive the equations that must be optimized to find the true LSE [Scharf 1991]. It is assumed that $p = q + 1$. First, show that the signal model can be written as

$$\begin{aligned} \mathbf{s} &= \begin{bmatrix} h[0] \\ h[1] \\ \vdots \\ h[N-1] \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} g[0] & 0 & \dots & 0 \\ g[1] & g[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g[q] & g[q-1] & \dots & g[0] \\ g[q+1] & g[q] & \dots & g[1] \\ \vdots & \vdots & \vdots & \vdots \\ g[N-1] & g[N-2] & \dots & g[N-1-q] \end{bmatrix}}_{\mathbf{G}} \underbrace{\begin{bmatrix} b[0] \\ b[1] \\ \vdots \\ b[q] \end{bmatrix}}_{\mathbf{b}}. \end{aligned}$$

where

$$g[n] = \mathcal{Z}^{-1} \left\{ \frac{1}{\mathcal{A}(z)} \right\}$$

and $\mathbf{G}$ is $N \times (q+1)$. Letting $\mathbf{x} = [h_d[0] \ h_d[1] \ \dots \ h_d[N-1]]^T$, show that

$$J(\mathbf{a}, \hat{\mathbf{b}}) = \mathbf{x}^T (\mathbf{I} - \mathbf{G}(\mathbf{G}^T \mathbf{G})^{-1} \mathbf{G}^T) \mathbf{x}$$

where $\hat{\mathbf{b}}$ is the LSE of $\mathbf{b}$ for a given $\mathbf{a}$. Finally, prove that

$$\mathbf{I} - \mathbf{G}(\mathbf{G}^T \mathbf{G})^{-1} \mathbf{G}^T = \mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \quad (8.66)$$

where $\mathbf{A}^T$ is the $(N-q-1) \times N = (N-p) \times N$ matrix

$$\mathbf{A}^T = \begin{bmatrix} a[p] & a[p-1] & \dots & 1 & 0 & 0 & \dots & 0 \\ 0 & a[p] & \dots & a[1] & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & a[p] & a[p-1] & \dots & 1 \end{bmatrix}.$$

Hence, the LSE of $\mathbf{a}$ is found by minimizing

$$\mathbf{x}^T \mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{x}.$$

Once this is found (using a nonlinear optimization), we have

$$\begin{aligned} \hat{\mathbf{b}} &= (\mathbf{G}^T \mathbf{G})^{-1} \mathbf{G}^T \mathbf{x} \\ &= (\mathbf{I} - \mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T) \mathbf{x} \end{aligned}$$

where the elements in $\mathbf{A}$ are replaced by the LSE of $\mathbf{a}$. Note that this problem is an example of a separable nonlinear LS problem. Hint: To prove (8.66) consider $\mathbf{L} = [\mathbf{A} \ \mathbf{G}]$, which is invertible (since it is full rank), and compute

$$\mathbf{L}(\mathbf{L}^T \mathbf{L})^{-1} \mathbf{L}^T = \mathbf{I}.$$

You will also need to observe that $\mathbf{A}^T \mathbf{G} = \mathbf{0}$, which follows from $a[n] \star g[n] = \delta[n]$.

- 8.29** In Example 8.14 assume that the phase-locked loop converges so that $f_{0_{k+1}} = f_{0_k}$ and $\phi_{k+1} = \phi_k$. For a high SNR so that $x[n] \approx \cos(2\pi f_0 n + \phi)$, show that the final iterates will be the true values of frequency and phase.

Appendix 8A

Derivation of Order-Recursive Least Squares

In this appendix we derive the order-recursive least squares formulas (8.28)–(8.31). Referring to (8.25), we have

$$\hat{\boldsymbol{\theta}}_{k+1} = (\mathbf{H}_{k+1}^T \mathbf{H}_{k+1})^{-1} \mathbf{H}_{k+1}^T \mathbf{x}.$$

But

$$\begin{aligned} \mathbf{H}_{k+1}^T \mathbf{H}_{k+1} &= \begin{bmatrix} \mathbf{H}_k^T \\ \mathbf{h}_{k+1}^T \end{bmatrix} \begin{bmatrix} \mathbf{H}_k & \mathbf{h}_{k+1} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{H}_k^T \mathbf{H}_k & \mathbf{H}_k^T \mathbf{h}_{k+1} \\ \mathbf{h}_{k+1}^T \mathbf{H}_k & \mathbf{h}_{k+1}^T \mathbf{h}_{k+1} \end{bmatrix} = \begin{bmatrix} k \times k & k \times 1 \\ 1 \times k & 1 \times 1 \end{bmatrix}. \end{aligned}$$

Using the inverse of a partitioned matrix formula (see Appendix 1)

$$\begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^T & c \end{bmatrix}^{-1} = \begin{bmatrix} \left(\mathbf{A} - \frac{\mathbf{b}\mathbf{b}^T}{c}\right)^{-1} & -\frac{1}{c} \left(\mathbf{A} - \frac{\mathbf{b}\mathbf{b}^T}{c}\right)^{-1} \mathbf{b} \\ -\frac{1}{c} \mathbf{b}^T \left(\mathbf{A} - \frac{\mathbf{b}\mathbf{b}^T}{c}\right)^{-1} & \frac{1}{c - \mathbf{b}^T \mathbf{A}^{-1} \mathbf{b}} \end{bmatrix}$$

for $\mathbf{A}$ a symmetric $k \times k$ matrix, $\mathbf{b}$ a $k \times 1$ vector, and c a scalar, and also Woodbury's identity (see Appendix 1)

$$\left(\mathbf{A} - \frac{\mathbf{b}\mathbf{b}^T}{c}\right)^{-1} = \mathbf{A}^{-1} + \frac{\mathbf{A}^{-1} \mathbf{b} \mathbf{b}^T \mathbf{A}^{-1}}{c - \mathbf{b}^T \mathbf{A}^{-1} \mathbf{b}}$$

we have

$$(\mathbf{H}_{k+1}^T \mathbf{H}_{k+1})^{-1} = \begin{bmatrix} \underbrace{\mathbf{D}_k + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1} - \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}}_{\mathbf{E}_k} & -\frac{\mathbf{E}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} \\ -\frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{E}_k}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} & \frac{1}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1} - \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}} \end{bmatrix}$$

where $\mathbf{D}_k = (\mathbf{H}_k^T \mathbf{H}_k)^{-1}$ or

$$\mathbf{D}_{k+1} = \begin{bmatrix} \mathbf{D}_k + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & -\frac{\mathbf{E}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} \\ -\frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{E}_k}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} & \frac{1}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix}.$$

Simplifying the off-diagonal block of $\mathbf{D}_{k+1}$, we have

$$\begin{aligned} \frac{\mathbf{E}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} &= \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} \\ &= \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} [\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1} + \mathbf{h}_{k+1}^T (\mathbf{I} - \mathbf{P}_k^\perp) \mathbf{h}_{k+1}]}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{h}_{k+1}} \\ &= \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}. \end{aligned}$$

Finally, we have

$$\begin{aligned} \mathbf{D}_{k+1} &= \begin{bmatrix} \mathbf{D}_k + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & -\frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ -\frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & \frac{1}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} \\ &= \begin{bmatrix} k \times k & k \times 1 \\ 1 \times k & 1 \times 1 \end{bmatrix}. \end{aligned}$$

To find $\hat{\boldsymbol{\theta}}_{k+1}$ we make use of the previous result

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{k+1} &= \mathbf{D}_{k+1} \mathbf{H}_{k+1}^T \mathbf{x} \\ &= \mathbf{D}_{k+1} \begin{bmatrix} \mathbf{H}_k^T \mathbf{x} \\ \mathbf{h}_{k+1}^T \mathbf{x} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{D}_k \mathbf{H}_k^T \mathbf{x} + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} - \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ -\frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} + \frac{\mathbf{h}_{k+1}^T \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} \end{aligned}$$

$$\begin{aligned}
&= \left[\hat{\theta}_k - \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T (\mathbf{x} - \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{x})}{\frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T (\mathbf{x} - \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{x})}} \right] \\
&= \left[\hat{\theta}_k - \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}} \right].
\end{aligned}$$

Finally, the update for the minimum LS error is

$$\begin{aligned}
J_{\min_{k+1}} &= (\mathbf{x} - \mathbf{H}_{k+1} \hat{\theta}_{k+1})^T (\mathbf{x} - \mathbf{H}_{k+1} \hat{\theta}_{k+1}) \\
&= \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \mathbf{H}_{k+1} \hat{\theta}_{k+1} \\
&= \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \begin{bmatrix} \mathbf{H}_k & \mathbf{h}_{k+1} \end{bmatrix} \begin{bmatrix} \hat{\theta}_k - \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}} \\ \frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} \\
&= \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \mathbf{H}_k \hat{\theta}_k + \frac{\mathbf{x}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\
&\quad - \mathbf{x}^T \mathbf{h}_{k+1} \frac{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\
&= J_{\min_k} - \frac{\mathbf{x}^T (\mathbf{I} - \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T) \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\
&= J_{\min_k} - \frac{\mathbf{x}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\
&= J_{\min_k} - \frac{(\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{x})^2}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}.
\end{aligned}$$

Appendix 8B

Derivation of Recursive Projection Matrix

We now derive the recursive update for the projection matrix given by (8.34). Since

$$\begin{aligned}\mathbf{P}_{k+1} &= \mathbf{H}_{k+1} (\mathbf{H}_{k+1}^T \mathbf{H}_{k+1})^{-1} \mathbf{H}_{k+1}^T \\ &= \mathbf{H}_{k+1} \mathbf{D}_{k+1} \mathbf{H}_{k+1}^T\end{aligned}$$

where $\mathbf{D}_{k+1} = (\mathbf{H}_{k+1}^T \mathbf{H}_{k+1})^{-1}$, we can use the update for $\mathbf{D}_k$ derived in Appendix 8A. As a result, we have

$$\begin{aligned}\mathbf{P}_{k+1} &= \begin{bmatrix} \mathbf{H}_k & \mathbf{h}_{k+1} \end{bmatrix} \begin{bmatrix} \mathbf{D}_k + \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} - \frac{\mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1}}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ - \frac{\mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} & \frac{1}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \end{bmatrix} \begin{bmatrix} \mathbf{H}_k^T \\ \mathbf{h}_{k+1}^T \end{bmatrix} \\ &= \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T + \frac{\mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ &\quad - \frac{\mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ &\quad - \frac{\mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{H}_k \mathbf{D}_k \mathbf{H}_k^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} + \frac{\mathbf{h}_{k+1} \mathbf{h}_{k+1}^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ &= \mathbf{P}_k + \frac{\mathbf{P}_k \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} - \frac{\mathbf{P}_k \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ &\quad - \frac{\mathbf{h}_{k+1} \mathbf{h}_{k+1}^T \mathbf{P}_k}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} + \frac{\mathbf{h}_{k+1} \mathbf{h}_{k+1}^T}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}} \\ &= \mathbf{P}_k + \frac{(\mathbf{I} - \mathbf{P}_k) \mathbf{h}_{k+1} \mathbf{h}_{k+1}^T (\mathbf{I} - \mathbf{P}_k)}{\mathbf{h}_{k+1}^T \mathbf{P}_k^\perp \mathbf{h}_{k+1}}.\end{aligned}$$

Appendix 8C

Derivation of Sequential Least Squares

In this appendix we derive (8.46)–(8.48).

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= (\mathbf{H}^T[n]\mathbf{C}^{-1}[n]\mathbf{H}[n])^{-1} \mathbf{H}^T[n]\mathbf{C}^{-1}[n]\mathbf{x}[n] \\ &= \left(\begin{bmatrix} \mathbf{H}^T[n-1] & \mathbf{h}[n] \end{bmatrix} \begin{bmatrix} \mathbf{C}[n-1] & \mathbf{0} \\ \mathbf{0}^T & \sigma_n^2 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{H}[n-1] \\ \mathbf{h}[n] \end{bmatrix} \right)^{-1} \\ &\quad \cdot \left(\begin{bmatrix} \mathbf{H}^T[n-1] & \mathbf{h}[n] \end{bmatrix} \begin{bmatrix} \mathbf{C}[n-1] & \mathbf{0} \\ \mathbf{0}^T & \sigma_n^2 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{x}[n-1] \\ x[n] \end{bmatrix} \right).\end{aligned}$$

Since the covariance matrix is diagonal, it is easily inverted to yield

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= \left(\mathbf{H}^T[n-1]\mathbf{C}^{-1}[n-1]\mathbf{H}[n-1] + \frac{1}{\sigma_n^2}\mathbf{h}[n]\mathbf{h}^T[n] \right)^{-1} \\ &\quad \cdot \left(\mathbf{H}^T[n-1]\mathbf{C}^{-1}[n-1]\mathbf{x}[n-1] + \frac{1}{\sigma_n^2}\mathbf{h}[n]x[n] \right).\end{aligned}$$

Let

$$\boldsymbol{\Sigma}[n-1] = (\mathbf{H}^T[n-1]\mathbf{C}^{-1}[n-1]\mathbf{H}[n-1])^{-1}$$

which is the covariance matrix of $\hat{\boldsymbol{\theta}}[n-1]$. Then,

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= \left(\boldsymbol{\Sigma}^{-1}[n-1] + \frac{1}{\sigma_n^2}\mathbf{h}[n]\mathbf{h}^T[n] \right)^{-1} \\ &\quad \cdot \left(\mathbf{H}^T[n-1]\mathbf{C}^{-1}[n-1]\mathbf{x}[n-1] + \frac{1}{\sigma_n^2}\mathbf{h}[n]x[n] \right).\end{aligned}$$

The first term in parentheses is just $\Sigma[n]$. We can use Woodbury's identity (see Appendix 1) to obtain

$$\begin{aligned}\Sigma[n] &= \left(\Sigma^{-1}[n-1] + \frac{1}{\sigma_n^2} \mathbf{h}[n] \mathbf{h}^T[n] \right)^{-1} \\ &= \Sigma[n-1] - \frac{\Sigma[n-1] \mathbf{h}[n] \mathbf{h}^T[n] \Sigma[n-1]}{\sigma_n^2 + \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n]} \\ &= (\mathbf{I} - \mathbf{K}[n] \mathbf{h}^T[n]) \Sigma[n-1]\end{aligned}$$

where

$$\mathbf{K}[n] = \frac{\Sigma[n-1] \mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n]}.$$

Applying these results to $\hat{\boldsymbol{\theta}}[n]$ yields

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= (\mathbf{I} - \mathbf{K}[n] \mathbf{h}^T[n]) \Sigma[n-1] \\ &\quad \cdot \left(\Sigma^{-1}[n-1] \hat{\boldsymbol{\theta}}[n-1] + \frac{1}{\sigma_n^2} \mathbf{h}[n] x[n] \right)\end{aligned}$$

since

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n-1] &= (\mathbf{H}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{H}[n-1])^{-1} \mathbf{H}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{x}[n-1] \\ &= \Sigma[n-1] \mathbf{H}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{x}[n-1].\end{aligned}$$

Continuing, we have

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= \hat{\boldsymbol{\theta}}[n-1] + \frac{1}{\sigma_n^2} \Sigma[n-1] \mathbf{h}[n] x[n] - \mathbf{K}[n] \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n-1] \\ &\quad - \frac{1}{\sigma_n^2} \mathbf{K}[n] \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n] x[n].\end{aligned}$$

But

$$\begin{aligned}&\frac{1}{\sigma_n^2} \Sigma[n-1] \mathbf{h}[n] - \frac{1}{\sigma_n^2} \mathbf{K}[n] \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n] \\ &= \frac{1}{\sigma_n^2} (\sigma_n^2 + \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n]) \mathbf{K}[n] \\ &\quad - \frac{1}{\sigma_n^2} \mathbf{K}[n] \mathbf{h}^T[n] \Sigma[n-1] \mathbf{h}[n] \\ &= \mathbf{K}[n]\end{aligned}$$

and therefore

$$\begin{aligned}\hat{\boldsymbol{\theta}}[n] &= \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n] x[n] - \mathbf{K}[n] \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n-1] \\ &= \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n] (x[n] - \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n-1]).\end{aligned}$$

Finally, to update the minimum LS error we have

$$\begin{aligned}
 J_{\min}[n] &= \left(\mathbf{x}[n] - \mathbf{H}[n]\hat{\boldsymbol{\theta}}[n] \right)^T \mathbf{C}^{-1}[n] \left(\mathbf{x}[n] - \mathbf{H}[n]\hat{\boldsymbol{\theta}}[n] \right) \\
 &= \mathbf{x}^T[n] \mathbf{C}^{-1}[n] \left(\mathbf{x}[n] - \mathbf{H}[n]\hat{\boldsymbol{\theta}}[n] \right) \\
 &= \begin{bmatrix} \mathbf{x}^T[n-1] & x[n] \end{bmatrix} \begin{bmatrix} \mathbf{C}^{-1}[n-1] & \mathbf{0} \\ \mathbf{0}^T & \frac{1}{\sigma_n^2} \end{bmatrix} \begin{bmatrix} \mathbf{x}[n-1] - \mathbf{H}[n-1]\hat{\boldsymbol{\theta}}[n] \\ x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n] \end{bmatrix} \\
 &= \mathbf{x}^T[n-1] \mathbf{C}^{-1}[n-1] \left(\mathbf{x}[n-1] - \mathbf{H}[n-1]\hat{\boldsymbol{\theta}}[n] \right) \\
 &\quad + \frac{1}{\sigma_n^2} x[n] \left(x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n] \right).
 \end{aligned}$$

Let $e[n] = x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n-1]$. Using the update for $\hat{\boldsymbol{\theta}}[n]$ produces

$$\begin{aligned}
 J_{\min}[n] &= \mathbf{x}^T[n-1] \mathbf{C}^{-1}[n-1] \left(\mathbf{x}[n-1] - \mathbf{H}[n-1]\hat{\boldsymbol{\theta}}[n-1] - \mathbf{H}[n-1]\mathbf{K}[n]e[n] \right) \\
 &\quad + \frac{1}{\sigma_n^2} x[n] \left(x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n-1] - \mathbf{h}^T[n]\mathbf{K}[n]e[n] \right) \\
 &= J_{\min}[n-1] - \mathbf{x}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{H}[n-1] \mathbf{K}[n] e[n] \\
 &\quad + \frac{1}{\sigma_n^2} x[n] \left(1 - \mathbf{h}^T[n] \mathbf{K}[n] \right) e[n].
 \end{aligned}$$

But

$$\mathbf{H}^T[n-1] \mathbf{C}^{-1}[n-1] \mathbf{x}[n-1] = \boldsymbol{\Sigma}^{-1}[n-1] \hat{\boldsymbol{\theta}}[n-1]$$

so that

$$\begin{aligned}
 J_{\min}[n] &= J_{\min}[n-1] - \hat{\boldsymbol{\theta}}^T[n-1] \boldsymbol{\Sigma}^{-1}[n-1] \mathbf{K}[n] e[n] \\
 &\quad + \frac{1}{\sigma_n^2} x[n] \left(1 - \mathbf{h}^T[n] \mathbf{K}[n] \right) e[n]
 \end{aligned}$$

and also

$$\begin{aligned}
 &\frac{1}{\sigma_n^2} x[n] \left(1 - \mathbf{h}^T[n] \mathbf{K}[n] \right) - \hat{\boldsymbol{\theta}}^T[n-1] \boldsymbol{\Sigma}^{-1}[n-1] \mathbf{K}[n] \\
 &= \frac{1}{\sigma_n^2} x[n] \left(1 - \frac{\mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]} \right) - \frac{\hat{\boldsymbol{\theta}}^T[n-1] \mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]} \\
 &= \frac{x[n] - \mathbf{h}^T[n] \hat{\boldsymbol{\theta}}[n-1]}{\sigma_n^2 + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]} \\
 &= \frac{e[n]}{\sigma_n^2 + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]}.
 \end{aligned}$$

Hence, we have finally

$$J_{\min}[n] = J_{\min}[n-1] + \frac{e^2[n]}{\sigma_n^2 + \mathbf{h}^T[n] \boldsymbol{\Sigma}[n-1] \mathbf{h}[n]}.$$

Chapter 9

Method of Moments

9.1 Introduction

The approach known as the *method of moments* is described in this chapter. It produces an estimator that is easy to determine and simple to implement. Although the estimator has no optimality properties, it is useful if the data record is long enough. This is because the method of moments estimator is usually consistent. If the performance is not satisfactory, then it can be used as an initial estimate, which is subsequently improved through a Newton-Raphson implementation of the MLE. After obtaining the method of moments estimator, we describe some approximate approaches for analyzing its statistical performance. The techniques are general enough to be used to evaluate the performance of many other estimators and are therefore useful in their own right.

9.2 Summary

The method of moments approach to estimation is illustrated in Section 9.3 with some examples. In general, the estimator is given by (9.11) for a vector parameter. The performance of the method of moments estimator may be partially characterized by the approximate mean of (9.15) and the approximate variance of (9.16). Also, any estimator that depends on a group of statistics whose PDF is concentrated about its mean can make use of the same expressions. Another case where an approximate mean and variance evaluation is useful involves signal in noise problems when the SNR is high. Then, the estimation performance can be partially described by the approximate mean of (9.18) and the approximate variance of (9.19).

9.3 Method of Moments

The method of moments approach to estimation is based on the solution of a theoretical equation involving the moments of a PDF. As an example, assume we observe $x[n]$ for $n = 0, 1, \dots, N - 1$, which are IID samples from the Gaussian mixture PDF (see also

Problem 6.14)

$$p(x[n]; \epsilon) = \frac{1-\epsilon}{\sqrt{2\pi\sigma_1^2}} \exp\left(-\frac{1}{2} \frac{x^2[n]}{\sigma_1^2}\right) + \frac{\epsilon}{\sqrt{2\pi\sigma_2^2}} \exp\left(-\frac{1}{2} \frac{x^2[n]}{\sigma_2^2}\right)$$

or in more succinct form

$$p(x[n]; \epsilon) = (1-\epsilon)\phi_1(x[n]) + \epsilon\phi_2(x[n])$$

where

$$\phi_i(x[n]) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{1}{2} \frac{x^2[n]}{\sigma_i^2}\right).$$

The parameter ϵ is termed the mixture parameter, which satisfies $0 < \epsilon < 1$, and σ_1^2, σ_2^2 are the variances of the individual Gaussian PDFs. The Gaussian mixture PDF may be thought of as the PDF of a random variable obtained from a $\mathcal{N}(0, \sigma_1^2)$ PDF with probability $1-\epsilon$ and from a $\mathcal{N}(0, \sigma_2^2)$ PDF with probability ϵ . Now if σ_1^2, σ_2^2 are known and ϵ is to be estimated, all our usual MVU estimation methods will fail. The MLE will require the maximization of a very nonlinear function of ϵ . Although this can be implemented using a grid search, the method of moments provides a much simpler estimator. Note that

$$\begin{aligned} E(x^2[n]) &= \int_{-\infty}^{\infty} x^2[n] [(1-\epsilon)\phi_1(x[n]) + \epsilon\phi_2(x[n])] dx[n] \\ &= (1-\epsilon)\sigma_1^2 + \epsilon\sigma_2^2 \end{aligned} \quad (9.1)$$

since the mean of $x[n]$ is zero. This theoretical equation relates the unknown parameter ϵ to the second moment. If we now replace $E(x^2[n])$ by its natural estimator $\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]$, we have

$$\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] = (1-\epsilon)\sigma_1^2 + \epsilon\sigma_2^2 \quad (9.2)$$

and solving for ϵ , we have our method of moments estimator

$$\hat{\epsilon} = \frac{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \sigma_1^2}{\sigma_2^2 - \sigma_1^2}. \quad (9.3)$$

The ease with which the estimator was found is due to the linear nature of the theoretical moment equation (9.1). The estimator is easily shown to be unbiased (although this is not generally the case). The variance of $\hat{\epsilon}$ is found as

$$\begin{aligned} \text{var}(\hat{\epsilon}) &= \frac{1}{(\sigma_2^2 - \sigma_1^2)^2} \text{var}\left(\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]\right) \\ &= \frac{1}{N(\sigma_2^2 - \sigma_1^2)^2} \text{var}(x^2[n]) \\ &= \frac{1}{N(\sigma_2^2 - \sigma_1^2)^2} [E(x^4[n]) - E^2(x^2[n])]. \end{aligned}$$

But it is easily shown that

$$E(x^4[n]) = (1 - \epsilon)3\sigma_1^4 + \epsilon 3\sigma_2^4$$

which when combined with (9.1) produces the variance

$$\text{var}(\hat{\epsilon}) = \frac{3(1 - \epsilon)\sigma_1^4 + 3\epsilon\sigma_2^4 - [(1 - \epsilon)\sigma_1^2 + \epsilon\sigma_2^2]^2}{N(\sigma_2^2 - \sigma_1^2)^2}. \quad (9.4)$$

To determine the loss in performance we could evaluate the CRLB (we would need to do this numerically) and compare it to (9.4). If the increase in variance were substantial, we could attempt to implement the MLE, which would attain the bound for large data records. It should be observed that the method of moments estimator is consistent in that $\hat{\epsilon} \rightarrow \epsilon$ as $N \rightarrow \infty$. This is because as $N \rightarrow \infty$, from (9.3) $E(\hat{\epsilon}) = \epsilon$ and from (9.4) $\text{var}(\hat{\epsilon}) \rightarrow 0$. Since the estimates of the moments that are substituted into the theoretical equation approach the true moments for large data records, the equation to be solved approaches the theoretical equation. As a result, the method of moments estimator will in general be consistent (see Problem 7.5). In (9.2) for example, as $N \rightarrow \infty$, we have

$$\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] \rightarrow E(x^2[n])$$

and the theoretical equation results. When solved for ϵ , the true value is obtained.

We now summarize the method of moments for a scalar parameter. Assume that the k th moment $\mu_k = E(x^k[n])$ depends upon the unknown parameter θ according to

$$\mu_k = h(\theta). \quad (9.5)$$

We first solve for θ as

$$\theta = h^{-1}(\mu_k)$$

assuming h^{-1} exists. We then replace the theoretical moment by its natural estimator

$$\hat{\mu}_k = \frac{1}{N} \sum_{n=0}^{N-1} x^k[n] \quad (9.6)$$

to yield the method of moments estimator

$$\hat{\theta} = h^{-1} \left(\frac{1}{N} \sum_{n=0}^{N-1} x^k[n] \right). \quad (9.7)$$

Some examples follow.

Example 9.1 - DC Level in WGN

If $x[n] = A + w[n]$ is observed for $n = 0, 1, \dots, N-1$, where $w[n]$ is WGN with variance σ^2 and A is to be estimated, then we know that

$$\mu_1 = E(x[n]) = A.$$

This is the theoretical equation of (9.5). According to (9.7), we replace μ_1 by its natural estimator, resulting in

$$\hat{A} = \frac{1}{N} \sum_{n=0}^{N-1} x[n].$$

In this case h is the identity transformation or $h(x) = x$. ◇

Example 9.2 - Exponential PDF

Consider N IID observations from the exponential PDF

$$p(x[n]; \lambda) = \begin{cases} \lambda \exp(-\lambda x[n]) & x[n] > 0 \\ 0 & x[n] < 0. \end{cases}$$

We wish to estimate the parameter λ , where $\lambda > 0$. The first moment is)

$$\begin{aligned} \mu_1 = E(x[n]) &= \int_0^{\infty} x[n] \lambda \exp(-\lambda x[n]) dx[n] \\ &= \frac{1}{\lambda} \int_0^{\infty} \xi \exp(-\xi) d\xi \\ &= \frac{1}{\lambda}. \end{aligned}$$

Solving for λ

$$\lambda = \frac{1}{\mu_1}$$

and substituting the natural estimator result in the method of moments estimator

$$\hat{\lambda} = \frac{1}{\frac{1}{N} \sum_{n=0}^{N-1} x[n]}.$$

◇

9.4 Extension to a Vector Parameter

Now consider a vector parameter θ of dimension $p \times 1$. It is obvious that to solve for θ requires p theoretical moment equations. Hence, we suppose

$$\begin{aligned} \mu_1 &= h_1(\theta_1, \theta_2, \dots, \theta_p) \\ \mu_2 &= h_2(\theta_1, \theta_2, \dots, \theta_p) \\ &\vdots \\ \mu_p &= h_p(\theta_1, \theta_2, \dots, \theta_p) \end{aligned} \tag{9.8}$$

or in matrix form

$$\boldsymbol{\mu} = \mathbf{h}(\boldsymbol{\theta}). \quad (9.9)$$

We then solve for $\boldsymbol{\theta}$ as

$$\boldsymbol{\theta} = \mathbf{h}^{-1}(\boldsymbol{\mu}) \quad (9.10)$$

and determine the method of moments estimator as

$$\hat{\boldsymbol{\theta}} = \mathbf{h}^{-1}(\hat{\boldsymbol{\mu}}) \quad (9.11)$$

where

$$\hat{\boldsymbol{\mu}} = \begin{bmatrix} \frac{1}{N} \sum_{n=0}^{N-1} x[n] \\ \frac{1}{N} \sum_{n=0}^{N-1} x^2[n] \\ \vdots \\ \frac{1}{N} \sum_{n=0}^{N-1} x^p[n] \end{bmatrix}.$$

It may occur that the first p moments are insufficient to determine all the parameters to be estimated (see Example 9.3). In this case we need to find some set of p moment equations allowing (9.9) to be solved for $\boldsymbol{\theta}$ to obtain (9.10). In practice, it is desirable to use the lowest order moments possible. This is because the variance of the moment estimator generally increases with order. Also, to be able to solve the resulting equations we would like them to be *linear* or at least mildly nonlinear. Otherwise, a nonlinear optimization may be needed, defeating the original motivation for the method of moments estimator as an easily implemented estimator. In the vector parameter case we may also need cross-moments, as the signal processing example in Section 9.6 illustrates (see also Problem 9.5). We now continue the Gaussian mixture example.

Example 9.3 - Gaussian Mixture PDF

Returning to the introductory example of a Gaussian mixture PDF, assume now that in addition to ϵ , the Gaussian variances σ_1^2 and σ_2^2 are unknown as well. To estimate all three parameters we require three moment equations. Noting that the PDF is an even function so that all odd order moments are zero, we utilize

$$\begin{aligned} \mu_2 &= E(x^2[n]) = (1 - \epsilon)\sigma_1^2 + \epsilon\sigma_2^2 \\ \mu_4 &= E(x^4[n]) = 3(1 - \epsilon)\sigma_1^4 + 3\epsilon\sigma_2^4 \\ \mu_6 &= E(x^6[n]) = 15(1 - \epsilon)\sigma_1^6 + 15\epsilon\sigma_2^6. \end{aligned}$$

Although nonlinear, these equations may be solved by letting [Rider 1961]

$$\begin{aligned} u &= \sigma_1^2 + \sigma_2^2 \\ v &= \sigma_1^2\sigma_2^2. \end{aligned} \quad (9.12)$$

Then it can be shown by direct substitution that

$$\begin{aligned}u &= \frac{\mu_6 - 5\mu_4\mu_2}{5\mu_4 - 15\mu_2^2} \\v &= \mu_2 u - \frac{\mu_4}{3}.\end{aligned}$$

Once u is found, v can be determined. Then, σ_1^2 and σ_2^2 are obtained by solving (9.12) to yield

$$\begin{aligned}\sigma_1^2 &= \frac{u + \sqrt{u^2 - 4v}}{2} \\ \sigma_2^2 &= \frac{v}{\sigma_1^2}\end{aligned}$$

and finally, the mixture parameter becomes (see (9.3))

$$\epsilon = \frac{\mu_2 - \sigma_1^2}{\sigma_2^2 - \sigma_1^2}.$$

The method of moments performs the same operations but with the theoretical moments $\{\mu_2, \mu_4, \mu_6\}$ replaced by their natural estimators.

◇

9.5 Statistical Evaluation of Estimators

In the method of moments we do not know beforehand whether the estimator will perform well. Once the estimator has been obtained we may be able, if we are lucky, to determine the statistical properties of mean and variance as in the introductory Gaussian mixture example. Since

$$\hat{\boldsymbol{\theta}} = \mathbf{h}^{-1}(\hat{\boldsymbol{\mu}}) = \mathbf{g}(\mathbf{x}) \quad (9.13)$$

we can in principle determine the PDF of $\hat{\boldsymbol{\theta}}$ using standard formulas for the transformation of random variables. In practice, this is usually impossible due to its mathematical intractability. Because (9.13) is the general form for any estimator, the methods to be described apply to other estimators in addition to the method of moments estimator. They allow us to assess the performance of an estimator by determining *approximate* expressions for the mean and variance. The approximation generally improves as the data record length increases, being an asymptotic approach. For determination of the *exact* mean and variance a Monte Carlo computer simulation as discussed in Chapter 7 must be employed.

Consider a scalar parameter that has been estimated using (9.13). To determine the approximate mean and variance of $\hat{\theta}$ we must assume that it depends upon the $r < N$ statistics $\{\mathbf{T}_1(\mathbf{x}), \mathbf{T}_2(\mathbf{x}), \dots, \mathbf{T}_r(\mathbf{x})\}$, whose variances and covariances are small. With

this latter assumption the PDF of $[T_1 T_2 \dots T_r]^T$ will be concentrated about its mean. In Example 9.2, for instance, the estimator of λ can be written as

$$\hat{\lambda} = g(T_1(\mathbf{x}))$$

where $T_1(\mathbf{x}) = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$ and $g(T_1) = 1/T_1$. For large N the PDF of T_1 will be heavily concentrated about its mean since $\text{var}(T_1) = 1/(N\lambda^2)$, as will be shown in Example 9.4. Using the statistical linearization argument of Chapter 3, we can use a first-order Taylor expansion of g about the mean of T_1 . In general, we assume that

$$\hat{\theta} = g(\mathbf{T})$$

where $\mathbf{T} = [T_1 T_2 \dots T_r]^T$. We then perform a first-order Taylor expansion of g about the point $\mathbf{T} = E(\mathbf{T}) = \boldsymbol{\mu}$ to yield

$$\hat{\theta} = g(\mathbf{T}) \approx g(\boldsymbol{\mu}) + \sum_{k=1}^r \left. \frac{\partial g}{\partial T_k} \right|_{\mathbf{T}=\boldsymbol{\mu}} (T_k - \mu_k). \quad (9.14)$$

Assuming this to hold, the mean becomes

$$E(\hat{\theta}) = g(\boldsymbol{\mu}). \quad (9.15)$$

To this degree of approximation $E(\hat{\theta}) = E(g(\mathbf{T})) = g(E(\mathbf{T}))$ or the expectation commutes over the nonlinear function g . Note that we require the first-order moment of $\mathbf{T}$ to determine the mean. To determine the variance we again use (9.14) to obtain

$$\text{var}(\hat{\theta}) = E \left\{ \left[g(\boldsymbol{\mu}) + \left. \frac{\partial g}{\partial \mathbf{T}} \right|_{\mathbf{T}=\boldsymbol{\mu}} (\mathbf{T} - \boldsymbol{\mu}) - E(\hat{\theta}) \right]^2 \right\}.$$

But from (9.15) we have that $E(\hat{\theta}) = g(\boldsymbol{\mu})$. Therefore,

$$\text{var}(\hat{\theta}) = \left. \frac{\partial g}{\partial \mathbf{T}} \right|_{\mathbf{T}=\boldsymbol{\mu}}^T \mathbf{C}_T \left. \frac{\partial g}{\partial \mathbf{T}} \right|_{\mathbf{T}=\boldsymbol{\mu}} \quad (9.16)$$

where $\mathbf{C}_T$ is the covariance matrix of $\mathbf{T}$. Now we require both the mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{C}_T$ to determine the variance. Similar derivations can be used to extend the Taylor expansion to higher order terms, although the algebra becomes increasingly tedious (see Problems 9.8 and 9.9). An example follows.

Example 9.4 - Exponential PDF (continued)

Continuing Example 9.2, suppose we have the method of moments estimator

$$\hat{\lambda} = \frac{1}{\frac{1}{N} \sum_{n=0}^{N-1} x[n]}$$

where the $x[n]$'s are IID and each has an exponential distribution. To find the approximate mean and variance we use (9.15) and (9.16). In this case we have

$$\hat{\lambda} = g(T_1)$$

where

$$T_1 = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$$

and

$$g(T_1) = \frac{1}{T_1}.$$

The mean of T_1 is

$$\begin{aligned} \mu_1 = E(T_1) &= \frac{1}{N} \sum_{n=0}^{N-1} E(x[n]) \\ &= E(x[n]) \\ &= \frac{1}{\lambda} \end{aligned}$$

using the results of Example 9.2. The variance of T_1 is

$$\begin{aligned} \text{var}(T_1) &= \text{var}\left(\frac{1}{N} \sum_{n=0}^{N-1} x[n]\right) \\ &= \frac{\text{var}(x[n])}{N}. \end{aligned}$$

But

$$\begin{aligned} \text{var}(x[n]) &= \int_0^{\infty} x^2[n] \lambda \exp(-\lambda x[n]) dx[n] - \frac{1}{\lambda^2} \\ &= \frac{2}{\lambda^2} - \frac{1}{\lambda^2} = \frac{1}{\lambda^2} \end{aligned}$$

so that

$$\text{var}(T_1) = \frac{1}{N\lambda^2}.$$

Since $g(T_1) = 1/T_1$,

$$\left. \frac{\partial g}{\partial T_1} \right|_{T_1=\mu_1} = -\frac{1}{\mu_1^2} = -\lambda^2.$$

From (9.15) we have the approximate mean

$$E(\hat{\lambda}) = g(\mu_1) = \frac{1}{\frac{1}{\lambda}} = \lambda$$

and from (9.16) the approximate variance is

$$\begin{aligned}\text{var}(\hat{\lambda}) &= \left. \frac{\partial g}{\partial T_1} \right|_{T_1=\mu_1} \text{var}(T_1) \left. \frac{\partial g}{\partial T_1} \right|_{T_1=\mu_1} \\ &= (-\lambda^2) \frac{1}{N\lambda^2} (-\lambda^2) \\ &= \frac{\lambda^2}{N}.\end{aligned}$$

The estimator is seen to be approximately unbiased and has an approximate variance decreasing with N . It should be emphasized that these expressions are only approximate. They are only as accurate as is the linearization of g . In fact, for this problem $\hat{\lambda}$ can be easily shown to be the MLE. As such, its asymptotic PDF (as $N \rightarrow \infty$) can be shown to be (see (7.8))

$$\hat{\lambda} \stackrel{a}{\sim} \mathcal{N}(\lambda, \lambda^2/N).$$

Hence, the approximation is accurate only as $N \rightarrow \infty$. (In Problem 9.10 *quadratic* expansions of g are employed to derive the *second-order* approximations to the mean and variance of $\hat{\lambda}$.) To determine how large N must be for the mean and variance expressions to hold requires a Monte Carlo computer simulation. $\diamond$

The basic premise behind the Taylor series approach is that the function g be approximately linear over the range of $\mathbf{T}$ for which $p(\mathbf{T}; \theta)$ is essentially nonzero. This occurs naturally when

1. The data record is large so that $p(\mathbf{T}; \theta)$ is concentrated about its mean, as in the previous example. There the PDF of $T_1 = \frac{1}{N} \sum_{n=0}^{N-1} x[n]$ becomes more concentrated about $E(T_1)$ as $N \rightarrow \infty$ since $\text{var}(T_1) = 1/(N\lambda^2) \rightarrow 0$ as $N \rightarrow \infty$.
2. The problem is to estimate the parameter of a signal in noise, and the SNR is high. In this case, by expanding the function about the signal we can obtain the approximate mean and variance. For a high SNR the results will be quite accurate. This is because at a high SNR the noise causes only a slight perturbation of the estimator value obtained in the case of no noise. We now develop this second case.

We consider the data

$$x[n] = s[n; \theta] + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is zero mean noise with covariance matrix $\mathbf{C}$. A general estimator of the scalar parameter θ is

$$\begin{aligned}\hat{\theta} &= g(\mathbf{x}) \\ &= g(s(\theta) + \mathbf{w}) \\ &= h(\mathbf{w}).\end{aligned}$$

In this case we may choose the statistic $\mathbf{T}$ as the original data since as the SNR becomes larger, the PDF of $\mathbf{T} = \mathbf{x}$ becomes more concentrated about its mean, which is the signal. We now use a first-order Taylor expansion of g about the mean of $\mathbf{x}$, which is $\boldsymbol{\mu} = \mathbf{s}(\theta)$, or, equivalently, of $\mathbf{h}$ about $\mathbf{w} = \mathbf{0}$. This results in

$$\hat{\theta} \approx h(\mathbf{0}) + \sum_{n=0}^{N-1} \left. \frac{\partial h}{\partial w[n]} \right|_{\mathbf{w}=\mathbf{0}} w[n]. \quad (9.17)$$

As before, it follows that approximately

$$E(\hat{\theta}) = h(\mathbf{0}) = g(s(\theta)) \quad (9.18)$$

$$\text{var}(\hat{\theta}) = \left. \frac{\partial h}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{0}}^T \mathbf{C} \left. \frac{\partial h}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{0}}. \quad (9.19)$$

In most cases if there is no noise, the estimator yields the true value since $g(s(\hat{\theta})) = \theta$. Hence, for a high SNR it will be unbiased. An example follows.

Example 9.5 - Exponential Signal in White Noise

For the data

$$x[n] = r^n + w[n] \quad n = 0, 1, 2$$

where $w[n]$ is zero mean uncorrelated noise with variance σ^2 , the damping factor r is to be estimated. Wishing to avoid a maximization of the likelihood to find the MLE (see Example 7.11), we propose the estimator

$$\hat{r} = \frac{x[2] + x[1]}{x[1] + x[0]}.$$

In terms of the signal and noise we have

$$\hat{r} = h(\mathbf{w}) = \frac{r^2 + w[2] + r + w[1]}{r + w[1] + 1 + w[0]}$$

and according to (9.18)

$$E(\hat{r}) = h(\mathbf{0}) = \frac{r^2 + r}{r + 1} = r$$

or the estimator is approximately unbiased. To find the variance

$$\begin{aligned} \left. \frac{\partial h}{\partial w[0]} \right|_{\mathbf{w}=\mathbf{0}} &= - \left. \frac{r^2 + w[2] + r + w[1]}{(r + w[1] + 1 + w[0])^2} \right|_{\mathbf{w}=\mathbf{0}} \\ &= - \frac{r^2 + r}{(r + 1)^2} \\ &= - \frac{r}{r + 1}. \end{aligned}$$

Similarly,

$$\begin{aligned}\left.\frac{\partial h}{\partial w[1]}\right|_{\mathbf{w}=\mathbf{0}} &= -\frac{r-1}{r+1} \\ \left.\frac{\partial h}{\partial w[2]}\right|_{\mathbf{w}=\mathbf{0}} &= \frac{1}{r+1}.\end{aligned}$$

Since $\mathbf{C} = \sigma^2 \mathbf{I}$, we have from (9.19)

$$\begin{aligned}\text{var}(\hat{r}) &= \sigma^2 \sum_{n=0}^2 \left(\left. \frac{\partial h}{\partial w[n]} \right|_{\mathbf{w}=\mathbf{0}} \right)^2 \\ &= \frac{\sigma^2}{(r+1)^2} [r^2 + (r-1)^2 + 1] \\ &= 2\sigma^2 \frac{r^2 - r + 1}{(r+1)^2}.\end{aligned}$$

◇

It is also worthwhile to remark that from (9.14) and (9.17) the estimators are approximately linear functions of $\mathbf{T}$ and $\mathbf{w}$, respectively. If these are Gaussian, then $\hat{\theta}$ will also be Gaussian — at least to within the approximation assumed by the Taylor expansion. Finally, we should note that for a *vector parameter* the approximate mean and variance for each component can be obtained by applying these techniques. Obtaining the *covariance matrix* of $\hat{\theta}$ is possible using a first-order Taylor expansion but can be extremely tedious.

9.6 Signal Processing Example

We now apply the method of moments and the approximate performance analysis to the problem of frequency estimation. Assume that we observe

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is zero mean white noise with variance σ^2 . The frequency f_0 is to be estimated. This problem was discussed in Example 7.16 in which the MLE of frequency was shown to be approximately given by the peak location of a periodogram. In an effort to reduce the computation involved in searching for the peak location, we now describe a method of moments estimator. To do so we depart slightly from the usual sinusoidal model to assume that the phase ϕ is a random variable independent of $w[n]$ and distributed as $\phi \sim \mathcal{U}[0, 2\pi]$. With this assumption the signal $s[n] = A \cos(2\pi f_0 n + \phi)$ can be viewed as the *realization* of a WSS random process. That this is the case is verified by determining the mean and ACF of $s[n]$. The mean is

$$\begin{aligned}
E(s[n]) &= E[A \cos(2\pi f_0 n + \phi)] \\
&= \int_0^{2\pi} A \cos(2\pi f_0 n + \phi) \frac{1}{2\pi} d\phi \\
&= 0
\end{aligned}$$

and the ACF is

$$\begin{aligned}
r_{ss}[k] &= E[s[n]s[n+k]] \\
&= E[A^2 \cos(2\pi f_0 n + \phi) \cos(2\pi f_0(n+k) + \phi)] \\
&= A^2 E \left[\frac{1}{2} \cos(4\pi f_0 n + 2\pi f_0 k + 2\phi) + \frac{1}{2} \cos 2\pi f_0 k \right] \\
&= \frac{A^2}{2} \cos 2\pi f_0 k.
\end{aligned}$$

The ACF of the observed process becomes

$$\begin{aligned}
r_{xx}[k] &= r_{ss}[k] + r_{ww}[k] \\
&= \frac{A^2}{2} \cos 2\pi f_0 k + \sigma^2 \delta[k].
\end{aligned}$$

To simplify the discussion and the subsequent estimator we will assume that the signal amplitude is known (see Problem 9.12 for the unknown amplitude case). We will let $A = \sqrt{2}$, so that the ACF becomes

$$r_{xx}[k] = \cos 2\pi f_0 k + \sigma^2 \delta[k].$$

To estimate the frequency by the method of moments approach we observe that

$$r_{xx}[1] = \cos 2\pi f_0,$$

and therefore without having to assume knowledge of σ^2 we can implement the method of moments estimator

$$\hat{f}_0 = \frac{1}{2\pi} \arccos \hat{r}_{xx}[1]$$

where $\hat{r}_{xx}[1]$ is an estimator of $r_{xx}[1]$. A reasonable estimator of the ACF for $k = 1$ is

$$\hat{r}_{xx}[1] = \frac{1}{N-1} \sum_{n=0}^{N-2} x[n]x[n+1]$$

so that finally we have for our frequency estimator

$$\hat{f}_0 = \frac{1}{2\pi} \arccos \left[\frac{1}{N-1} \sum_{n=0}^{N-2} x[n]x[n+1] \right]. \quad (9.20)$$

The argument of the arccos function may exceed 1 in magnitude, usually occurring when the SNR is low. If this is the case, the estimate is meaningless and it is indicative of an unreliable estimator.

Although motivated by the random phase sinusoidal model, this estimator remains valid if ϕ is deterministic and unknown (see Problem 9.11). In determining the mean and variance of $\hat{f}_0$ we assume that ϕ is deterministic. To assess the performance of our estimator we use the first-order Taylor expansion approach and expand about the point $\mathbf{w} = \mathbf{0}$. Thus, the obtained performance is valid for high enough SNR. From (9.18) we have

$$\begin{aligned} E(\hat{f}_0) &= \frac{1}{2\pi} \arccos \left[\frac{1}{N-1} \sum_{n=0}^{N-2} \sqrt{2} \cos(2\pi f_0 n + \phi) \sqrt{2} \cos(2\pi f_0(n+1) + \phi) \right] \\ &= \frac{1}{2\pi} \arccos \left[\frac{1}{N-1} \sum_{n=0}^{N-2} (\cos(4\pi f_0 n + 2\pi f_0 + 2\phi) + \cos 2\pi f_0) \right]. \end{aligned}$$

The double-frequency term when summed is approximately zero, so that at a high SNR

$$\begin{aligned} E(\hat{f}_0) &= \frac{1}{2\pi} \arccos [\cos 2\pi f_0] \\ &= f_0. \end{aligned} \tag{9.21}$$

To find the variance we note that

$$\hat{f}_0 = h(\mathbf{w}) = \frac{1}{2\pi} \arccos \left[\underbrace{\frac{1}{N-1} \sum_{n=0}^{N-2} (s[n] + w[n])(s[n+1] + w[n+1])}_u \right]$$

so that the first-order partials are

$$\left. \frac{\partial h}{\partial w[i]} \right|_{\mathbf{w}=\mathbf{0}} = -\frac{1}{2\pi} \frac{1}{\sqrt{1-u^2}} \left. \frac{\partial u}{\partial w[i]} \right|_{\mathbf{w}=\mathbf{0}}.$$

But

$$\left. \frac{\partial u}{\partial w[i]} \right|_{\mathbf{w}=\mathbf{0}} = \frac{1}{N-1} \frac{\partial}{\partial w[i]} \sum_{n=0}^{N-2} (s[n]w[n+1] + w[n]s[n+1]) \bigg|_{\mathbf{w}=\mathbf{0}}$$

with the other terms in the sum resulting in zero after differentiating and letting $\mathbf{w} = \mathbf{0}$. Continuing, we have

$$\left. \frac{\partial u}{\partial w[i]} \right|_{\mathbf{w}=\mathbf{0}} = \begin{cases} \frac{1}{N-1} s[1] & i = 0 \\ \frac{1}{N-1} (s[i-1] + s[i+1]) & i = 1, 2, \dots, N-2 \\ \frac{1}{N-1} s[N-2] & i = N-1 \end{cases}$$

Also,

$$\begin{aligned} u|_{\mathbf{w}=0} &= \frac{1}{N-1} \sum_{n=0}^{N-2} s[n]s[n+1] \\ &\approx \cos 2\pi f_0 \end{aligned}$$

as already shown. Thus,

$$\left. \frac{\partial h}{\partial w[i]} \right|_{\mathbf{w}=0} = -\frac{1}{2\pi \sin 2\pi f_0} \left. \frac{\partial u}{\partial w[i]} \right|_{\mathbf{w}=0}.$$

From (9.19) with $\mathbf{C} = \sigma^2 \mathbf{I}$ we obtain

$$\begin{aligned} \text{var}(\hat{f}_0) &= \sigma^2 \sum_{n=0}^{N-1} \left(\left. \frac{\partial h}{\partial w[n]} \right|_{\mathbf{w}=0} \right)^2 \\ &= \frac{\sigma^2}{(2\pi)^2 (N-1)^2 \sin^2 2\pi f_0} \left[s^2[1] \right. \\ &\quad \left. + \sum_{n=1}^{N-2} (s[n-1] + s[n+1])^2 + s^2[N-2] \right]. \end{aligned}$$

But $s[n-1] + s[n+1] = 2 \cos 2\pi f_0 s[n]$, as can easily be verified, so that finally the variance of the frequency estimator is for a high SNR

$$\text{var}(\hat{f}_0) = \frac{\sigma^2}{(2\pi)^2 (N-1)^2 \sin^2 2\pi f_0} \left[s^2[1] + 4 \cos^2 2\pi f_0 \sum_{n=1}^{N-2} s^2[n] + s^2[N-2] \right] \quad (9.22)$$

where $s[n] = \sqrt{2} \cos(2\pi f_0 n + \phi)$. The evaluation of the variance is straightforward, although tedious. Note that the variance decreases most rapidly as $1/N^2$ (for $f_0 = 1/4$) versus the $1/N^3$ dependence for the CRLB (see (3.41)). Additionally, for f_0 near 0 or $1/2$ the variance increases rapidly. This is due to the arccos operation required to determine $\hat{f}_0$. As shown in Figure 9.1, a slight perturbation in the argument for x near ± 1 (f_0 near 0 or $1/2$) will cause large fluctuations in the function value due to the steep slope.

As an example of the accuracy of the approximate mean and variance expressions, we consider the data set with parameters $\text{SNR} = A^2/2\sigma^2 = 1/\sigma^2$ for $A = \sqrt{2}$ and $f_0 = 0.2$, $\phi = 0$, and $N = 50$. For $w[n]$ being WGN a Monte Carlo computer simulation in which 1000 realizations of $\hat{f}_0$ as given by (9.20) were generated for each of several SNRs. In Figure 9.2a the actual mean $E(\hat{f}_0)$ obtained, as well as the theoretical mean of $f_0 = 0.2$, are shown versus SNR. As expected for a high enough SNR, the estimator is essentially unbiased. The slight discrepancy is due to the approximations made in deriving the estimator, principally that N was large enough to neglect the double-frequency term. For larger N it can be verified that the mean equals the true value






Figure 9.1 Explanation for increased variance of frequency estimator



(a)



(b)

Figure 9.2 Method of moments frequency estimator

for a high SNR. Correspondingly, in Figure 9.2b the actual variance obtained and the theoretical variance as determined from (9.22) are shown versus SNR. The same behavior is evident. Some further discussions of this example can be found in Problem 9.13. For further discussions of this estimation approach the reader should consult [Lank et al. 1973, Kay 1989].

References

- Kay, S., "A Fast and Accurate Single Frequency Estimator," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 37, pp. 1987-1990, Dec. 1989.
- Lank, G.W., I.S. Reed, G.E. Pollon, "A Semicoherent Detection and Doppler Estimation Statistic," *IEEE Trans. Aerosp. Electron. Systems*, Vol. 9, pp. 151-165, March 1973.
- Rider, P.R., "Estimating the Parameters of Mixed Poisson, Binomial, and Weibull Distributions by the Method of Moments," *Bull. Int. Statist. Inst.*, Vol. 38, pp. 1-8, 1961.

Problems

- 9.1 If N IID observations $\{x[0], x[1], \dots, x[N-1]\}$ are made from the Rayleigh PDF

$$p(x; \sigma^2) = \begin{cases} \frac{x}{\sigma^2} \exp\left(-\frac{1}{2} \frac{x^2}{\sigma^2}\right) & x > 0 \\ 0 & x < 0 \end{cases}$$

find a method of moments estimator for σ^2 .

- 9.2 If N IID observations $\{x[0], x[1], \dots, x[N-1]\}$ are made from the Laplacian PDF

$$p(x; \sigma) = \frac{1}{\sqrt{2}\sigma} \exp\left(-\frac{\sqrt{2}|x|}{\sigma}\right)$$

find a method of moments estimator for σ .

- 9.3 Assume that N IID samples from a bivariate Gaussian PDF are observed or $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{N-1}\}$, where each $\mathbf{x}$ is a 2×1 random vector with PDF $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$. If

$$\mathbf{C} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

find a method of moments estimator for ρ . Also, determine a cubic equation to be solved for the MLE of ρ . Comment on the ease of implementation of the different estimators.

- 9.4 If N IID observations $\{x[0], x[1], \dots, x[N-1]\}$ are made from the $\mathcal{N}(\mu, \sigma^2)$ PDF, find a method of moments estimator for $\boldsymbol{\theta} = [\mu \sigma^2]^T$.

- 9.5 In Example 8.12 we determined that the ACF of an ARMA process satisfies the recursive difference equation

$$r_{xx}[n] = - \sum_{k=1}^p a[k] r_{xx}[n-k] \quad n > q$$

where $\{a[1], a[2], \dots, a[p]\}$ denote the AR filter parameters and q is the MA order. Propose a method of moments estimator for the AR filter parameters.

- 9.6 For a DC level in WGN or

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 , the parameter A^2 is to be estimated. It is proposed to use

$$\widehat{A^2} = (\bar{x})^2.$$

For this estimator find the approximate mean and variance using a first-order Taylor expansion approach.

- 9.7 For the observed data

$$x[n] = \cos \phi + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ^2 , find a method of moments estimator for ϕ . Assuming the SNR is high, determine the approximate mean and variance for your estimator using the first-order Taylor expansion approach.

- 9.8 In this problem we determine the mean of an estimator to *second order* by using a second-order Taylor expansion. To do so expand $\hat{\theta} = g(\mathbf{T})$ about the mean of $\mathbf{T}$ or $\mathbf{T} = \boldsymbol{\mu}$, retaining the terms up to and including the quadratic term. Show that the approximate mean of $\hat{\theta}$ is

$$E(\hat{\theta}) = g(\boldsymbol{\mu}) + \frac{1}{2} \text{tr}[\mathbf{G}(\boldsymbol{\mu}) \mathbf{C}_T]$$

where

$$[\mathbf{G}(\boldsymbol{\mu})]_{ij} = \left. \frac{\partial^2 g}{\partial T_i \partial T_j} \right|_{\mathbf{T}=\boldsymbol{\mu}}.$$

Hint: Use the result

$$E(\mathbf{x}^T \mathbf{y}) = E(\text{tr}(\mathbf{y} \mathbf{x}^T)) = \text{tr}(E(\mathbf{y} \mathbf{x}^T)).$$

- 9.9 In this problem we determine the approximate variance of an estimator to *second order* or by using a second-order Taylor expansion approach. To do so first expand $\hat{\theta} = g(\mathbf{T})$ about the mean of $\mathbf{T}$ or $\mathbf{T} = \boldsymbol{\mu}$, retaining the terms up to and including

the quadratic term. Next, use the results of Problem 9.8 and finally make the assumption that $\mathbf{T} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}_T)$. You should be able to verify the following result:

$$\text{var}(\hat{\theta}) = \left. \frac{\partial g}{\partial \mathbf{T}} \right|_{\mathbf{T}=\boldsymbol{\mu}}^T \mathbf{C}_T \left. \frac{\partial g}{\partial \mathbf{T}} \right|_{\mathbf{T}=\boldsymbol{\mu}} + \frac{1}{2} \text{tr} [(\mathbf{G}(\boldsymbol{\mu}) \mathbf{C}_T)^2].$$

The $r \times r$ $\mathbf{G}(\boldsymbol{\mu})$ is defined in Problem 9.8. Hint: You will need the following result, which is valid for a symmetric $N \times N$ matrix $\mathbf{A}$,

$$\text{var}(\mathbf{x}^T \mathbf{A} \mathbf{x}) = 2 \text{tr} [(\mathbf{A} \mathbf{C}_x)^2]$$

if $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_x)$.

- 9.10** Using the results of Problems 9.8 and 9.9 find the approximate mean and variance to second order for the estimator of λ discussed in Example 9.4. How do your results compare to the first-order result? Be sure to justify the approximate Gaussian PDF of $\bar{x}$, required to apply the variance expression. Also, compare the results to those predicted from asymptotic MLE theory (recall that the estimator is also the MLE).

- 9.11** Prove that for a sinusoidal signal

$$s[n] = A \cos(2\pi f_0 n + \phi)$$

where the phase ϕ is deterministic but unknown, that

$$\frac{1}{N-1} \sum_{n=0}^{N-2} s[n]s[n+1] \rightarrow \frac{A^2}{2} \cos 2\pi f_0$$

as $N \rightarrow \infty$. Assume that f_0 is not near 0 or 1/2. Hence, comment on the use of the method of moments estimator proposed in the signal processing example in Section 9.6 for the deterministic phase sinusoid.

- 9.12** To extend the applicability of the frequency estimator proposed in the signal processing example in Section 9.6 now assume that the signal amplitude A is unknown. Why can't the proposed estimator (9.20) be used in this case? Consider the estimator

$$\hat{f}_0 = \frac{1}{2\pi} \arccos \left[\frac{\frac{1}{N-1} \sum_{n=0}^{N-2} x[n]x[n+1]}{\frac{1}{N} \sum_{n=0}^{N-1} x^2[n]} \right].$$

Show that at a high SNR and for large N , $E(\hat{f}_0) = f_0$. Justify this estimator as a method of moments estimator based on the ACF. Would you expect this estimator to work well at a lower SNR?

- 9.13** For the signal processing example in Section 9.6 assume that the parameters are $A = \sqrt{2}$, $f_0 = 0.25$, $\phi = 0$, and that N is odd. Show that from (9.22) $\text{var}(\hat{f}_0) = 0$. To understand this result note that

$$\left. \frac{\partial h}{\partial w[i]} \right|_{\mathbf{w}=\mathbf{0}} = 0$$

for $i = 0, 1, \dots, N-1$. What does this say about the first-order Taylor expansion approximation? How could you improve this situation?

Chapter 10

The Bayesian Philosophy

10.1 Introduction

We now depart from the *classical approach* to statistical estimation in which the parameter θ of interest is assumed to be a *deterministic but unknown* constant. Instead, we assume that θ is a *random variable* whose particular *realization* we must estimate. This is the *Bayesian approach*, so named because its implementation is based directly on Bayes' theorem. The motivation for doing so is twofold. First, if we have available some prior knowledge about θ , we can incorporate it into our estimator. The mechanism for doing this requires us to assume that θ is a random variable with a given *prior* PDF. Classical estimation, on the other hand, finds it difficult to make use of any prior knowledge. The Bayesian approach, when applicable, can therefore improve the estimation accuracy. Second, Bayesian estimation is useful in situations where an MVU estimator cannot be found, as for example, when the variance of an unbiased estimator may not be uniformly less than that of all other unbiased estimators. In this instance, it may be true that for *most* values of the parameter an estimator can be found whose mean square error may be less than that of all other estimators. By assigning a PDF to θ we can devise strategies to find that estimator. The resultant estimator can then be said to be optimal "on the average," or with respect to the assumed prior PDF of θ . In this chapter we attempt to motivate the Bayesian approach and to discuss some of the issues surrounding its use. The reader should be aware that this approach to estimation has had a long and controversial history. For a more definitive account [Box and Tiao 1973] is recommended.

10.2 Summary

The Bayesian MSE is defined in (10.2) and is minimized by the estimator of (10.5), which is the mean of the posterior PDF. The example of a DC level in WGN with a Gaussian prior PDF is described in Section 10.4. The minimum MSE estimator for this example is given by (10.11) and represents a weighting between the data knowledge and prior knowledge. The corresponding minimum MSE is given by (10.14). The ability to




Figure 10.1 Improvement of estimator via prior knowledge

estimate the realization of a random variable based on data is described in Section 10.5 as being due to the correlation between the random variables. Theorem 10.2 summarizes the result that a jointly Gaussian PDF yields a conditional PDF that is also Gaussian, having a mean (10.24) and a covariance (10.25). This is then applied to the Bayesian linear model of (10.26) to yield the posterior PDF as summarized in Theorem 10.3. As will be shown in Chapter 11, the mean of the posterior PDF (10.28) is the minimum MSE estimator for a vector parameter. Section 10.7 discusses nuisance parameters from a Bayesian viewpoint, while Section 10.8 describes the potential difficulties of using a Bayesian estimator in a classical estimation problem.

10.3 Prior Knowledge and Estimation

It is a fundamental rule of estimation theory that the use of prior knowledge will lead to a more accurate estimator. For example, if a parameter is constrained to lie in a known interval, then any good estimator should produce only estimates within that interval. In Example 3.1 it was shown that the MVU estimator of A is the sample mean $\bar{x}$. However, this assumed that A could take on any value in the interval $-\infty < A < \infty$. Due to physical constraints it may be more reasonable to assume that A can take on only values in the finite interval $-A_0 \leq A \leq A_0$. To retain $\hat{A} = \bar{x}$ as the best estimator would be undesirable since $\hat{A}$ may yield values outside the known interval. As shown in Figure 10.1a, this is due to noise effects. Certainly, we would expect to improve our estimation if we used the *truncated* sample mean estimator

$$\hat{A} = \begin{cases} -A_0 & \bar{x} < -A_0 \\ \bar{x} & -A_0 \leq \bar{x} \leq A_0 \\ A_0 & \bar{x} > A_0 \end{cases}$$

which would be consistent with the known constraints. Such an estimator would have the PDF

$$\begin{aligned} p_{\hat{A}}(\xi; A) &= \Pr\{\bar{x} \leq -A_0\}\delta(\xi + A_0) \\ &\quad + p_{\bar{x}}(\xi; A)[u(\xi + A_0) - u(\xi - A_0)] \\ &\quad + \Pr\{\bar{x} \geq A_0\}\delta(\xi - A_0) \end{aligned} \quad (10.1)$$

where $u(x)$ is the unit step function. This is shown in Figure 10.1b. It is seen that $\hat{A}$ is a biased estimator. However, if we compare the MSE of the two estimators, we note that for any A in the interval $-A_0 \leq A \leq A_0$

$$\begin{aligned}
 \text{mse}(\hat{A}) &= \int_{-\infty}^{\infty} (\xi - A)^2 p_{\hat{A}}(\xi; A) d\xi \\
 &= \int_{-\infty}^{-A_0} (\xi - A)^2 p_{\hat{A}}(\xi; A) d\xi + \int_{-A_0}^{A_0} (\xi - A)^2 p_{\hat{A}}(\xi; A) d\xi \\
 &\quad + \int_{A_0}^{\infty} (\xi - A)^2 p_{\hat{A}}(\xi; A) d\xi \\
 &> \int_{-\infty}^{-A_0} (-A_0 - A)^2 p_{\hat{A}}(\xi; A) d\xi + \int_{-A_0}^{A_0} (\xi - A)^2 p_{\hat{A}}(\xi; A) d\xi \\
 &\quad + \int_{A_0}^{\infty} (A_0 - A)^2 p_{\hat{A}}(\xi; A) d\xi \\
 &= \text{mse}(\tilde{A}).
 \end{aligned}$$

Hence, $\tilde{A}$, the truncated sample mean estimator, is better than the sample mean estimator in terms of MSE. Although $\hat{A}$ is still the MVU estimator, we have been able to reduce the *mean square error* by allowing the estimator to be biased. In as much as we have been able to produce a better estimator, the question arises as to whether an *optimal* estimator exists for this problem. (The reader may recall that in the classical case the MSE criterion of optimality usually led to unrealizable estimators. We shall see that this is not a problem in the Bayesian approach.) We can answer affirmatively but only after reformulating the data model. Knowing that A must lie in a known interval, we suppose that the true value of A has been chosen from that interval. We then model the process of choosing a value as a random event to which a PDF can be assigned. With knowledge only of the interval and no inclination as to whether A should be nearer any particular value, it makes sense to assign a $\mathcal{U}[-A_0, A_0]$ PDF to the *random variable* A . The overall data model then appears as in Figure 10.2. As shown there, the act of choosing A according to the given PDF represents the departure of the Bayesian approach from the classical approach. The problem, as always, is to estimate the *value* of A or the *realization* of the random variable. However, now we can incorporate our knowledge of *how* A was chosen. For example, we might attempt to find an estimator $\hat{A}$ that would minimize the *Bayesian* MSE defined as

$$\text{Bmse}(\hat{A}) = E[(A - \hat{A})^2]. \quad (10.2)$$

We choose to define the error as $A - \hat{A}$ in contrast to the classical estimation error of $\hat{A} - A$. This definition will be useful later when we discuss a vector space interpretation of the Bayesian estimator. In (10.2) we emphasize that since A is a random variable, the expectation operator is with respect to the *joint PDF* $p(\mathbf{x}, A)$. This is a fundamentally different MSE than in the classical case. We distinguish it by using the Bmse notation. To appreciate the difference compare the classical MSE

$$\text{mse}(\hat{A}) = \int (\hat{A} - A)^2 p(\mathbf{x}; A) d\mathbf{x} \quad (10.3)$$




Figure 10.2 Bayesian approach to data modeling

to the Bayesian MSE

$$\text{Bmse}(\hat{A}) = \iint (A - \hat{A})^2 p(\mathbf{x}, A) d\mathbf{x} dA. \quad (10.4)$$

Even the underlying experiments are different, as is reflected in the averaging PDFs. If we were to assess the MSE performance using a Monte Carlo computer simulation, then for the classical approach we would choose a realization of $w[n]$ and add it to a *given* A . This procedure would be repeated M times. Each time we would add a new realization of $w[n]$ to the *same* A . In the Bayesian approach, for each realization we would choose A according to its PDF $\mathcal{U}[-A_0, A_0]$ and then generate $w[n]$ (assuming that $w[n]$ is independent of A). We would then repeat this procedure M times. In the classical case we would obtain a MSE for each assumed value of A , while in the Bayesian case the single MSE figure obtained would be an average over the PDF of A . Note that whereas the classical MSE will depend on A , and hence estimators that attempt to minimize the MSE will usually depend on A (see Section 2.4), the Bayesian MSE will not. In effect, we have integrated the parameter dependence away! It should be clear that comparing classical and Bayesian estimators is like comparing “apples and oranges.” The reader who is tempted to do so may become thoroughly confused. Nonetheless, at times, the *forms* of the estimators will be identical (see Problem 10.1).

To complete our example we now derive the estimator that minimizes the Bayesian MSE. First, we use Bayes’ theorem to write

$$p(\mathbf{x}, A) = p(A|\mathbf{x})p(\mathbf{x})$$

so that

$$\text{Bmse}(\hat{A}) = \int \left[\int (A - \hat{A})^2 p(A|\mathbf{x}) dA \right] p(\mathbf{x}) d\mathbf{x}.$$

Now since $p(\mathbf{x}) \geq 0$ for all $\mathbf{x}$, if the integral in brackets can be minimized for each $\mathbf{x}$, then the Bayesian MSE will be minimized. Hence, fixing $\mathbf{x}$ so that $\hat{A}$ is a scalar variable




Figure 10.3 Comparison of prior and posterior PDFs

(as opposed to a general function of $\mathbf{x}$), we have

$$\begin{aligned}
 \frac{\partial}{\partial \hat{A}} \int (A - \hat{A})^2 p(A|\mathbf{x}) dA &= \int \frac{\partial}{\partial \hat{A}} (A - \hat{A})^2 p(A|\mathbf{x}) dA \\
 &= \int -2(A - \hat{A}) p(A|\mathbf{x}) dA \\
 &= -2 \int A p(A|\mathbf{x}) dA + 2\hat{A} \int p(A|\mathbf{x}) dA
 \end{aligned}$$

which when set equal to zero results in

$$\hat{A} = \int A p(A|\mathbf{x}) dA$$

or finally

$$\hat{A} = E(A|\mathbf{x}) \quad (10.5)$$

since the conditional PDF must integrate to 1. It is seen that the optimal estimator in terms of minimizing the Bayesian MSE is the *mean* of the *posterior* PDF $p(A|\mathbf{x})$ (see also Problem 10.5 for an alternative derivation). The posterior PDF refers to the PDF of A *after* the data have been observed. In contrast, $p(A)$ or

$$p(A) = \int p(\mathbf{x}, A) d\mathbf{x}$$

may be thought of as the prior PDF of A , indicating the PDF *before* the data are observed. We will henceforth term the estimator that minimizes the Bayesian MSE the minimum mean square error (MMSE) estimator. Intuitively, the effect of observing data will be to concentrate the PDF of A as shown in Figure 10.3 (see also Problem 10.15). This is because knowledge of the data should reduce our uncertainty about A . We will return to this idea later.

In determining the MMSE estimator we first require the posterior PDF. We can use Bayes' rule to determine it as

$$\begin{aligned} p(A|\mathbf{x}) &= \frac{p(\mathbf{x}|A)p(A)}{p(\mathbf{x})} \\ &= \frac{p(\mathbf{x}|A)p(A)}{\int p(\mathbf{x}|A)p(A) dA}. \end{aligned} \quad (10.6)$$

Note that the denominator is just a normalizing factor, independent of A , needed to ensure that $p(A|\mathbf{x})$ integrates to 1. If we continue our example, we recall that the prior PDF $p(A)$ is $\mathcal{U}[-A_0, A_0]$. To specify the conditional PDF $p(\mathbf{x}|A)$ we need to further assume that the choice of A via $p(A)$ does not affect the PDF of the noise samples or that $w[n]$ is independent of A . Then, for $n = 0, 1, \dots, N-1$

$$\begin{aligned} p_x(x[n]|A) &= p_w(x[n] - A|A) \\ &= p_w(x[n] - A) \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}(x[n] - A)^2\right] \end{aligned}$$

and therefore

$$p(\mathbf{x}|A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2\right]. \quad (10.7)$$

It is apparent that the PDF is identical *in form* to the usual classical PDF $p(\mathbf{x}; A)$. In the Bayesian case, however, the PDF is a *conditional* PDF, hence the “|” separator, while in the classical case, it represents an unconditional PDF, albeit parameterized by A , hence the separator “;” (see also Problem 10.6). Using (10.6) and (10.7), the posterior PDF becomes

$$p(A|\mathbf{x}) = \begin{cases} \frac{\frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2\right]}{\int_{-A_0}^{A_0} \frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2\right] dA} & |A| \leq A_0 \\ 0 & |A| > A_0. \end{cases}$$

But

$$\begin{aligned} \sum_{n=0}^{N-1} (x[n] - A)^2 &= \sum_{n=0}^{N-1} x^2[n] - 2NA\bar{x} + NA^2 \\ &= N(A - \bar{x})^2 + \sum_{n=0}^{N-1} x^2[n] - N\bar{x}^2 \end{aligned}$$

so that we have

$$p(A|\mathbf{x}) = \begin{cases} \frac{1}{c\sqrt{2\pi\frac{\sigma^2}{N}}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right] & |A| \leq A_0 \\ 0 & |A| > A_0 \end{cases} \quad (10.8)$$

The factor c is determined by the requirement that $p(A|\mathbf{x})$ integrate to 1, resulting in

$$c = \int_{-A_0}^{A_0} \frac{1}{\sqrt{2\pi\frac{\sigma^2}{N}}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right] dA.$$

The PDF is seen to be a truncated Gaussian, as shown in Figure 10.3b. The MMSE estimator, which is the mean of $p(A|\mathbf{x})$, is

$$\begin{aligned} \hat{A} &= E(A|\mathbf{x}) \\ &= \int_{-\infty}^{\infty} A p(A|\mathbf{x}) dA \\ &= \frac{\int_{-A_0}^{A_0} A \frac{1}{\sqrt{2\pi\frac{\sigma^2}{N}}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right] dA}{\int_{-A_0}^{A_0} \frac{1}{\sqrt{2\pi\frac{\sigma^2}{N}}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right] dA}. \end{aligned} \quad (10.9)$$

Although this cannot be evaluated in closed form, we note that $\hat{A}$ will be a function of $\bar{x}$ as well as of A_0 and σ^2 (see Problem 10.7). The MMSE estimator will not be $\bar{x}$ due to the truncation shown in Figure 10.3b unless A_0 is so large that there is effectively no truncation. This will occur if $A_0 \gg \sqrt{\sigma^2/N}$. Otherwise, the estimator will be “biased” towards zero as opposed to being equal to $\bar{x}$. This is because the prior knowledge embodied in $p(A)$ would in the absence of the data $\mathbf{x}$ produce the MMSE estimator (see Problem 10.8)

$$\hat{A} = E(A) = 0.$$

The effect of the data is to position the posterior mean between $A = 0$ and $A = \bar{x}$ in a compromise between the prior knowledge and that contributed by the data. To further appreciate this weighting consider what happens as N becomes large so that the data knowledge becomes more important. As shown in Figure 10.4, as N increases, we have from (10.8) that the posterior PDF becomes more concentrated about $\bar{x}$ (since σ^2/N decreases). Hence, it becomes nearly Gaussian, and its mean becomes just $\bar{x}$. The MMSE estimator relies less and less on the prior knowledge and more on the data. It is said that the data “swamps out” the prior knowledge.




Figure 10.4 Effect of increasing data record on posterior PDF

The results of this example are true in general and are now summarized. The Bayesian approach to parameter estimation assumes that the parameter to be estimated is a realization of the random variable θ . As such, we assign a prior PDF $p(\theta)$ to it. After the data are observed, our state of knowledge about the parameter is summarized by the posterior PDF $p(\theta|\mathbf{x})$. An optimal estimator is defined to be the one that minimizes the MSE when averaged over all realizations of θ and $\mathbf{x}$, the so-called Bayesian MSE. This estimator is the mean of the posterior PDF or $\hat{\theta} = E(\theta|\mathbf{x})$. The estimator is determined explicitly as

$$\hat{\theta} = E(\theta|\mathbf{x}) = \int \theta p(\theta|\mathbf{x}) d\theta. \quad (10.10)$$

The MMSE estimator will in general depend on the prior knowledge as well as the data. If the prior knowledge is weak relative to that of the data, then the estimator will ignore the prior knowledge. Otherwise, the estimator will be “biased” towards the prior mean. As expected, the use of prior information always improves the estimation accuracy (see Example 10.1).

The choice of a prior PDF is critical in Bayesian estimation. The wrong choice will result in a poor estimator, similar to the problems of a classical estimator designed with an incorrect data model. Much of the controversy surrounding the use of Bayesian estimators stems from the inability in practice to be able to justify the prior PDF. Suffice it to say that unless the prior PDF can be based on the physical constraints of the problem, then classical estimation is more appropriate.

10.4 Choosing a Prior PDF

As shown in the previous section, once a prior PDF has been chosen, the MMSE estimator follows directly from (10.10). There is no question of existence as there is with the MVU estimator in the classical approach. The only practical stumbling block that remains, however, is whether or not $E(\theta|\mathbf{x})$ can be determined in closed form. In

the introductory example the posterior PDF $p(A|\mathbf{x})$ as given by (10.8) could not be found explicitly due to the need to normalize $p(\mathbf{x}|A)p(A)$ so that it integrates to 1. Additionally, the posterior mean could not be found, as evidenced by (10.9). We would have to resort to numerical integration to actually implement the MMSE estimator. This problem is compounded considerably in the vector parameter case. There the posterior PDF becomes

$$p(\boldsymbol{\theta}|\mathbf{x}) = \frac{p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}) d\boldsymbol{\theta}}$$

which requires a p -dimensional integration over $\boldsymbol{\theta}$. Additionally, the mean needs to be evaluated (see Chapter 11), requiring further integration. For practical MMSE estimators we need to be able to express them in closed form. The next example illustrates an important case where this is possible.

Example 10.1 - DC Level in WGN - Gaussian Prior PDF

We now modify our prior knowledge for the introductory example. Instead of assuming the uniform prior PDF

$$p(A) = \begin{cases} \frac{1}{2A_0} & |A| \leq A_0 \\ 0 & |A| > A_0 \end{cases}$$

which led to an intractable integration, consider the Gaussian prior PDF

$$p(A) = \frac{1}{\sqrt{2\pi\sigma_A^2}} \exp\left[-\frac{1}{2\sigma_A^2}(A - \mu_A)^2\right].$$

The two prior PDFs clearly express different prior knowledge about A , although with $\mu_A = 0$ and $3\sigma_A = A_0$ the Gaussian prior PDF could be thought of as incorporating the knowledge that $|A| \leq A_0$. Of course, values of A near zero are thought to be more probable with the Gaussian prior PDF. Now if

$$\begin{aligned} p(\mathbf{x}|A) &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2\right] \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n]\right] \exp\left[-\frac{1}{2\sigma^2}(NA^2 - 2NA\bar{x})\right] \end{aligned}$$

we have

$$\begin{aligned}
p(A|\mathbf{x}) &= \frac{p(\mathbf{x}|A)p(A)}{\int p(\mathbf{x}|A)p(A) dA} \\
&= \frac{\frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}\sqrt{2\pi\sigma_A^2}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n]\right] \exp\left[-\frac{1}{2\sigma^2}(NA^2 - 2NA\bar{x})\right]}{\int_{-\infty}^{\infty} \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}\sqrt{2\pi\sigma_A^2}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n]\right] \exp\left[-\frac{1}{2\sigma^2}(NA^2 - 2NA\bar{x})\right] \\
&\quad \cdot \frac{\exp\left[-\frac{1}{2\sigma_A^2}(A - \mu_A)^2\right]}{\exp\left[-\frac{1}{2\sigma_A^2}(A - \mu_A)^2\right]} dA} \\
&= \frac{\exp\left[-\frac{1}{2}\left(\frac{1}{\sigma^2}(NA^2 - 2NA\bar{x}) + \frac{1}{\sigma_A^2}(A - \mu_A)^2\right)\right]}{\int_{-\infty}^{\infty} \exp\left[-\frac{1}{2}\left(\frac{1}{\sigma^2}(NA^2 - 2NA\bar{x}) + \frac{1}{\sigma_A^2}(A - \mu_A)^2\right)\right] dA} \\
&= \frac{\exp\left[-\frac{1}{2}Q(A)\right]}{\int_{-\infty}^{\infty} \exp\left[-\frac{1}{2}Q(A)\right] dA}.
\end{aligned}$$

Note, however, that the denominator does not depend on A , being a normalizing factor, and the argument of the exponential is quadratic in A . Hence, $p(A|\mathbf{x})$ must be a Gaussian PDF whose mean and variance depend on $\mathbf{x}$. Continuing, we have for $Q(A)$

$$\begin{aligned}
Q(A) &= \frac{N}{\sigma^2}A^2 - \frac{2NA\bar{x}}{\sigma^2} + \frac{A^2}{\sigma_A^2} - \frac{2\mu_A A}{\sigma_A^2} + \frac{\mu_A^2}{\sigma_A^2} \\
&= \left(\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}\right)A^2 - 2\left(\frac{N}{\sigma^2}\bar{x} + \frac{\mu_A}{\sigma_A^2}\right)A + \frac{\mu_A^2}{\sigma_A^2}.
\end{aligned}$$

Let

$$\begin{aligned}
\sigma_{A|x}^2 &= \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \\
\mu_{A|x} &= \left(\frac{N}{\sigma^2}\bar{x} + \frac{\mu_A}{\sigma_A^2}\right)\sigma_{A|x}^2.
\end{aligned}$$

Then, by completing the square we have

$$Q(A) = \frac{1}{\sigma_{A|x}^2} \left(A^2 - 2\mu_{A|x}A + \mu_{A|x}^2\right) - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} + \frac{\mu_A^2}{\sigma_A^2}$$

$$= \frac{1}{\sigma_{A|x}^2} (A - \mu_{A|x})^2 - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} + \frac{\mu_A^2}{\sigma_A^2}$$

so that

$$\begin{aligned} p(A|x) &= \frac{\exp \left[-\frac{1}{2\sigma_{A|x}^2} (A - \mu_{A|x})^2 \right] \exp \left[-\frac{1}{2} \left(\frac{\mu_A^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} \right) \right]}{\int_{-\infty}^{\infty} \exp \left[-\frac{1}{2\sigma_{A|x}^2} (A - \mu_{A|x})^2 \right] \exp \left[-\frac{1}{2} \left(\frac{\mu_A^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} \right) \right] dA} \\ &= \frac{1}{\sqrt{2\pi\sigma_{A|x}^2}} \exp \left[-\frac{1}{2\sigma_{A|x}^2} (A - \mu_{A|x})^2 \right] \end{aligned}$$

where the last step follows from the requirement that $p(A|x)$ integrate to 1. The posterior PDF is also Gaussian, as claimed. (This result could also have been obtained by using Theorem 10.2 since $A, \mathbf{x}$ are jointly Gaussian.) In this form the MMSE estimator is readily found as

$$\begin{aligned} \hat{A} &= E(A|x) \\ &= \mu_{A|x} \\ &= \frac{N}{\sigma^2} \bar{x} + \frac{\mu_A}{\frac{\sigma^2}{\sigma_A^2}} \\ &= \frac{N}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \end{aligned}$$

or finally, the MMSE estimator is

$$\begin{aligned} \hat{A} &= \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{x} + \frac{\frac{\sigma^2}{N}}{\sigma_A^2 + \frac{\sigma^2}{N}} \mu_A \\ &= \alpha \bar{x} + (1 - \alpha) \mu_A \end{aligned} \quad (10.11)$$

where

$$\alpha = \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}}.$$

Note that α is a weighting factor since $0 < \alpha < 1$. Using a Gaussian prior PDF, we are able to determine the MMSE estimator explicitly. It is interesting to examine the interplay between the prior knowledge and the data. When there is little data so that $\sigma_A^2 \ll \sigma^2/N$, α is small and $\hat{A} \approx \mu_A$, but as more data are observed so that $\sigma_A^2 \gg \sigma^2/N$, $\alpha \approx 1$ and $\hat{A} \approx \bar{x}$. The weighting factor α depends directly on our confidence in the prior knowledge or σ_A^2 and the data knowledge or σ^2/N . (The quantity σ^2/N is interpreted as the conditional variance or $E[(\bar{x} - A)^2|A]$). Alternatively, we may view this process by examining the posterior PDF as N increases. Referring to Figure 10.5, as the data record length N increases, the posterior PDF becomes narrower. This is because the




Figure 10.5 Effect of increasing data record length on posterior PDF

posterior variance

$$\text{var}(A|\mathbf{x}) = \sigma_{A|\mathbf{x}}^2 = \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \quad (10.12)$$

will decrease. Also, the posterior mean (10.11) or $\hat{A}$ will also change with increasing N . For small N it will be approximately μ_A but will approach $\bar{x}$ for increasing N . In fact, as $N \rightarrow \infty$, we will have $\hat{A} \rightarrow \bar{x}$, which in turn approaches the true value of A chosen. Observe that if there is no prior knowledge, which can be modeled by letting $\sigma_A^2 \rightarrow \infty$, then $\hat{A} \rightarrow \bar{x}$ for any data record length. The “classical” estimator is obtained. Finally, it was originally claimed that by using prior knowledge we could improve the estimation accuracy. To see why this is so recall that

$$\text{Bmse}(\hat{A}) = E[(A - \hat{A})^2]$$

where we evaluate the expectation with respect to $p(\mathbf{x}, A)$. But

$$\begin{aligned} \text{Bmse}(\hat{A}) &= \iint (A - \hat{A})^2 p(\mathbf{x}, A) d\mathbf{x} dA \\ &= \iint (A - \hat{A})^2 p(A|\mathbf{x}) dA p(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

Since $\hat{A} = E(A|\mathbf{x})$, we have

$$\begin{aligned} \text{Bmse}(\hat{A}) &= \int \int [A - E(A|\mathbf{x})]^2 p(A|\mathbf{x}) dA p(\mathbf{x}) d\mathbf{x} \\ &= \int \text{var}(A|\mathbf{x}) p(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (10.13)$$

We see that *the Bayesian MSE is just the variance of the posterior PDF when averaged over the PDF of $\mathbf{x}$* . As such, we have

$$\begin{aligned} \text{Bmse}(\hat{A}) &= \int \sigma_{A|\mathbf{x}}^2 p(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \end{aligned}$$

since $\sigma_{A|x}^2$ does not depend on $\mathbf{x}$. This can be rewritten as

$$\text{Bmse}(\hat{A}) = \frac{\sigma^2}{N} \left(\frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \right) \quad (10.14)$$

so that finally we see that

$$\text{Bmse}(\hat{A}) < \frac{\sigma^2}{N}$$

where σ^2/N is the minimum MSE obtained when no prior knowledge is available (let $\sigma_A^2 \rightarrow \infty$). Clearly, any prior knowledge *when modeled in the Bayesian sense* will improve our *Bayesian estimator*. $\diamond$

Gaussian prior PDFs are quite useful in practice due to their mathematical tractability, as illustrated by the previous example. The basic property that makes this so is the *reproducing property*. If $p(\mathbf{x}, A)$ is Gaussian, then $p(A)$, being the marginal PDF, is Gaussian, as is the posterior PDF $p(A|\mathbf{x})$. Hence, the form of the PDF remains the same, as it is conditioned on $\mathbf{x}$. Only the mean and variance change. Another example of a PDF sharing this property is given in Problem 10.10. Furthermore, Gaussian prior PDFs occur naturally in many practical problems. For the previous example we can envision the problem of measuring the DC voltage of a power source by means of a DC voltmeter. If we set the power source to 10 volts, for example, we would probably be willing to assume that the true voltage is close to 10 volts. Our prior knowledge then might be modeled as $A \sim \mathcal{N}(10, \sigma_A^2)$, where σ_A^2 would be small for a precision power source and large for a less reliable one. Next, we could take N measurements of the voltage. Our model for the measurements could be $x[n] = A + w[n]$, where the voltmeter error $w[n]$ is modeled by WGN with variance σ^2 . The value of σ^2 would reflect our confidence in the quality of the voltmeter. The MMSE estimator of the true voltage would be given by (10.11). If we repeated the procedure for an ensemble of power sources and voltmeters with the same error characteristics, then our estimator would minimize the Bayesian MSE.

10.5 Properties of the Gaussian PDF

We now generalize the results of the previous section by examining the properties of the Gaussian PDF. The results of this section will be needed for the derivations of Bayesian estimators in the next chapter. The bivariate Gaussian PDF is first investigated to illustrate the important properties. Then, the corresponding results for the general multivariate Gaussian PDF are described. The remarkable property that we shall exploit is that the posterior PDF is also Gaussian, although with a different mean and variance. Some physical interpretations are stressed.

Consider a jointly Gaussian random vector $[x y]^T$ whose PDF is

$$p(x, y) = \frac{1}{2\pi \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} \begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix}^T \mathbf{C}^{-1} \begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix} \right]. \quad (10.15)$$




Figure 10.6 Contours of constant density for bivariate Gaussian PDF

This is also termed the *bivariate* Gaussian PDF. The mean vector and covariance matrix are

$$E \left(\begin{bmatrix} x \\ y \end{bmatrix} \right) = \begin{bmatrix} E(x) \\ E(y) \end{bmatrix}$$

$$\mathbf{C} = \begin{bmatrix} \text{var}(x) & \text{cov}(x, y) \\ \text{cov}(y, x) & \text{var}(y) \end{bmatrix}.$$

Note that the marginal PDFs $p(x)$ and $p(y)$ are also Gaussian, as can be verified by the integrations

$$p(x) = \int_{-\infty}^{\infty} p(x, y) dy = \frac{1}{\sqrt{2\pi\text{var}(x)}} \exp \left[-\frac{1}{2\text{var}(x)} (x - E(x))^2 \right]$$

$$p(y) = \int_{-\infty}^{\infty} p(x, y) dx = \frac{1}{\sqrt{2\pi\text{var}(y)}} \exp \left[-\frac{1}{2\text{var}(y)} (y - E(y))^2 \right].$$

The contours along which the PDF $p(x, y)$ is constant are those values of x and y for which

$$\begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix}^T \mathbf{C}^{-1} \begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix}$$

is a constant. They are shown in Figure 10.6 as elliptical contours. Once x , say x_0 , is observed, the conditional PDF of y becomes

$$p(y|x_0) = \frac{p(x_0, y)}{p(x_0)} = \frac{p(x_0, y)}{\int_{-\infty}^{\infty} p(x_0, y) dy}$$

so that the conditional PDF of y is that of the cross section shown in Figure 10.6 when suitably normalized to integrate to 1. It is readily seen that since $p(x_0, y)$ (where x_0 is a fixed number) has the Gaussian form in y (from (10.15) the exponential argument is quadratic in y), the conditional PDF must also be Gaussian. Since $p(y)$ is also Gaussian, we may view this property as saying that if x and y are jointly Gaussian, the prior PDF $p(y)$ and posterior PDF $p(y|x)$ are both Gaussian. In Appendix 10A we derive the exact PDF as summarized in the following theorem.

Theorem 10.1 (Conditional PDF of Bivariate Gaussian) If x and y are distributed according to a bivariate Gaussian PDF with mean vector $[E(x) E(y)]^T$ and covariance matrix

$$\mathbf{C} = \begin{bmatrix} \text{var}(x) & \text{cov}(x, y) \\ \text{cov}(y, x) & \text{var}(y) \end{bmatrix}$$

so that

$$p(x, y) = \frac{1}{2\pi \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} \begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix}^T \mathbf{C}^{-1} \begin{bmatrix} x - E(x) \\ y - E(y) \end{bmatrix} \right],$$

then the conditional PDF $p(y|x)$ is also Gaussian and

$$E(y|x) = E(y) + \frac{\text{cov}(x, y)}{\text{var}(x)} (x - E(x)) \quad (10.16)$$

$$\text{var}(y|x) = \text{var}(y) - \frac{\text{cov}^2(x, y)}{\text{var}(x)}. \quad (10.17)$$

We can view this result in the following way. Before observing x , the random variable y is distributed according to the prior PDF $p(y)$ or $y \sim \mathcal{N}(E(y), \text{var}(y))$. After observing x , the random variable y is distributed according to the posterior PDF $p(y|x)$ given in Theorem 10.1. Only the mean and variance have changed. Assuming that x and y are not independent and hence $\text{cov}(x, y) \neq 0$, the posterior PDF becomes more concentrated since there is less uncertainty about y . To verify this, note from (10.17) that

$$\begin{aligned} \text{var}(y|x) &= \text{var}(y) \left[1 - \frac{\text{cov}^2(x, y)}{\text{var}(x)\text{var}(y)} \right] \\ &= \text{var}(y)(1 - \rho^2) \end{aligned} \quad (10.18)$$

where

$$\rho = \frac{\text{cov}(x, y)}{\sqrt{\text{var}(x)\text{var}(y)}} \quad (10.19)$$

is the correlation coefficient satisfying $|\rho| \leq 1$. From our previous discussions we also realize that $E(y|x)$ is the MMSE estimator of y after observing x , so that from (10.16)

$$\hat{y} = E(y) + \frac{\text{cov}(x, y)}{\text{var}(x)} (x - E(x)). \quad (10.20)$$

In normalized form (a random variable with zero mean and unity variance) this becomes

$$\frac{\hat{y} - E(y)}{\sqrt{\text{var}(y)}} = \frac{\text{cov}(x, y)}{\sqrt{\text{var}(x)\text{var}(y)}} \frac{x - E(x)}{\sqrt{\text{var}(x)}}$$

or

$$\hat{y}_n = \rho x_n. \quad (10.21)$$




Figure 10.7

Contours of constant density of normalized bivariate PDF

The correlation coefficient then acts to scale the normalized observation x_n to obtain the MMSE estimator of the normalized realization of the random variable y_n . If the random variables are already normalized ($E(x) = E(y) = 0$, $\text{var}(x) = \text{var}(y) = 1$), the constant PDF contours appear as in Figure 10.7. The locations of the peaks of $p(x, y)$, when considered as a function of y for each x , is the dashed line $y = \rho x$, and it is readily shown that $\hat{y} = E(y|x) = \rho x$ (see Problem 10.12). *The MMSE estimator therefore exploits the correlation between the random variables to estimate the realization of one based on the realization of the other.*

The minimum MSE is, from (10.13) and (10.18),

$$\begin{aligned} \text{Bmse}(\hat{y}) &= \int \text{var}(y|x) p(x) dx \\ &= \text{var}(y|x) \\ &= \text{var}(y)(1 - \rho^2) \end{aligned} \quad (10.22)$$

since the posterior variance does not depend on x ($\text{var}(y)$ and ρ depend on the covariance matrix only). Hence, the quality of our estimator also depends on the correlation coefficient, which is a measure of the statistical dependence between x and y .

To generalize these results consider a jointly Gaussian vector $[\mathbf{x}^T \mathbf{y}^T]^T$, where $\mathbf{x}$ is $k \times 1$ and $\mathbf{y}$ is $l \times 1$. In other words, $[\mathbf{x}^T \mathbf{y}^T]^T$ is distributed according to a multivariate Gaussian PDF. Then, the conditional PDF of $\mathbf{y}$ for a given $\mathbf{x}$ is also Gaussian, as summarized in the following theorem (see Appendix 10A for proof).

Theorem 10.2 (Conditional PDF of Multivariate Gaussian) *If $\mathbf{x}$ and $\mathbf{y}$ are jointly Gaussian, where $\mathbf{x}$ is $k \times 1$ and $\mathbf{y}$ is $l \times 1$, with mean vector $[E(\mathbf{x})^T E(\mathbf{y})^T]^T$ and partitioned covariance matrix*

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_{xx} & \mathbf{C}_{xy} \\ \mathbf{C}_{yx} & \mathbf{C}_{yy} \end{bmatrix} = \begin{bmatrix} k \times k & k \times l \\ l \times k & l \times l \end{bmatrix} \quad (10.23)$$

so that

$$p(\mathbf{x}, \mathbf{y}) = \frac{1}{(2\pi)^{\frac{k+l}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} \left(\begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix} \right)^T \mathbf{C}^{-1} \left(\begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix} \right) \right],$$

then the conditional PDF $p(\mathbf{y}|\mathbf{x})$ is also Gaussian and

$$E(\mathbf{y}|\mathbf{x}) = E(\mathbf{y}) + \mathbf{C}_{yx}\mathbf{C}_{xx}^{-1}(\mathbf{x} - E(\mathbf{x})) \quad (10.24)$$

$$\mathbf{C}_{y|x} = \mathbf{C}_{yy} - \mathbf{C}_{yx}\mathbf{C}_{xx}^{-1}\mathbf{C}_{xy}. \quad (10.25)$$

Note that the covariance matrix of the conditional PDF does not depend on $\mathbf{x}$, although this property is not generally true. This will be useful later. As in the bivariate case, the prior PDF $p(\mathbf{y})$ is Gaussian, as well as the posterior PDF $p(\mathbf{y}|\mathbf{x})$. The question may arise as to when the jointly Gaussian assumption may be made. In the next section we examine an important data model for which this holds, the Bayesian linear model.

10.6 Bayesian Linear Model

Recall that in Example 10.1 the data model was

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $A \sim \mathcal{N}(\mu_A, \sigma_A^2)$, and $w[n]$ is WGN independent of A . In terms of vectors we have the equivalent data model

$$\mathbf{x} = \mathbf{1}A + \mathbf{w}.$$

This appears to be of the form of the linear model described in Chapter 4 except for the assumption that A is a random variable. It should not then be surprising that a Bayesian equivalent of the general linear model can be defined. In particular let the data be modeled as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \quad (10.26)$$

where $\mathbf{x}$ is an $N \times 1$ data vector, $\mathbf{H}$ is a known $N \times p$ matrix, $\boldsymbol{\theta}$ is a $p \times 1$ random vector with prior PDF $\mathcal{N}(\boldsymbol{\mu}_\theta, \mathbf{C}_\theta)$, and $\mathbf{w}$ is an $N \times 1$ noise vector with PDF $\mathcal{N}(\mathbf{0}, \mathbf{C}_w)$ and independent of $\boldsymbol{\theta}$. This data model is termed the *Bayesian general linear model*. It differs from the classical general linear model in that $\boldsymbol{\theta}$ is modeled as a random variable with a Gaussian prior PDF. It will be of interest in deriving Bayesian estimators to have an explicit expression for the posterior PDF $p(\boldsymbol{\theta}|\mathbf{x})$. From Theorem 10.2 we know that if $\mathbf{x}$ and $\boldsymbol{\theta}$ are jointly Gaussian, then the posterior PDF is also Gaussian. Hence, it only remains to verify that this is indeed the case. Let $\mathbf{z} = [\mathbf{x}^T \boldsymbol{\theta}^T]^T$, so that from (10.26) we have

$$\begin{aligned} \mathbf{z} &= \begin{bmatrix} \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \\ \boldsymbol{\theta} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{H} & \mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{w} \end{bmatrix} \end{aligned}$$

where the identity matrices are of dimension $N \times N$ (upper right) and $p \times p$ (lower left), and $\mathbf{0}$ is an $N \times N$ matrix of zeros. Since $\boldsymbol{\theta}$ and $\mathbf{w}$ are independent of each other and each one is Gaussian, they are jointly Gaussian. Furthermore, because $\mathbf{z}$ is a linear

transformation of a Gaussian vector, it too is Gaussian. Hence, Theorem 10.2 applies directly, and we need only determine the mean and covariance of the posterior PDF. We identify $\mathbf{x}$ as $\mathbf{H}\boldsymbol{\theta} + \mathbf{w}$ and $\mathbf{y}$ as $\boldsymbol{\theta}$ to obtain the means

$$\begin{aligned} E(\mathbf{x}) &= E(\mathbf{H}\boldsymbol{\theta} + \mathbf{w}) = \mathbf{H}E(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\mu}_\theta \\ E(\mathbf{y}) &= E(\boldsymbol{\theta}) = \boldsymbol{\mu}_\theta \end{aligned}$$

and covariances

$$\begin{aligned} \mathbf{C}_{xx} &= E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{x} - E(\mathbf{x}))^T] \\ &= E[(\mathbf{H}\boldsymbol{\theta} + \mathbf{w} - \mathbf{H}\boldsymbol{\mu}_\theta)(\mathbf{H}\boldsymbol{\theta} + \mathbf{w} - \mathbf{H}\boldsymbol{\mu}_\theta)^T] \\ &= E[(\mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \mathbf{w})(\mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \mathbf{w})^T] \\ &= \mathbf{H}E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)^T]\mathbf{H}^T + E(\mathbf{w}\mathbf{w}^T) \\ &= \mathbf{H}\mathbf{C}_\theta\mathbf{H}^T + \mathbf{C}_w \end{aligned}$$

recalling that $\boldsymbol{\theta}$ and $\mathbf{w}$ are independent. Also, the cross-covariance matrix is

$$\begin{aligned} \mathbf{C}_{yx} &= E[(\mathbf{y} - E(\mathbf{y}))(\mathbf{x} - E(\mathbf{x}))^T] \\ &= E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \mathbf{w})^T] \\ &= E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta))^T] \\ &= \mathbf{C}_\theta\mathbf{H}^T. \end{aligned}$$

We can now summarize our results for the Bayesian general linear model.

Theorem 10.3 (Posterior PDF for the Bayesian General Linear Model) *If the observed data $\mathbf{x}$ can be modeled as*

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w} \quad (10.27)$$

where $\mathbf{x}$ is an $N \times 1$ data vector, $\mathbf{H}$ is a known $N \times p$ matrix, $\boldsymbol{\theta}$ is a $p \times 1$ random vector with prior PDF $\mathcal{N}(\boldsymbol{\mu}_\theta, \mathbf{C}_\theta)$, and $\mathbf{w}$ is an $N \times 1$ noise vector with PDF $\mathcal{N}(\mathbf{0}, \mathbf{C}_w)$ and independent of $\boldsymbol{\theta}$, then the posterior PDF $p(\boldsymbol{\theta}|\mathbf{x})$ is Gaussian with mean

$$E(\boldsymbol{\theta}|\mathbf{x}) = \boldsymbol{\mu}_\theta + \mathbf{C}_\theta\mathbf{H}^T(\mathbf{H}\mathbf{C}_\theta\mathbf{H}^T + \mathbf{C}_w)^{-1}(\mathbf{x} - \mathbf{H}\boldsymbol{\mu}_\theta) \quad (10.28)$$

and covariance

$$\mathbf{C}_{\theta|\mathbf{x}} = \mathbf{C}_\theta - \mathbf{C}_\theta\mathbf{H}^T(\mathbf{H}\mathbf{C}_\theta\mathbf{H}^T + \mathbf{C}_w)^{-1}\mathbf{H}\mathbf{C}_\theta. \quad (10.29)$$

In contrast to the classical general linear model, $\mathbf{H}$ need not be full rank to ensure the invertibility of $\mathbf{H}\mathbf{C}_\theta\mathbf{H}^T + \mathbf{C}_w$. We illustrate the use of these formulas by applying them to Example 10.1.

Example 10.2 - DC Level in WGN - Gaussian Prior PDF (continued)

Since $x[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$ with $A \sim \mathcal{N}(\mu_A, \sigma_A^2)$ and $w[n]$ is WGN with variance σ^2 and independent of A , we have the Bayesian general linear model

$$\mathbf{x} = \mathbf{1}A + \mathbf{w}.$$

According to Theorem 10.3, $p(A|\mathbf{x})$ is Gaussian and

$$E(A|\mathbf{x}) = \mu_A + \sigma_A^2 \mathbf{1}^T (\mathbf{1} \sigma_A^2 \mathbf{1}^T + \sigma^2 \mathbf{I})^{-1} (\mathbf{x} - \mathbf{1} \mu_A).$$

Using Woodbury's identity (see Appendix 1)

$$\left(\mathbf{I} + \frac{\sigma_A^2}{\sigma^2} \mathbf{1} \mathbf{1}^T \right)^{-1} = \mathbf{I} - \frac{\frac{\sigma_A^2}{\sigma^2} \mathbf{1} \mathbf{1}^T}{1 + N \frac{\sigma_A^2}{\sigma^2}} \quad (10.30)$$

so that

$$\begin{aligned} E(A|\mathbf{x}) &= \mu_A + \frac{\sigma_A^2}{\sigma^2} \mathbf{1}^T \left(\mathbf{I} - \frac{\mathbf{1} \mathbf{1}^T}{N + \frac{\sigma_A^2}{\sigma^2}} \right) (\mathbf{x} - \mathbf{1} \mu_A) \\ &= \mu_A + \frac{\sigma_A^2}{\sigma^2} \left(\mathbf{1}^T - \frac{N}{N + \frac{\sigma_A^2}{\sigma^2}} \mathbf{1}^T \right) (\mathbf{x} - \mathbf{1} \mu_A) \\ &= \mu_A + \frac{\sigma_A^2}{\sigma^2} \left(1 - \frac{N}{N + \frac{\sigma_A^2}{\sigma^2}} \right) (N \bar{x} - N \mu_A) \\ &= \mu_A + \frac{N}{N + \frac{\sigma_A^2}{\sigma^2}} (\bar{x} - \mu_A) \\ &= \mu_A + \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} (\bar{x} - \mu_A). \end{aligned} \quad (10.31)$$

It is interesting to note that in this form the MMSE estimator resembles a “sequential”-type estimator (see Section 8.7). The estimator with no data or $\hat{A} = \mu_A$ is corrected by the error between the data estimator $\bar{x}$ and the “previous” estimate μ_A . The “gain factor” $\sigma_A^2/(\sigma_A^2 + \sigma^2/N)$ depends on our confidence in the previous estimate and the current data. We will see later that a sequential MMSE estimator can be defined and will have just these properties. Finally, with some more algebra (10.11) can be obtained.

To find the posterior variance we use (10.29), so that

$$\text{var}(A|\mathbf{x}) = \sigma_A^2 - \sigma_A^2 \mathbf{1}^T (\mathbf{1} \sigma_A^2 \mathbf{1}^T + \sigma^2 \mathbf{I})^{-1} \mathbf{1} \sigma_A^2.$$

Using (10.30) once again, we obtain

$$\begin{aligned}
 \text{var}(A|\mathbf{x}) &= \sigma_A^2 - \frac{\sigma_A^2}{\sigma^2} \mathbf{1}^T \left(\mathbf{I} - \frac{\mathbf{1}\mathbf{1}^T}{N + \frac{\sigma^2}{\sigma_A^2}} \right) \mathbf{1} \sigma_A^2 \\
 &= \sigma_A^2 - \frac{\sigma_A^4}{\sigma^2} \left(N - \frac{N^2}{N + \frac{\sigma^2}{\sigma_A^2}} \right) \\
 &= \frac{\frac{\sigma^2}{N} \sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}}
 \end{aligned}$$

which is just (10.12). $\diamond$

In the succeeding chapters we will make extensive use of the Bayesian linear model. For future reference we point out that the mean (10.28) and covariance (10.29) of the posterior PDF can be expressed in alternative forms as (see Problem 10.13)

$$E(\boldsymbol{\theta}|\mathbf{x}) = \boldsymbol{\mu}_\theta + (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} (\mathbf{x} - \mathbf{H} \boldsymbol{\mu}_\theta) \quad (10.32)$$

and

$$\mathbf{C}_{\theta|x} = (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \quad (10.33)$$

or

$$\mathbf{C}_{\theta|x}^{-1} = \mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H}. \quad (10.34)$$

The latter expression is particularly interesting. For the previous example we would have

$$\begin{aligned}
 \frac{1}{\text{var}(A|\mathbf{x})} &= \frac{1}{\sigma_A^2} + \mathbf{1}^T (\sigma^2 \mathbf{I})^{-1} \mathbf{1} \\
 &= \frac{1}{\sigma_A^2} + \frac{1}{\frac{\sigma^2}{N}}.
 \end{aligned}$$

This form lends itself to the interpretation that the “information” or reciprocal of the variance of the prior knowledge $1/\sigma_A^2$ and the “information” of the data $1/(\sigma^2/N)$ add to yield the information embodied in the posterior PDF.

10.7 Nuisance Parameters

Many estimation problems are characterized by a set of unknown parameters, of which we are really interested only in a subset. The remaining parameters, which serve only

to complicate the problem, are referred to as *nuisance* parameters. Such would be the case if for a DC level in WGN, we were interested in estimating σ^2 but A was unknown. The DC level A would be the nuisance parameter. If we assume that the parameters are deterministic, as in the classical estimation approach, then in general we have no alternative but to estimate σ^2 and A . In the Bayesian approach we can rid ourselves of nuisance parameters by “integrating them out.” Suppose the unknown parameters to be estimated are θ and some additional nuisance parameters α are present. Then, if $p(\theta, \alpha | \mathbf{x})$ denotes the posterior PDF, we can determine the posterior PDF of θ only as

$$p(\theta | \mathbf{x}) = \int p(\theta, \alpha | \mathbf{x}) d\alpha. \quad (10.35)$$

We can also express this as

$$p(\theta | \mathbf{x}) = \frac{p(\mathbf{x} | \theta) p(\theta)}{\int p(\mathbf{x} | \theta) p(\theta) d\theta} \quad (10.36)$$

where

$$p(\mathbf{x} | \theta) = \int p(\mathbf{x} | \theta, \alpha) p(\alpha | \theta) d\alpha. \quad (10.37)$$

If we furthermore assume that the nuisance parameters are independent of the desired parameters, then (10.37) reduces to

$$p(\mathbf{x} | \theta) = \int p(\mathbf{x} | \theta, \alpha) p(\alpha) d\alpha. \quad (10.38)$$

We observe that the nuisance parameters are first integrated out of the conditional PDF $p(\mathbf{x} | \theta, \alpha)$ and then the posterior PDF is found as usual by Bayes' theorem. If a MMSE estimator is desired, we need only determine the mean of the posterior PDF. The nuisance parameters no longer enter into the problem. Of course, their presence will affect the final estimator since from (10.38) $p(\mathbf{x} | \theta)$ depends on $p(\alpha)$. From a theoretical viewpoint, the Bayesian approach does not suffer from the problems of classical estimators in which nuisance parameters may invalidate an estimator. We now illustrate the approach with an example.

Example 10.3 - Scaled Covariance Matrix

Assume that we observe the $N \times 1$ data vector $\mathbf{x}$ whose conditional PDF $p(\mathbf{x} | \theta, \sigma^2)$ is $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{C}(\theta))$. (The reader should not confuse $\mathbf{C}(\theta)$, the scaled covariance matrix of $\mathbf{x}$, with $\mathbf{C}_\theta$, the covariance matrix of θ .) The parameter θ is to be estimated, and σ^2 is to be regarded as a nuisance parameter. The covariance matrix depends on θ in some unspecified manner. We assign the prior PDF to σ^2 of

$$p(\sigma^2) = \begin{cases} \frac{\lambda \exp(-\frac{\lambda}{\sigma^2})}{\sigma^4} & \sigma^2 > 0 \\ 0 & \sigma^2 < 0 \end{cases} \quad (10.39)$$

where $\lambda > 0$, and assume σ^2 to be independent of θ . The prior PDF is a special case of the *inverted gamma* PDF. Then, from (10.38) we have

$$\begin{aligned} p(\mathbf{x}|\theta) &= \int p(\mathbf{x}|\theta, \sigma^2) p(\sigma^2) d\sigma^2 \\ &= \int_0^\infty \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}[\sigma^2 \mathbf{C}(\theta)]} \exp\left[-\frac{1}{2} \mathbf{x}^T (\sigma^2 \mathbf{C}(\theta))^{-1} \mathbf{x}\right] \frac{\lambda \exp(-\frac{\lambda}{\sigma^2})}{\sigma^4} d\sigma^2 \\ &= \int_0^\infty \frac{1}{(2\pi)^{\frac{N}{2}} \sigma^N \det^{\frac{1}{2}}[\mathbf{C}(\theta)]} \exp\left[-\frac{1}{2\sigma^2} \mathbf{x}^T \mathbf{C}^{-1}(\theta) \mathbf{x}\right] \frac{\lambda \exp(-\frac{\lambda}{\sigma^2})}{\sigma^4} d\sigma^2. \end{aligned}$$

Letting $\xi = 1/\sigma^2$, we have

$$p(\mathbf{x}|\theta) = \frac{\lambda}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}[\mathbf{C}(\theta)]} \int_0^\infty \xi^{\frac{N}{2}} \exp\left[-\left(\lambda + \frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1}(\theta) \mathbf{x}\right) \xi\right] d\xi,$$

But

$$\int_0^\infty x^{m-1} \exp(-ax) dx = a^{-m} \Gamma(m)$$

for $a > 0$ and $m > 0$, resulting from the properties of the gamma integral. Hence, the integral can be evaluated to yield

$$p(\mathbf{x}|\theta) = \frac{\lambda \Gamma(\frac{N}{2} + 1)}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}[\mathbf{C}(\theta)] \left(\lambda + \frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1}(\theta) \mathbf{x}\right)^{\frac{N}{2} + 1}}.$$

The posterior PDF may be found by substituting this into (10.36) (at least in theory!) $\diamond$

10.8 Bayesian Estimation for Deterministic Parameters

Although strictly speaking the Bayesian approach can be applied only when θ is random, in practice it is often used for deterministic parameter estimation. By this we mean that the Bayesian assumptions are made to obtain an estimator, the MMSE estimator for example, and then used as if θ were nonrandom. Such might be the case if no MVU estimator existed. For instance, we may not be able to find an unbiased estimator that is uniformly better in terms of variance than all others (see Example 2.3). Within the Bayesian framework, however, the MMSE estimator *always exists* and thus provides an estimator which at least “on the average” (as different values of θ are chosen) works well. Of course, for a *particular* θ it may not perform well, and this is the risk we take by applying it for a deterministic parameter. To illustrate this potential pitfall consider Example 10.1 for which (see (10.11))




Figure 10.8 Mean square error of MVU and Bayesian estimators for deterministic DC level in WGN

$$\begin{aligned}\hat{A} &= \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2/N} \bar{x} + \frac{\sigma^2/N}{\sigma_A^2 + \sigma^2/N} \mu_A \\ &= \alpha \bar{x} + (1 - \alpha) \mu_A\end{aligned}$$

and $0 < \alpha < 1$. If A is a *deterministic* parameter, we can evaluate the MSE using (2.6) as

$$\text{mse}(\hat{A}) = \text{var}(\hat{A}) + b^2(\hat{A})$$

where $b(\hat{A}) = E(\hat{A}) - A$ is the bias. Then,

$$\begin{aligned}\text{mse}(\hat{A}) &= \alpha^2 \text{var}(\bar{x}) + [\alpha A + (1 - \alpha) \mu_A - A]^2 \\ &= \alpha^2 \frac{\sigma^2}{N} + (1 - \alpha)^2 (A - \mu_A)^2.\end{aligned}\quad (10.40)$$

It is seen that the use of the Bayesian estimator reduces the variance, since $0 < \alpha < 1$, but may substantially increase the bias component of the MSE. As shown further in Figure 10.8, the Bayesian estimator exhibits less MSE than the MVU estimator $\bar{x}$ only if A is close to the prior mean μ_A . Otherwise, it is a poorer estimator. It does not have the desirable property of being *uniformly* less in MSE than all other estimators but only “on the average” or in a Bayesian sense. Hence, only if A is random and $\text{mse}(\hat{A})$ can thus be considered to be the MSE *conditioned* on a known value of A , do we obtain

$$\begin{aligned}\text{Bmse}(\hat{A}) &= E_A[\text{mse}(\hat{A})] \\ &= \alpha^2 \frac{\sigma^2}{N} + (1 - \alpha)^2 E_A[(A - \mu_A)^2] \\ &= \alpha^2 \frac{\sigma^2}{N} + (1 - \alpha)^2 \sigma_A^2 \\ &= \frac{\sigma^2}{N} \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \\ &< \frac{\sigma^2}{N} = \text{Bmse}(\bar{x}),\end{aligned}$$

or that the *Bayesian* MSE is less. In effect, the Bayesian MMSE estimator trades off bias for variance in an attempt to reduce the overall MSE. In doing so it benefits from the prior knowledge that $A \sim \mathcal{N}(\mu_A, \sigma_A^2)$. This allows it to adjust α so that the overall MSE, on the average, is less. The large MSEs shown in Figure 10.8 for A not near μ_A are of no consequence to the choice of α since they occur infrequently. Of course, the classical estimation approach will not have this advantage, being required to produce an estimator that has good performance for *all* A .

A second concern is the choice of the prior PDF. If no prior knowledge is available, as we assume in classical estimation, then we do not want to apply a Bayesian estimator based on a highly concentrated prior PDF. Again considering the same example, if A is deterministic but we apply the Bayesian estimator anyway, we see that

$$E(\hat{A}) = \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2/N} A + \frac{\sigma^2/N}{\sigma_A^2 + \sigma^2/N} \mu_A.$$

There is a bias and it is towards μ_A . To alleviate this bias we would need to specify a prior PDF that is nearly flat, or for this example, by letting $\sigma_A^2 \rightarrow \infty$. Otherwise, as already described, the MSE can be quite large for A not near μ_A . (Also, note that as $\sigma_A^2 \rightarrow \infty$, we have $\hat{A} \rightarrow \bar{x}$ and $\text{mse}(\hat{A}) \rightarrow \text{mse}(\bar{x})$, so that the two curves in Figure 10.8 become identical.) In general, such a prior PDF is termed a *noninformative* prior PDF. The use of a Bayesian estimator for a deterministic parameter is often justified on the basis that the noninformative prior PDF does not add any information to the problem. How the prior PDF is chosen for a given problem is beyond the scope of our discussions (see Problems 10.15–10.17 for one possible approach). Suffice it to say that there is a good deal of controversy surrounding this idea. The interested reader should consult [Box and Tiao 1973] for further details and philosophy.

References

- Ash, R., *Information Theory*, J. Wiley, New York, 1965.
 Box, G.E.P., G.C. Tiao, *Bayesian Inference in Statistical Analysis*, Addison-Wesley, Reading, Mass., 1973.
 Gallager, R.G., *Information Theory and Reliable Communication*, J. Wiley, New York, 1968.
 Zacks, S., *Parametric Statistical Inference*, Pergamon, New York, 1981.

Problems

- 10.1** In this problem we attempt to apply the Bayesian approach to the estimation of a deterministic parameter. Since the parameter is deterministic, we assign the prior PDF $p(\theta) = \delta(\theta - \theta_0)$, where θ_0 is the true value. Find the MMSE estimator for this prior PDF and explain your results.
- 10.2** Two random variables x and y are said to be conditionally independent of each other if the joint conditional PDF factors as

$$p(x, y|z) = p(x|z)p(y|z)$$

where z is the conditioning random variable. We observe $x[n] = A + w[n]$ for $n = 0, 1$, where A , $w[0]$, and $w[1]$ are random variables. If $A, w[0], w[1]$ are all independent, prove that $x[0]$ and $x[1]$ are conditionally independent. The conditioning random variable is A . Are $x[0]$ and $x[1]$ independent unconditionally or is it true that

$$p(x[0], x[1]) = p(x[0])p(x[1])?$$

To answer this consider the case where $A, w[0], w[1]$ are independent and each random variable has the PDF $\mathcal{N}(0, 1)$.

- 10.3** The data $x[n]$ for $n = 0, 1, \dots, N-1$ are observed, each sample having the conditional PDF

$$p(x[n]|\theta) = \begin{cases} \exp[-(x[n] - \theta)] & x[n] > \theta \\ 0 & x[n] < \theta, \end{cases}$$

and conditioned on θ the observations are independent. The prior PDF is

$$p(\theta) = \begin{cases} \exp(-\theta) & \theta > 0 \\ 0 & \theta < 0. \end{cases}$$

Find the MMSE estimator of θ . Note: See Problem 10.2 for the definition of conditional independence.

- 10.4** Repeat Problem 10.3 but with the conditional PDF

$$p(x[n]|\theta) = \begin{cases} \frac{1}{\theta} & 0 \leq x[n] \leq \theta \\ 0 & \text{otherwise} \end{cases}$$

and the uniform prior PDF $\theta \sim \mathcal{U}[0, \beta]$. What happens if β is very large so that there is little prior knowledge?

- 10.5** Rederive the MMSE estimator by letting

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= E_{x,\theta} [(\theta - \hat{\theta})^2] \\ &= E_{x,\theta} \left\{ \left[(\theta - E(\theta|\mathbf{x})) + (E(\theta|\mathbf{x}) - \hat{\theta}) \right]^2 \right\} \end{aligned}$$

and evaluating. Hint: Use the result $E_{x,\theta}(\cdot) = E_x[E_{\theta|x}(\cdot)]$.

- 10.6** In Example 10.1 modify the data model as follows:

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $w[n]$ is WGN with variance σ_+^2 if $A \geq 0$ and σ_-^2 if $A < 0$. Find the PDF $p(\mathbf{x}|A)$. Compare this to $p(\mathbf{x}; A)$ for the classical case in which A is deterministic

and $w[n]$ is WGN with variance σ^2 for

a. $\sigma_+^2 = \sigma_-^2$

b. $\sigma_+^2 \neq \sigma_-^2$.

10.7 Plot $\hat{A}$ as given by (10.9) as a function of $\bar{x}$ if $\sqrt{\sigma^2/N} = 1$ for $A_0 = 3$ and $A_0 = 10$. Compare your results to the estimator $\hat{A} = \bar{x}$. Hint: Note that the numerator can be evaluated in closed form and the denominator is related to the cumulative distribution function of a Gaussian random variable.

10.8 A random variable θ has the PDF $p(\theta)$. It is desired to estimate a realization of θ without the availability of any data. To do so a MMSE estimator is proposed that minimizes $E[(\theta - \hat{\theta})^2]$, where the expectation is with respect to $p(\theta)$ only. Prove that the MMSE estimator is $\hat{\theta} = E(\theta)$. Apply your results to Example 10.1 to show that the minimum Bayesian MSE is reduced when data are incorporated into the estimator.

10.9 A quality assurance inspector has the job of monitoring the resistance values of manufactured resistors. He does so by choosing a resistor from a batch and measuring its resistance with an ohmmeter. He knows that the ohmmeter is of poor quality and imparts an error to the measurement which he models as a $\mathcal{N}(0, 1)$ random variable. Hence, he takes N independent measurements. Also, he knows that the resistors should be 100 ohms. Due to manufacturing tolerances, however, they generally are in error by ϵ , where $\epsilon \sim \mathcal{N}(0, 0.011)$. If the inspector chooses a resistor, how many ohmmeter measurements are necessary to ensure that a MMSE estimator of the resistance R yields the correct resistance to 0.1 ohms “on the average” or as he continues to choose resistors throughout the day? How many measurements would he need if he did not have any prior knowledge about the manufacturing tolerances?

10.10 In this problem we discuss reproducing PDFs. Recall that

$$p(\theta|\mathbf{x}) = \frac{p(\mathbf{x}|\theta)p(\theta)}{\int p(\mathbf{x}|\theta)p(\theta) d\theta}$$

where the denominator does not depend on θ . If $p(\theta)$ is chosen so that when multiplied by $p(\mathbf{x}|\theta)$ we obtain the same form of PDF in θ , then the posterior PDF $p(\theta|\mathbf{x})$ will have the same form as $p(\theta)$. Such was the case in Example 10.1 for the Gaussian PDF. Now assume that the PDF of $x[n]$ conditioned on θ is the exponential PDF

$$p(x[n]|\theta) = \begin{cases} \theta \exp(-\theta x[n]) & x[n] > 0 \\ 0 & x[n] < 0 \end{cases}$$




Figure 10.9 Height-weight data

where the $x[n]$'s are conditionally independent (see Problem 10.2). Next, assume the gamma prior PDF

$$p(\theta) = \begin{cases} \frac{\lambda^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} \exp(-\lambda\theta) & \theta > 0 \\ 0 & \theta < 0 \end{cases}$$

where $\lambda > 0, \alpha > 0$, and find the posterior PDF. Compare it to the prior PDF. Such a PDF, in this case the gamma, is termed a *conjugate prior* PDF.

10.11 It is desired to estimate a person's weight based on his height. To see if this is feasible, data were taken for $N = 100$ people to generate the ordered pairs (h, w) , where h denotes the height and w the weight. The data that were obtained appear as shown in Figure 10.9a. Explain how you might be able to guess someone's weight based on his height using a MMSE estimator. What modeling assumptions would you have to make about the data? Next, the same experiment was performed for people on a planet far away. The data obtained are shown in Figure 10.9b. What would the MMSE estimator of weight be now?

10.12 If $[x y]^T \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, where

$$\mathbf{C} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

let $g(y) = p(x_0, y)$ for some $x = x_0$. Prove that $g(y)$ is maximized for $y = \rho x_0$. Also, show that $E(y|x_0) = \rho x_0$. Why are they the same? If $\rho = 0$, what is the MMSE estimator of y based on x ?

10.13 Verify (10.32) and (10.33) by using the matrix inversion lemma. Hint: Verify (10.33) first and use it to verify (10.32).

10.14 The data

$$x[n] = Ar^n + w[n] \quad n = 0, 1, \dots, N-1$$

where r is known, $w[n]$ is WGN with variance σ^2 , and $A \sim \mathcal{N}(0, \sigma_A^2)$ independent of $w[n]$ are observed. Find the MMSE estimator of A as well as the minimum Bayesian MSE.

10.15 A measure of the randomness of a random variable θ is its *entropy* defined as

$$H(\theta) = E(-\ln p(\theta)) = - \int \ln p(\theta) p(\theta) d\theta.$$

If $\theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, find the entropy and relate it to the concentration of the PDF. After observing the data, the entropy of the posterior PDF can be defined as

$$H(\theta|\mathbf{x}) = E(-\ln p(\theta|\mathbf{x})) = - \iint \ln p(\theta|\mathbf{x}) p(\mathbf{x}, \theta) d\mathbf{x} d\theta$$

and should be less than $H(\theta)$. Hence, a measure of the information gained by observing the data is

$$I = H(\theta) - H(\theta|\mathbf{x}).$$

Prove that $I \geq 0$. Under what conditions will $I = 0$? Hint: Express $H(\theta)$ as

$$H(\theta) = - \iint \ln p(\theta) p(\mathbf{x}, \theta) d\mathbf{x} d\theta.$$

You will also need the inequality

$$\int \ln \frac{p_1(\mathbf{u})}{p_2(\mathbf{u})} p_1(\mathbf{u}) d\mathbf{u} \geq 0$$

for PDFs $p_1(\mathbf{u}), p_2(\mathbf{u})$. Equality holds if and only if $p_1(\mathbf{u}) = p_2(\mathbf{u})$. (This approach is described more fully in [Zacks 1981] and relies upon the standard concept of mutual information in information theory [Ash 1965, Gallager 1968].)

10.16 For Example 10.1 show that the information gained by observing the data is

$$I = \frac{1}{2} \ln \left(1 + \frac{\sigma_A^2}{\sigma^2/N} \right).$$

10.17 In choosing a prior PDF that is noninformative or does not assume any prior knowledge, the argument is made that we should do so to gain maximum information from the data. In this way the data are the principal contributor to our state of knowledge about the unknown parameter. Using the results of Problem 10.15, this approach may be implemented by choosing $p(\theta)$ for which I is maximum. For the Gaussian prior PDF for Example 10.1 how should μ_A and σ_A^2 be chosen so that $p(A)$ is noninformative?

Appendix 10A

Derivation of Conditional Gaussian PDF

In this appendix we derive the conditional PDF of a multivariate Gaussian PDF as summarized in Theorem 10.2. The reader will note that the result is the vector extension of that given in Example 10.1. Also, the posterior mean and covariance expressions are valid even if the covariance matrix $\mathbf{C}$ is singular. Such was the case in Appendix 7C in which $\mathbf{x}$ was a linear function of $\mathbf{y}$. The derivation in this appendix, however, assumes an invertible covariance matrix.

Using the notation of Theorem 10.2, we have

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}) &= \frac{p(\mathbf{x}, \mathbf{y})}{p(\mathbf{x})} \\ &= \frac{\frac{1}{(2\pi)^{\frac{k+l}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} \begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix}^T \mathbf{C}^{-1} \begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix} \right]}{\frac{1}{(2\pi)^{\frac{k}{2}} \det^{\frac{1}{2}}(\mathbf{C}_{xx})} \exp \left[-\frac{1}{2} (\mathbf{x} - E(\mathbf{x}))^T \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \right]} \end{aligned}$$

That $p(\mathbf{x})$ can be written this way or that $\mathbf{x}$ is also multivariate Gaussian with the given mean and covariance follows from the fact that $\mathbf{x}$ and $\mathbf{y}$ are jointly Gaussian. To verify this let

$$\mathbf{x} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}$$

and apply the property of a Gaussian vector that a linear transformation also produces a Gaussian vector. We next examine the determinant of the partitioned covariance matrix. Since the determinant of a partitioned matrix may be evaluated as

$$\det \left(\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \right) = \det(\mathbf{A}_{11}) \det(\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12})$$

it follows that

$$\det(\mathbf{C}) = \det(\mathbf{C}_{xx}) \det(\mathbf{C}_{yy} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy})$$

and thus

$$\frac{\det(\mathbf{C})}{\det(\mathbf{C}_{xx})} = \det(\mathbf{C}_{yy} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy}).$$

We thus have that

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi)^{\frac{1}{2}} \det^{\frac{1}{2}}(\mathbf{C}_{yy} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy})} \exp\left(-\frac{1}{2}Q\right)$$

where

$$Q = \begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix}^T \mathbf{C}^{-1} \begin{bmatrix} \mathbf{x} - E(\mathbf{x}) \\ \mathbf{y} - E(\mathbf{y}) \end{bmatrix} - (\mathbf{x} - E(\mathbf{x}))^T \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})).$$

To evaluate Q we use the matrix inversion formula for a partitioned symmetric matrix

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} (\mathbf{A}_{11} - \mathbf{A}_{12} \mathbf{A}_{22}^{-1} \mathbf{A}_{21})^{-1} & -\mathbf{A}_{11}^{-1} \mathbf{A}_{12} (\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12})^{-1} \\ -(\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12})^{-1} \mathbf{A}_{21} \mathbf{A}_{11}^{-1} & (\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12})^{-1} \end{bmatrix}.$$

Using this form, the off-diagonal blocks are transposes of each other so that the inverse matrix must be symmetric. This is because $\mathbf{C}$ is symmetric, and hence $\mathbf{C}^{-1}$ is symmetric. By using the matrix inversion lemma we have

$$(\mathbf{A}_{11} - \mathbf{A}_{12} \mathbf{A}_{22}^{-1} \mathbf{A}_{21})^{-1} = \mathbf{A}_{11}^{-1} + \mathbf{A}_{11}^{-1} \mathbf{A}_{12} (\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12})^{-1} \mathbf{A}_{21} \mathbf{A}_{11}^{-1}$$

so that

$$\mathbf{C}^{-1} = \begin{bmatrix} \mathbf{C}_{xx}^{-1} + \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \mathbf{B}^{-1} \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} & -\mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \mathbf{B}^{-1} \\ -\mathbf{B}^{-1} \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} & \mathbf{B}^{-1} \end{bmatrix}$$

where

$$\mathbf{B} = \mathbf{C}_{yy} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy}.$$

The inverse can be written in factored form as

$$\mathbf{C}^{-1} = \begin{bmatrix} \mathbf{I} & -\mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{xx}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} & \mathbf{I} \end{bmatrix}$$

so that upon letting $\tilde{\mathbf{x}} = \mathbf{x} - E(\mathbf{x})$ and $\tilde{\mathbf{y}} = \mathbf{y} - E(\mathbf{y})$ we have

$$Q = \begin{bmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{y}} \end{bmatrix}^T \begin{bmatrix} \mathbf{I} & -\mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{xx}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{y}} \end{bmatrix} - \tilde{\mathbf{x}}^T \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}}$$

$$\begin{aligned}
&= \begin{bmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{y}} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}} \end{bmatrix}^T \begin{bmatrix} \mathbf{C}_{xx}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^{-1} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{y}} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}} \end{bmatrix} - \tilde{\mathbf{x}}^T \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}} \\
&= (\tilde{\mathbf{y}} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}})^T \mathbf{B}^{-1} (\tilde{\mathbf{y}} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \tilde{\mathbf{x}})
\end{aligned}$$

or finally

$$\begin{aligned}
Q &= [\mathbf{y} - (E(\mathbf{y}) + \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})))]^T \\
&\quad [\mathbf{C}_{yy} - \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy}]^{-1} [\mathbf{y} - (E(\mathbf{y}) + \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})))].
\end{aligned}$$

The mean of the posterior PDF is therefore given by (10.24), and the covariance by (10.25).

Chapter 11

General Bayesian Estimators

11.1 Introduction

Having introduced the Bayesian approach to parameter estimation in the last chapter, we now study more general Bayesian estimators and their properties. To do so, the concept of the Bayesian risk function is discussed. Minimization of this criterion results in a variety of estimators. The ones that we will concentrate on are the MMSE estimator and the maximum a posteriori estimator. These two estimators are the principal ones used in practice. Also, the performance of Bayesian estimators is discussed, leading to the concept of error ellipses. The use of Bayesian estimators in signal processing is illustrated by a deconvolution problem. As a special case we consider in detail the noise filtering problem, in which we use the MMSE estimator in conjunction with the Bayesian linear model to yield the important Wiener filter. In the succeeding chapter we will describe more fully the properties and extensions of the Wiener filter.

11.2 Summary

The Bayes risk is defined in (11.1). For a quadratic cost function the mean of the posterior PDF or the usual MMSE estimator minimizes the risk. A proportional cost function (11.2) results in an optimal estimator which is the median of the posterior PDF. For a “hit-or-miss” cost function (11.3) the mode or maximum location of the posterior PDF is the optimal estimator. The latter is termed the maximum a posteriori (MAP) estimator. The MMSE estimator for a vector parameter is given in (11.10), and the corresponding minimum Bayesian MSE by (11.12). Some examples of the computation of the MAP estimator are included in Section 11.5, followed by the definition for a vector MAP estimator, (11.23) or (11.24). The vector MAP estimator is not a simple extension of the scalar MAP estimator but minimizes a slightly different Bayes risk. It is pointed out that the MMSE and MAP estimators commute over linear transformations, but this property does not carry over to nonlinear ones. The performance of the MMSE estimator is characterized by the PDF of its error (11.26). For the Bayesian linear model this may be determined explicitly from (11.29) and (11.30) and leads to the concept of




Figure 11.1 Examples of cost function

an error ellipse as discussed in Example 11.7. Finally, Theorem 11.1 summarizes the MMSE estimator and its performance for the important Bayesian linear model.

11.3 Risk Functions

Previously, we had derived the MMSE estimator by minimizing $E[(\theta - \hat{\theta})^2]$, where the expectation is with respect to the PDF $p(\mathbf{x}, \theta)$. If we let $\epsilon = \theta - \hat{\theta}$ denote the error of the estimator for a particular realization of $\mathbf{x}$ and θ , and also let $C(\epsilon) = \epsilon^2$, then the MSE criterion minimizes $E[C(\epsilon)]$. The deterministic function $C(\epsilon)$ as shown in Figure 11.1a is termed the *cost function*. It is noted that large errors are particularly costly. Also, the average cost or $E[C(\epsilon)]$ is termed the *Bayes risk* $\mathcal{R}$ or

$$\mathcal{R} = E[C(\epsilon)] \quad (11.1)$$

and measures the performance of a given estimator. If $C(\epsilon) = \epsilon^2$, then the cost function is quadratic and the Bayes risk is just the MSE. Of course, there is no need to restrict ourselves to quadratic cost functions, although from a mathematical tractability standpoint, they are highly desirable. Other possible cost functions are shown in Figures 11.1b and 11.1c. In Figure 11.1b we have

$$C(\epsilon) = |\epsilon|. \quad (11.2)$$

This cost function penalizes errors proportionally. In Figure 11.1c the “hit-or-miss” cost function is displayed. It assigns no cost for small errors and a cost of 1 for all

errors in excess of a threshold error or

$$C(\epsilon) = \begin{cases} 0 & |\epsilon| < \delta \\ 1 & |\epsilon| > \delta \end{cases} \quad (11.3)$$

where $\delta > 0$. If δ is small, we can think of this cost function as assigning the same penalty for any error (a "miss") and no penalty for no error (a "hit"). Note that in all three cases the cost function is symmetric in ϵ , reflecting the implicit assumption that positive errors are just as bad as negative errors. Of course, in general this need not be the case.

We already have seen that the Bayes risk is minimized for a quadratic cost function by the MMSE estimator $\hat{\theta} = E(\theta|\mathbf{x})$. We now determine the optimal estimators for the other cost functions. The Bayes risk $\mathcal{R}$ is

$$\begin{aligned} \mathcal{R} &= E[C(\epsilon)] \\ &= \int \int C(\theta - \hat{\theta}) p(\mathbf{x}, \theta) d\mathbf{x} d\theta \\ &= \int \left[\int C(\theta - \hat{\theta}) p(\theta|\mathbf{x}) d\theta \right] p(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (11.4)$$

As we did for the MMSE case in Chapter 10, we will attempt to minimize the inner integral for each $\mathbf{x}$. By holding $\mathbf{x}$ fixed $\hat{\theta}$ becomes a scalar variable. First, considering the absolute error cost function, we have for the inner integral of (11.4)

$$\begin{aligned} g(\hat{\theta}) &= \int |\theta - \hat{\theta}| p(\theta|\mathbf{x}) d\theta \\ &= \int_{-\infty}^{\hat{\theta}} (\hat{\theta} - \theta) p(\theta|\mathbf{x}) d\theta + \int_{\hat{\theta}}^{\infty} (\theta - \hat{\theta}) p(\theta|\mathbf{x}) d\theta. \end{aligned}$$

To differentiate with respect to $\hat{\theta}$ we make use of Leibnitz's rule:

$$\begin{aligned} \frac{\partial}{\partial u} \int_{\phi_1(u)}^{\phi_2(u)} h(u, v) dv &= \int_{\phi_1(u)}^{\phi_2(u)} \frac{\partial h(u, v)}{\partial u} dv + \frac{d\phi_2(u)}{du} h(u, \phi_2(u)) \\ &\quad - \frac{d\phi_1(u)}{du} h(u, \phi_1(u)). \end{aligned}$$

Letting $h(\hat{\theta}, \theta) = (\hat{\theta} - \theta)p(\theta|\mathbf{x})$ for the first integral, we have

$$h(u, \phi_2(u)) = h(\hat{\theta}, \hat{\theta}) = (\hat{\theta} - \hat{\theta})p(\hat{\theta}|\mathbf{x}) = 0$$

and $d\phi_1(u)/du = 0$ since the lower limit does not depend on u . Similarly, for the second integral the corresponding terms are zero. Hence, we can differentiate the integrand only to yield

$$\frac{dg(\hat{\theta})}{d\hat{\theta}} = \int_{-\infty}^{\hat{\theta}} p(\theta|\mathbf{x}) d\theta - \int_{\hat{\theta}}^{\infty} p(\theta|\mathbf{x}) d\theta = 0$$

or

$$\int_{-\infty}^{\hat{\theta}} p(\theta|\mathbf{x}) d\theta = \int_{\hat{\theta}}^{\infty} p(\theta|\mathbf{x}) d\theta.$$

By definition $\hat{\theta}$ is the *median* of the posterior PDF or the point for which $\Pr\{\theta \leq \hat{\theta}|\mathbf{x}\} = 1/2$.

For the “hit-or-miss” cost function we have $\mathcal{C}(\epsilon) = 1$ for $\epsilon > \delta$ and $\epsilon < -\delta$ or for $\theta > \hat{\theta} + \delta$ and $\theta < \hat{\theta} - \delta$, so that the inner integral in (11.4) is

$$g(\hat{\theta}) = \int_{-\infty}^{\hat{\theta}-\delta} 1 \cdot p(\theta|\mathbf{x}) d\theta + \int_{\hat{\theta}+\delta}^{\infty} 1 \cdot p(\theta|\mathbf{x}) d\theta.$$

But

$$\int_{-\infty}^{\infty} p(\theta|\mathbf{x}) d\theta = 1,$$

yielding

$$g(\hat{\theta}) = 1 - \int_{\hat{\theta}-\delta}^{\hat{\theta}+\delta} p(\theta|\mathbf{x}) d\theta.$$

This is minimized by maximizing

$$\int_{\hat{\theta}-\delta}^{\hat{\theta}+\delta} p(\theta|\mathbf{x}) d\theta.$$

For δ arbitrarily small this is maximized by choosing $\hat{\theta}$ to correspond to the location of the *maximum* of $p(\theta|\mathbf{x})$. The estimator that minimizes the Bayes risk for the “hit-or-miss” cost function is therefore the *mode* (location of the maximum) of the posterior PDF. It is termed the *maximum a posteriori* (MAP) estimator and will be described in more detail later.

In summary, the estimators that minimize the Bayes risk for the cost functions of Figure 11.1 are the *mean*, *median*, and *mode* of the posterior PDF. This is illustrated in Figure 11.2a. For some posterior PDFs these three estimators are identical. A notable example is the Gaussian posterior PDF

$$p(\theta|\mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma_{\theta|\mathbf{x}}^2}} \exp \left[-\frac{1}{2\sigma_{\theta|\mathbf{x}}^2} (\theta - \mu_{\theta|\mathbf{x}})^2 \right].$$

The mean $\mu_{\theta|\mathbf{x}}$ is identical to the median (due to the symmetry) and the mode, as illustrated in Figure 11.2b. (See also Problem 11.2.)

11.4 Minimum Mean Square Error Estimators

In Chapter 10 the MMSE estimator was determined to be $E(\theta|\mathbf{x})$ or the mean of the posterior PDF. For this reason it is also commonly referred to as the *conditional mean*




Figure 11.2 Estimators for different cost functions

estimator. We continue our discussion of this important estimator by first extending it to the vector parameter case and then studying some of its properties.

If θ is a vector parameter of dimension $p \times 1$, then to estimate θ_1 , for example, we may view the remaining parameters as nuisance parameters (see Chapter 10). If $p(\mathbf{x}|\theta)$ is the conditional PDF of the data and $p(\theta)$ the prior PDF of the vector parameter, we may obtain the posterior PDF for θ_1 as

$$p(\theta_1|\mathbf{x}) = \int \cdots \int p(\theta|\mathbf{x}) d\theta_2 \cdots d\theta_p \quad (11.5)$$

where

$$p(\theta|\mathbf{x}) = \frac{p(\mathbf{x}|\theta)p(\theta)}{\int p(\mathbf{x}|\theta)p(\theta) d\theta}. \quad (11.6)$$

Then, by the same reasoning as in Chapter 10 we have

$$\begin{aligned}\hat{\theta}_1 &= E(\theta_1|\mathbf{x}) \\ &= \int \theta_1 p(\theta_1|\mathbf{x}) d\theta_1\end{aligned}$$

or in general

$$\hat{\theta}_i = \int \theta_i p(\theta_i|\mathbf{x}) d\theta_i \quad i = 1, 2, \dots, p. \quad (11.7)$$

This is the MMSE estimator that minimizes

$$E[(\theta_i - \hat{\theta}_i)^2] = \int (\theta_i - \hat{\theta}_i)^2 p(\mathbf{x}, \theta_i) d\mathbf{x} d\theta_i \quad (11.8)$$

or the squared error when averaged with respect to the *marginal* PDF $p(\mathbf{x}, \theta_i)$. Thus, the MMSE estimator for a vector parameter does not entail anything new but only a need to determine the posterior PDF for each parameter. Alternatively, we can express the MMSE estimator for the first parameter from (11.5) as

$$\begin{aligned}\hat{\theta}_1 &= \int \theta_1 p(\theta_1|\mathbf{x}) d\theta_1 \\ &= \int \theta_1 \left[\int \dots \int p(\boldsymbol{\theta}|\mathbf{x}) d\theta_2 \dots d\theta_p \right] d\theta_1 \\ &= \int \theta_1 p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta}\end{aligned}$$

or in general

$$\hat{\theta}_i = \int \theta_i p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \quad i = 1, 2, \dots, p.$$

In vector form we have

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \begin{bmatrix} \int \theta_1 p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \\ \int \theta_2 p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \\ \vdots \\ \int \theta_p p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \end{bmatrix} \\ &= \int \boldsymbol{\theta} p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \quad (11.9)\end{aligned}$$

$$= E(\boldsymbol{\theta}|\mathbf{x}) \quad (11.10)$$

where the expectation is with respect to the posterior PDF of the *vector parameter* or $p(\boldsymbol{\theta}|\mathbf{x})$. Note that the vector MMSE estimator $E(\boldsymbol{\theta}|\mathbf{x})$ *minimizes the MSE for each component of the unknown vector parameter*, or $[\hat{\boldsymbol{\theta}}]_i = [E(\boldsymbol{\theta}|\mathbf{x})]_i$ minimizes $E[(\theta_i - \hat{\theta}_i)^2]$. This follows from the derivation.

As discussed in Chapter 10, the minimum Bayesian MSE for a scalar parameter is the posterior PDF variance when averaged over the PDF of $\mathbf{x}$ (see (10.13)). This is because

$$\text{Bmse}(\hat{\theta}_1) = E[(\theta_1 - \hat{\theta}_1)^2] = \int (\theta_1 - \hat{\theta}_1)^2 p(\mathbf{x}, \theta_1) d\theta_1 d\mathbf{x}$$

and since $\hat{\theta}_1 = E(\theta_1|\mathbf{x})$, we have

$$\begin{aligned} \text{Bmse}(\hat{\theta}_1) &= \int \left[\int (\theta_1 - E(\theta_1|\mathbf{x}))^2 p(\theta_1|\mathbf{x}) d\theta_1 \right] p(\mathbf{x}) d\mathbf{x} \\ &= \int \text{var}(\theta_1|\mathbf{x}) p(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

Now, however, the posterior PDF can be written as

$$p(\theta_1|\mathbf{x}) = \int \cdots \int p(\boldsymbol{\theta}|\mathbf{x}) d\theta_2 \cdots d\theta_p$$

so that

$$\text{Bmse}(\hat{\theta}_1) = \int \left[\int (\theta_1 - E(\theta_1|\mathbf{x}))^2 p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \right] p(\mathbf{x}) d\mathbf{x}. \quad (11.11)$$

The inner integral in (11.11) is the variance of θ_1 for the posterior PDF $p(\boldsymbol{\theta}|\mathbf{x})$. This is just the $[1,1]$ element of $\mathbf{C}_{\theta|\mathbf{x}}$, the covariance matrix of the posterior PDF. Hence, in general we have that the minimum Bayesian MSE is

$$\text{Bmse}(\hat{\theta}_i) = \int [\mathbf{C}_{\theta|\mathbf{x}}]_{ii} p(\mathbf{x}) d\mathbf{x} \quad i = 1, 2, \dots, p \quad (11.12)$$

where

$$\mathbf{C}_{\theta|\mathbf{x}} = E_{\theta|\mathbf{x}} [(\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x}))(\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x}))^T]. \quad (11.13)$$

An example follows.

Example 11.1 - Bayesian Fourier Analysis

We reconsider Example 4.2, but to simplify the calculations we let $M = 1$ so that our data model becomes

$$x[n] = a \cos 2\pi f_0 n + b \sin 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where f_0 is a multiple of $1/N$, excepting 0 or $1/2$ (for which $\sin 2\pi f_0 n$ is identically zero), and $w[n]$ is WGN with variance σ^2 . It is desired to estimate $\boldsymbol{\theta} = [a \ b]^T$. We depart from the classical model by assuming a, b are random variables with prior PDF

$$\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \sigma_\theta^2 \mathbf{I})$$

and $\boldsymbol{\theta}$ is independent of $w[n]$. This type of model is referred to as a *Rayleigh fading sinusoid* [Van Trees 1968] and is frequently used to represent a sinusoid that has

propagated through a dispersive medium (see also Problem 11.6). To find the MMSE estimator we need to evaluate $E(\boldsymbol{\theta}|\mathbf{x})$. The data model is rewritten as

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where

$$\mathbf{H} = \begin{bmatrix} 1 & 0 \\ \cos 2\pi f_0 & \sin 2\pi f_0 \\ \vdots & \vdots \\ \cos[2\pi f_0(N-1)] & \sin[2\pi f_0(N-1)] \end{bmatrix}$$

which is recognized as the Bayesian linear model. From Theorem 10.3 we can obtain the mean as well as the covariance of the posterior PDF. To do so we let

$$\begin{aligned} \mu_{\theta} &= \mathbf{0} \\ \mathbf{C}_{\theta} &= \sigma_{\theta}^2 \mathbf{I} \\ \mathbf{C}_w &= \sigma^2 \mathbf{I} \end{aligned}$$

to obtain

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= E(\boldsymbol{\theta}|\mathbf{x}) = \sigma_{\theta}^2 \mathbf{H}^T (\mathbf{H} \sigma_{\theta}^2 \mathbf{H}^T + \sigma^2 \mathbf{I})^{-1} \mathbf{x} \\ \mathbf{C}_{\theta|x} &= \sigma_{\theta}^2 \mathbf{I} - \sigma_{\theta}^2 \mathbf{H}^T (\mathbf{H} \sigma_{\theta}^2 \mathbf{H}^T + \sigma^2 \mathbf{I})^{-1} \mathbf{H} \sigma_{\theta}^2. \end{aligned}$$

A somewhat more convenient form is given by (10.32) and (10.33) as

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= E(\boldsymbol{\theta}|\mathbf{x}) = \left(\frac{1}{\sigma_{\theta}^2} \mathbf{I} + \mathbf{H}^T \frac{1}{\sigma^2} \mathbf{H} \right)^{-1} \mathbf{H}^T \frac{1}{\sigma^2} \mathbf{x} \\ \mathbf{C}_{\theta|x} &= \left(\frac{1}{\sigma_{\theta}^2} \mathbf{I} + \mathbf{H}^T \frac{1}{\sigma^2} \mathbf{H} \right)^{-1}. \end{aligned}$$

Now, because the columns of $\mathbf{H}$ are orthogonal (due to the choice of frequency), we have (see Example 4.2)

$$\mathbf{H}^T \mathbf{H} = \frac{N}{2} \mathbf{I}$$

and

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= \left(\frac{1}{\sigma_{\theta}^2} \mathbf{I} + \frac{N}{2\sigma^2} \mathbf{I} \right)^{-1} \frac{\mathbf{H}^T \mathbf{x}}{\sigma^2} \\ &= \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_{\theta}^2} + \frac{N}{2\sigma^2}} \mathbf{H}^T \mathbf{x} \end{aligned}$$

or the MMSE estimator is

$$\begin{aligned} \hat{a} &= \frac{1}{1 + \frac{2\sigma^2/N}{\sigma_{\theta}^2}} \left[\frac{2}{N} \sum_{n=0}^{N-1} x[n] \cos 2\pi f_0 n \right] \\ \hat{b} &= \frac{1}{1 + \frac{2\sigma^2/N}{\sigma_{\theta}^2}} \left[\frac{2}{N} \sum_{n=0}^{N-1} x[n] \sin 2\pi f_0 n \right]. \end{aligned}$$

The results differ from the classical case only in the scale factor, and if $\sigma_\theta^2 \gg 2\sigma^2/N$, the two results are identical. This corresponds to little prior knowledge compared to the data knowledge. The posterior covariance matrix is

$$\mathbf{C}_{\theta|x} = \frac{1}{\frac{1}{\sigma_\theta^2} + \frac{N}{2\sigma^2}} \mathbf{I}$$

which does not depend on $\mathbf{x}$. Hence, from (11.12)

$$\begin{aligned} \text{Bmse}(\hat{a}) &= \frac{1}{\frac{1}{\sigma_\theta^2} + \frac{1}{2\sigma^2/N}} \\ \text{Bmse}(\hat{b}) &= \frac{1}{\frac{1}{\sigma_\theta^2} + \frac{1}{2\sigma^2/N}}. \end{aligned}$$

◇

It is interesting to note that in the absence of prior knowledge in the Bayesian linear model the MMSE estimator yields the same form as the MVU estimator for the classical linear model. Many fortuitous circumstances enter into making this so, as described in Problem 11.7. To verify this result note that from (10.32)

$$\hat{\theta} = E(\theta|x) = \mu_\theta + (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} (\mathbf{x} - \mathbf{H}\mu_\theta).$$

For no prior knowledge $\mathbf{C}_\theta^{-1} \rightarrow \mathbf{0}$, and therefore,

$$\hat{\theta} \rightarrow (\mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{x} \quad (11.14)$$

which is recognized as the MVU estimator for the general linear model (see Chapter 4).

The MMSE estimator has several useful properties that will be exploited in our study of Kalman filters in Chapter 13 (see also Problems 11.8 and 11.9). First, it commutes over linear (actually affine) transformations. Assume that we wish to estimate α for

$$\alpha = \mathbf{A}\theta + \mathbf{b} \quad (11.15)$$

where $\mathbf{A}$ is a known $r \times p$ matrix and $\mathbf{b}$ is a known $r \times 1$ vector. Then, α is a random vector for which the MMSE estimator is

$$\hat{\alpha} = E(\alpha|x).$$

Because of the linearity of the expectation operator

$$\begin{aligned} \hat{\alpha} &= E(\mathbf{A}\theta + \mathbf{b}|x) \\ &= \mathbf{A}E(\theta|x) + \mathbf{b} \\ &= \mathbf{A}\hat{\theta} + \mathbf{b}. \end{aligned} \quad (11.16)$$

This holds regardless of the joint PDF $p(\mathbf{x}, \boldsymbol{\theta})$. A second important property focuses on the MMSE estimator based on two data vectors $\mathbf{x}_1, \mathbf{x}_2$. We assume that $\boldsymbol{\theta}, \mathbf{x}_1, \mathbf{x}_2$ are jointly Gaussian and the data vectors are independent. The MMSE estimator is

$$\hat{\boldsymbol{\theta}} = E(\boldsymbol{\theta} | \mathbf{x}_1, \mathbf{x}_2).$$

Letting $\mathbf{x} = [\mathbf{x}_1^T \mathbf{x}_2^T]^T$, we have from Theorem 10.2

$$\hat{\boldsymbol{\theta}} = E(\boldsymbol{\theta} | \mathbf{x}) = E(\boldsymbol{\theta}) + \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}} \mathbf{C}_{\mathbf{x}\mathbf{x}}^{-1} (\mathbf{x} - E(\mathbf{x})). \quad (11.17)$$

Since $\mathbf{x}_1, \mathbf{x}_2$ are independent,

$$\begin{aligned} \mathbf{C}_{\mathbf{x}\mathbf{x}}^{-1} &= \begin{bmatrix} \mathbf{C}_{\mathbf{x}_1\mathbf{x}_1} & \mathbf{C}_{\mathbf{x}_1\mathbf{x}_2} \\ \mathbf{C}_{\mathbf{x}_2\mathbf{x}_1} & \mathbf{C}_{\mathbf{x}_2\mathbf{x}_2} \end{bmatrix}^{-1} \\ &= \begin{bmatrix} \mathbf{C}_{\mathbf{x}_1\mathbf{x}_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_{\mathbf{x}_2\mathbf{x}_2} \end{bmatrix}^{-1} \\ &= \begin{bmatrix} \mathbf{C}_{\mathbf{x}_1\mathbf{x}_1}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_{\mathbf{x}_2\mathbf{x}_2}^{-1} \end{bmatrix} \end{aligned}$$

and also

$$\mathbf{C}_{\boldsymbol{\theta}\mathbf{x}} = E \left[\boldsymbol{\theta} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}^T \right] = \begin{bmatrix} \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_1} & \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_2} \end{bmatrix}.$$

It follows from (11.17) that

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= E(\boldsymbol{\theta}) + \begin{bmatrix} \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_1} & \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_2} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{\mathbf{x}_1\mathbf{x}_1}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_{\mathbf{x}_2\mathbf{x}_2}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 - E(\mathbf{x}_1) \\ \mathbf{x}_2 - E(\mathbf{x}_2) \end{bmatrix} \\ &= E(\boldsymbol{\theta}) + \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_1} \mathbf{C}_{\mathbf{x}_1\mathbf{x}_1}^{-1} (\mathbf{x}_1 - E(\mathbf{x}_1)) + \mathbf{C}_{\boldsymbol{\theta}\mathbf{x}_2} \mathbf{C}_{\mathbf{x}_2\mathbf{x}_2}^{-1} (\mathbf{x}_2 - E(\mathbf{x}_2)). \end{aligned}$$

We may interpret the estimator as composed of the prior estimator $E(\boldsymbol{\theta})$ as well as that due to the independent data sets. The MMSE is seen to have an additivity property for *independent* data sets. This result is useful in deriving the sequential MMSE estimator (see Chapter 13).

Finally, in the jointly Gaussian case the MMSE is linear in the data, as can be seen from (11.17). This will allow us to easily determine the PDF of the error, as described in Section 11.6.

11.5 Maximum A Posteriori Estimators

In the MAP estimation approach we choose $\hat{\boldsymbol{\theta}}$ to maximize the posterior PDF or

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta} | \mathbf{x}).$$

This was shown to minimize the Bayes risk for a “hit-or-miss” cost function. In finding the maximum of $p(\theta|\mathbf{x})$ we observe that

$$p(\theta|\mathbf{x}) = \frac{p(\mathbf{x}|\theta)p(\theta)}{p(\mathbf{x})}$$

so an equivalent maximization is of $p(\mathbf{x}|\theta)p(\theta)$. This is reminiscent of the MLE except for the presence of the prior PDF. Hence, the MAP estimator is

$$\hat{\theta} = \arg \max_{\theta} p(\mathbf{x}|\theta)p(\theta) \quad (11.18)$$

or, equivalently,

$$\hat{\theta} = \arg \max_{\theta} [\ln p(\mathbf{x}|\theta) + \ln p(\theta)]. \quad (11.19)$$

Before extending the MAP estimator to the vector case we give some examples.

Example 11.2 - Exponential PDF

Assume that

$$p(x[n]|\theta) = \begin{cases} \theta \exp(-\theta x[n]) & x[n] > 0 \\ 0 & x[n] < 0 \end{cases}$$

where the $x[n]$'s are conditionally IID, or

$$p(\mathbf{x}|\theta) = \prod_{n=0}^{N-1} p(x[n]|\theta)$$

and the prior PDF is

$$p(\theta) = \begin{cases} \lambda \exp(-\lambda\theta) & \theta > 0 \\ 0 & \theta < 0. \end{cases}$$

Then, the MAP estimator is found by maximizing

$$\begin{aligned} g(\theta) &= \ln p(\mathbf{x}|\theta) + \ln p(\theta) \\ &= \ln \left[\theta^N \exp \left(-\theta \sum_{n=0}^{N-1} x[n] \right) \right] + \ln [\lambda \exp(-\lambda\theta)] \\ &= N \ln \theta - N\theta \bar{x} + \ln \lambda - \lambda\theta \end{aligned}$$

for $\theta > 0$. Differentiating with respect to θ produces

$$\frac{dg(\theta)}{d\theta} = \frac{N}{\theta} - N\bar{x} - \lambda$$

and setting it equal to zero yields the MAP estimator

$$\hat{\theta} = \frac{1}{\bar{x} + \frac{\lambda}{N}}.$$




Figure 11.3 Domination of prior PDF by conditional PDF in MAP estimator

Note that as $N \rightarrow \infty$, $\hat{\theta} \rightarrow 1/\bar{x}$. Also, recall that $E(x[n]|\theta) = 1/\theta$ (see Example 9.2), so that

$$\theta = \frac{1}{E(x[n]|\theta)}$$

confirming the reasonableness of the MAP estimator. Also, if $\lambda \rightarrow 0$ so that, the prior PDF is nearly uniform, we obtain the estimator $1/\bar{x}$. In fact, this is the *Bayesian MLE* (the estimator obtained by maximizing $p(\mathbf{x}|\theta)$) since as $\lambda \rightarrow 0$ we have the situation in Figure 11.3 in which the conditional PDF *dominates* the prior PDF. The maximum of g is then unaffected by the prior PDF. $\diamond$




Example 11.3 - DC Level in WGN - Uniform Prior PDF

Recall the introductory example in Chapter 10. There we discussed the MMSE estimator of A for a DC level in WGN with a uniform prior PDF. The MMSE estimator as given by (10.9) could not be obtained in explicit form due to the need to evaluate the integrals. The posterior PDF was given as

$$p(A|\mathbf{x}) = \begin{cases} \frac{\frac{1}{\sqrt{2\pi\sigma^2/N}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right]}{\int_{-A_0}^{A_0} \frac{1}{\sqrt{2\pi\sigma^2/N}} \exp\left[-\frac{1}{2\frac{\sigma^2}{N}}(A - \bar{x})^2\right] dA} & |A| \leq A_0 \\ 0 & |A| > A_0. \end{cases}$$

In determining the MAP estimator of A we first note that the denominator or normalizing factor does not depend on A . This observation formed the basis for (11.18). The posterior PDF is shown in Figure 11.4 for various values of $\bar{x}$. The dashed curves indicate the portion of the original Gaussian PDF that was truncated to yield the posterior PDF. It should be evident from the figure that the maximum is $\bar{x}$ if $|\bar{x}| \leq A_0$, A_0 if $\bar{x} > A_0$, and $-A_0$ if $\bar{x} < -A_0$. Thus, the MAP estimator is explicitly given as

$$\hat{A} = \begin{cases} -A_0 & \bar{x} < -A_0 \\ \bar{x} & -A_0 \leq \bar{x} \leq A_0 \\ A_0 & \bar{x} > A_0. \end{cases} \quad (11.20)$$

(a) $-A_0 \leq \bar{x} \leq A_0$ (b) $\bar{x} > A_0$ (c) $\bar{x} < -A_0$ **Figure 11.4** Determination of MAP estimator

The “classical estimator” $\bar{x}$ is used unless it falls outside the known constraint interval, in which case we discard the data and rely only on the prior knowledge. Note that the MAP estimator is usually easier to determine since it does not involve any integration but only a maximization. Since $p(\theta)$ is the prior PDF, which we must specify, and $p(\mathbf{x}|\theta)$ is the PDF of the data, which we also assume known in formulating our data model, we see that integration is not required. In the case of the MMSE estimator we saw that integration was a necessary part of the estimator implementation. Of course, the required maximization is not without its potential problems, as evidenced by the difficulties encountered in determining the MLE (see Section 7.7). $\diamond$

To extend the MAP estimator to the vector parameter case in which the posterior PDF is now $p(\boldsymbol{\theta}|\mathbf{x})$, we again employ the result

$$p(\theta_1|\mathbf{x}) = \int \cdots \int p(\boldsymbol{\theta}|\mathbf{x}) d\theta_2 \cdots d\theta_p. \quad (11.21)$$

Then, the MAP estimator is given by

$$\hat{\theta}_1 = \arg \max_{\theta_1} p(\theta_1|\mathbf{x})$$

or in general

$$\hat{\theta}_i = \arg \max_{\theta_i} p(\theta_i|\mathbf{x}) \quad i = 1, 2, \dots, p. \quad (11.22)$$

This estimator minimizes the average "hit-or-miss" cost function

$$\mathcal{R}_i = E[C(\theta_i - \hat{\theta}_i)]$$

for each i , where the expectation is over $p(\mathbf{x}, \theta_i)$.

One of the advantages of the MAP estimator for a scalar parameter is that to numerically determine it we need only maximize $p(\mathbf{x}|\theta)p(\theta)$. No integration is required. This desirable property of the MAP estimator for a scalar parameter does not carry over to the vector parameter case due to the need to obtain $p(\theta_i|\mathbf{x})$ as per (11.21). However, we might propose the following *vector* MAP estimator

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\mathbf{x}) \quad (11.23)$$

in which the posterior PDF for the vector parameter $\boldsymbol{\theta}$ is maximized to find the estimator. Now we no longer need to determine the marginal PDFs, eliminating the integration steps, since, equivalently,

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}) \quad (11.24)$$

That this estimator is not in general the same as (11.22) is illustrated by the example shown in Figure 11.5. Note that $p(\theta_1, \theta_2|\mathbf{x})$ is constant and equal to $1/6$ on the rectangular regions shown (single cross-hatched) and equal to $1/3$ on the square region (double cross-hatched). It is clear from (11.23) that the *vector* MAP estimator, although not unique, is any value of $\boldsymbol{\theta}$ within the square so that $0 < \hat{\theta}_2 < 1$. The MAP estimator of θ_2 is found as the value maximizing $p(\theta_2|\mathbf{x})$. But

$$\begin{aligned} p(\theta_2|\mathbf{x}) &= \int p(\theta_1, \theta_2|\mathbf{x}) d\theta_1 \\ &= \begin{cases} \int_2^3 \frac{1}{3} d\theta_1 & 0 < \theta_2 < 1 \\ \int_0^2 \frac{1}{6} d\theta_1 + \int_3^5 \frac{1}{6} d\theta_1 & 1 < \theta_2 < 2 \end{cases} \\ &= \begin{cases} \frac{1}{3} & 0 < \theta_2 < 1 \\ \frac{2}{3} & 1 < \theta_2 < 2 \end{cases} \end{aligned}$$




Figure 11.5 Comparison of scalar and vector MAP estimators

and is shown in Figure 11.5b. Clearly, the MAP estimator is any value $1 < \hat{\theta}_2 < 2$, which differs from the vector MAP estimator. It can be shown, however, that the vector MAP estimator does indeed minimize a Bayes risk, although a different one than the “hit-or-miss” cost function (see Problem 11.11). We will henceforth refer to the vector MAP estimator as just the MAP estimator.

Example 11.4 - DC Level in WGN - Unknown Variance

We observe

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1.$$

In contrast to the usual example in which only A is to be estimated, we now assume the variance σ^2 of the WGN $w[n]$ is also unknown. The vector parameter is $\theta = [A \ \sigma^2]^T$. We assume the conditional PDF

$$p(\mathbf{x}|A, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n] - A)^2 \right]$$

for the data, so that conditioned on knowing θ , we have a DC level in WGN. The prior PDF is assumed to be

$$\begin{aligned} p(A, \sigma^2) &= p(A|\sigma^2)p(\sigma^2) \\ &= \frac{1}{\sqrt{2\pi\alpha\sigma^2}} \exp \left[-\frac{1}{2\alpha\sigma^2} (A - \mu_A)^2 \right] \frac{\lambda \exp(-\frac{\lambda}{\sigma^2})}{\sigma^4}. \end{aligned}$$

The prior PDF of A conditioned on σ^2 is the usual $\mathcal{N}(\mu_A, \sigma_A^2)$ PDF except that $\sigma_A^2 = \alpha\sigma^2$. The prior PDF for σ^2 is the same one used in Example 10.3 and is a special case of the *inverted gamma* PDF. Note that A and σ^2 are not independent. This prior PDF is the conjugate prior PDF (see Problem 10.10). To find the MAP estimator we need to maximize

$$\begin{aligned} g(A, \sigma^2) &= p(\mathbf{x}|A, \sigma^2)p(A, \sigma^2) \\ &= p(\mathbf{x}|A, \sigma^2)p(A|\sigma^2)p(\sigma^2) \end{aligned}$$

over A and σ^2 . To find the value of A that maximizes g we can equivalently maximize just $p(\mathbf{x}|A, \sigma^2)p(A|\sigma^2)$. But in Example 10.1 we showed that

$$\begin{aligned} h(A) &= p(\mathbf{x}|A, \sigma^2)p(A|\sigma^2) \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}\sqrt{2\pi\sigma_A^2}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n]\right] \exp\left[-\frac{1}{2}Q(A)\right] \end{aligned}$$

where

$$Q(A) = \frac{1}{\sigma_{A|x}^2} (A - \mu_{A|x})^2 - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} + \frac{\mu_A^2}{\sigma_A^2}$$

and

$$\begin{aligned} \mu_{A|x} &= \frac{\frac{N}{\sigma^2} \bar{x} + \frac{\mu_A}{\sigma_A^2}}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \\ \sigma_{A|x}^2 &= \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}}. \end{aligned}$$

Maximizing $h(A)$ over A is equivalent to minimizing $Q(A)$, so that

$$\hat{A} = \mu_{A|x} = \frac{\frac{N}{\sigma^2} \bar{x} + \frac{\mu_A}{\sigma_A^2}}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}}$$

and since $\sigma_A^2 = \alpha\sigma^2$, we have

$$\hat{A} = \frac{N\bar{x} + \frac{\mu_A}{\alpha}}{N + \frac{1}{\alpha}}.$$

Since this does not depend on σ^2 , it is the MAP estimator. We next find the MAP estimator for σ^2 by maximizing $g(\hat{A}, \sigma^2) = h(\hat{A})p(\sigma^2)$. First, note that

$$h(\hat{A}) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}\sqrt{2\pi\sigma_A^2}} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n]\right] \exp\left[-\frac{1}{2}Q(\hat{A})\right]$$

where

$$Q(\hat{A}) = \frac{\mu_A^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2}.$$

Letting $\sigma_A^2 = \alpha\sigma^2$, we have

$$\begin{aligned} Q(\hat{A}) &= \frac{\mu_A^2}{\alpha\sigma^2} - \hat{A}^2 \left(\frac{N}{\sigma^2} + \frac{1}{\alpha\sigma^2} \right) \\ &= \frac{1}{\sigma^2} \left[\underbrace{\frac{\mu_A^2}{\alpha} - \hat{A}^2 \left(N + \frac{1}{\alpha} \right)}_{\xi} \right]. \end{aligned}$$

Therefore, we must maximize over σ^2

$$\begin{aligned} g(\hat{A}, \sigma^2) &= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}} \sqrt{2\pi\alpha\sigma^2}} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} x^2[n] - \frac{1}{2\sigma^2} \xi \right] \frac{\lambda \exp(-\frac{\lambda}{\sigma^2})}{\sigma^4} \\ &= \frac{c}{(\sigma^2)^{\frac{N+5}{2}}} \exp \left(-\frac{a}{\sigma^2} \right) \end{aligned}$$

where c is a constant independent of σ^2 and

$$a = \frac{1}{2} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{2} \xi + \lambda.$$

Differentiating $\ln g$ produces

$$\frac{\partial \ln g(\hat{A}, \sigma^2)}{\partial \sigma^2} = -\frac{\frac{N+5}{2}}{\sigma^2} + \frac{a}{\sigma^4}$$

and setting it to zero yields

$$\begin{aligned} \hat{\sigma}^2 &= \frac{1}{\frac{N+5}{2}} a \\ &= \frac{1}{\frac{N+5}{2}} \left[\frac{1}{2} \sum_{n=0}^{N-1} x^2[n] + \frac{1}{2} \xi + \lambda \right] \\ &= \frac{1}{N+5} \left[\sum_{n=0}^{N-1} x^2[n] + \frac{\mu_A^2}{\alpha} - \hat{A}^2 \left(N + \frac{1}{\alpha} \right) + 2\lambda \right] \\ &= \frac{N}{N+5} \left[\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \hat{A}^2 \right] + \frac{1}{(N+5)\alpha} (\mu_A^2 - \hat{A}^2) + \frac{2\lambda}{N+5}. \end{aligned}$$

The MAP estimator is therefore

$$\begin{aligned} \hat{\theta} &= \begin{bmatrix} \hat{A} \\ \hat{\sigma}^2 \end{bmatrix} \\ &= \begin{bmatrix} \frac{N\bar{x} + \frac{\mu_A}{\alpha}}{N + \frac{1}{\alpha}} \\ \frac{N}{N+5} \left[\frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \hat{A}^2 \right] + \frac{1}{(N+5)\alpha} (\mu_A^2 - \hat{A}^2) + \frac{2\lambda}{N+5} \end{bmatrix}. \end{aligned}$$

As expected, if $N \rightarrow \infty$ so that the conditional or data PDF dominates the prior PDF, we have

$$\begin{aligned} \hat{A} &\rightarrow \bar{x} \\ \hat{\sigma}^2 &\rightarrow \frac{1}{N} \sum_{n=0}^{N-1} x^2[n] - \bar{x}^2 = \frac{1}{N} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 \end{aligned}$$

which is the Bayesian MLE. Recall that the Bayesian MLE is the value of θ that maximizes $p(\mathbf{x}|\theta)$. $\diamond$

It is true in general that as $N \rightarrow \infty$, the MAP estimator becomes the Bayesian MLE. To find the MAP estimator we must maximize $p(\mathbf{x}|\theta)p(\theta)$. If the prior $p(\theta)$ is uniform over the range of θ for which $p(\mathbf{x}|\theta)$ is essentially nonzero (as will be the case if $N \rightarrow \infty$ so that the data PDF dominates the prior PDF), then the maximization will be equivalent to a maximization of $p(\mathbf{x}|\theta)$. If $p(\mathbf{x}|\theta)$ has the same form as the PDF family $p(\mathbf{x};\theta)$ (as it does for the previous example), then the Bayesian MLE and the classical MLE will have the same form. The reader is cautioned, however, that the estimators are inherently different due to the contrasting underlying experiments.

Some properties of the MAP estimator are of importance. First, if the posterior PDF is Gaussian, as in the Bayesian linear model for example, the mode or peak location is identical to the mean. Hence, *the MAP estimator is identical to the MMSE estimator if $\mathbf{x}$ and θ are jointly Gaussian*. Second, the invariance property encountered in maximum likelihood theory *does not hold* for the MAP estimator. The next example illustrates this.

Example 11.5 - Exponential PDF

Assume in Example 11.2 that we wish to estimate $\alpha = 1/\theta$. It might be supposed that the MAP estimator is

$$\hat{\alpha} = \frac{1}{\hat{\theta}}$$

where $\hat{\theta}$ is the MAP estimator of θ , so that

$$\hat{\alpha} = \bar{x} + \frac{\lambda}{N}. \quad (11.25)$$

We now show that *this is not true*. As before

$$p(x[n]|\theta) = \begin{cases} \theta \exp(-\theta x[n]) & x[n] > 0 \\ 0 & x[n] < 0. \end{cases}$$

The conditional PDF based on observing α is

$$p(x[n]|\alpha) = \begin{cases} \frac{1}{\alpha} \exp\left(-\frac{x[n]}{\alpha}\right) & x[n] > 0 \\ 0 & x[n] < 0 \end{cases}$$

since knowing α is equivalent to knowing θ . The prior PDF

$$p(\theta) = \begin{cases} \lambda \exp(-\lambda\theta) & \theta > 0 \\ 0 & \theta < 0 \end{cases}$$

cannot be transformed to $p(\alpha)$ by just letting $\theta = 1/\alpha$. This is because θ is a *random variable*, not a deterministic parameter. The PDF of α must account for the derivative

of the transformation or

$$\begin{aligned} p_{\alpha}(\alpha) &= \frac{p_{\theta}(\theta(\alpha))}{\left| \frac{d\alpha}{d\theta} \right|} \\ &= \begin{cases} \frac{\lambda \exp(-\lambda/\alpha)}{\alpha^2} & \alpha > 0 \\ 0 & \alpha < 0. \end{cases} \end{aligned}$$

Otherwise, the prior PDF would not integrate to 1. Completing the problem

$$\begin{aligned} g(\alpha) &= \ln p(\mathbf{x}|\alpha) + \ln p(\alpha) \\ &= \ln \left[\left(\frac{1}{\alpha} \right)^N \exp \left(-\frac{1}{\alpha} \sum_{n=0}^{N-1} x[n] \right) \right] + \ln \frac{\lambda \exp(-\lambda/\alpha)}{\alpha^2} \\ &= -N \ln \alpha - N \frac{\bar{x}}{\alpha} + \ln \lambda - \frac{\lambda}{\alpha} - 2 \ln \alpha \\ &= -(N+2) \ln \alpha - \frac{N\bar{x} + \lambda}{\alpha} + \ln \lambda. \end{aligned}$$

Differentiating with respect to α

$$\frac{dg}{d\alpha} = -\frac{N+2}{\alpha} + \frac{N\bar{x} + \lambda}{\alpha^2}$$

and setting it equal to zero yields the MAP estimator

$$\hat{\alpha} = \frac{N\bar{x} + \lambda}{N+2}$$

which is not the same as (11.25). Hence, it is seen that that MAP estimator does not commute over nonlinear transformations, although it does so for linear transformations (see Problem 11.12).

◇

11.6 Performance Description

In the classical estimation problem we were interested in the mean and variance of an estimator. Assuming the estimator was Gaussian, we could then immediately obtain the PDF. If the PDF was concentrated about the true value of the parameter, then we could say that the estimator performed well. In the case of a random parameter the same approach cannot be used. Now the randomness of the parameter results in a different PDF of the estimator for each *realization* of θ . We denote this PDF by $p(\hat{\theta}|\theta)$. To perform well the estimate should be close to θ for every possible value of θ or the error

$$\epsilon = \theta - \hat{\theta}$$




Figure 11.6 Dependence of estimation error on realization of parameter

should be small. This is illustrated in Figure 11.6 where several conditional PDFs are shown. Each PDF corresponds to that obtained for a given realization of θ . For $\hat{\theta}$ to be a good estimator $p(\hat{\theta}|\theta)$ should be concentrated about θ for all possible θ . This reasoning led to the MMSE estimator, which minimizes $E_{x,\theta}[(\theta - \hat{\theta})^2]$. For an arbitrary Bayesian estimator it thus makes sense to assess the performance by determining the PDF of the error. This PDF, which now accounts for the randomness of θ , should be concentrated about zero. For the MMSE estimator $\hat{\theta} = E(\theta|\mathbf{x})$, the error is

$$\epsilon = \theta - E(\theta|\mathbf{x})$$

and the mean of the error is

$$\begin{aligned} E_{x,\theta}(\epsilon) &= E_{x,\theta}[\theta - E(\theta|\mathbf{x})] \\ &= E_x[E_{\theta|\mathbf{x}}(\theta) - E_{\theta|\mathbf{x}}(\theta|\mathbf{x})] \\ &= E_x[E(\theta|\mathbf{x}) - E(\theta|\mathbf{x})] = 0 \end{aligned}$$

so that the error of the MMSE estimator is on the average (with respect to $p(\mathbf{x}, \theta)$) zero. The variance of the error for the MMSE estimator is

$$\text{var}(\epsilon) = E_{x,\theta}(\epsilon^2)$$

since $E(\epsilon) = 0$, and therefore

$$\text{var}(\epsilon) = E_{x,\theta}[(\theta - \hat{\theta})^2]$$

which is just the minimum Bayesian MSE. Finally, if ϵ is Gaussian, we have that

$$\epsilon \sim \mathcal{N}(0, \text{Bmse}(\hat{\theta})). \quad (11.26)$$

Example 11.6 - DC Level in WGN - Gaussian Prior PDF

Referring to Example 10.1, we have that

$$\hat{A} = \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{x} + \frac{\frac{\sigma^2}{N}}{\sigma_A^2 + \frac{\sigma^2}{N}} \mu_A$$

and

$$\text{Bmse}(\hat{A}) = \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}}.$$

In this case the error is $\epsilon = A - \hat{A}$. Because $\hat{A}$ depends linearly on $\mathbf{x}$, and $\mathbf{x}$ and A are jointly Gaussian, ϵ is Gaussian. As a result, we have from (11.26)

$$\epsilon \sim \mathcal{N}\left(0, \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}}\right).$$

As $N \rightarrow \infty$, the PDF collapses about zero and the estimator can be said to be *consistent in the Bayesian sense*. Consistency means that for a large enough data record $\hat{A}$ will always be close to the *realization* of A , regardless of the realization value (see also Problem 11.10). $\diamond$

For a vector parameter $\boldsymbol{\theta}$ the error can be defined as

$$\boldsymbol{\epsilon} = \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}.$$

As before, if $\hat{\boldsymbol{\theta}}$ is the MMSE estimator, then $\boldsymbol{\epsilon}$ will have a zero mean. Its covariance matrix is

$$E_{\mathbf{x},\boldsymbol{\theta}}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T) = E_{\mathbf{x},\boldsymbol{\theta}} \left[(\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x})) (\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x}))^T \right].$$

The expectation is with respect to the PDF $p(\mathbf{x}, \boldsymbol{\theta})$. Note that

$$\begin{aligned} [E(\boldsymbol{\theta}|\mathbf{x})]_i &= E(\theta_i|\mathbf{x}) = \int \theta_i p(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} \\ &= \int \theta_i p(\theta_i|\mathbf{x}) d\theta_i \end{aligned}$$

and depends on $\mathbf{x}$ only. Hence, $E(\boldsymbol{\theta}|\mathbf{x})$ is a function of $\mathbf{x}$, as it should be for a valid estimator. As in the scalar parameter case, the diagonal elements of the covariance matrix of the error represent the minimum Bayesian MSE since the $[i, i]$ element is

$$[E_{\mathbf{x},\boldsymbol{\theta}}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T)]_{ii} = \iint [\theta_i - E(\theta_i|\mathbf{x})]^2 p(\mathbf{x}, \boldsymbol{\theta}) d\mathbf{x} d\boldsymbol{\theta}.$$

Integrating with respect to $\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_p$ produces

$$[E_{\mathbf{x},\boldsymbol{\theta}}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T)]_{ii} = \iint [\theta_i - E(\theta_i|\mathbf{x})]^2 p(\mathbf{x}, \theta_i) d\mathbf{x} d\theta_i$$

which is the minimum Bayesian MSE for $\hat{\theta}_i$. Hence, the diagonal elements of the error covariance matrix $E_{\mathbf{x},\boldsymbol{\theta}}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T)$ are just the minimum Bayesian MSEs. The error covariance matrix is sometimes called the *Bayesian mean square error matrix* and is

termed $\mathbf{M}_{\hat{\theta}}$. To determine the PDF of ϵ we now derive the entire covariance matrix. We specialize our discussion to the Bayesian linear model (see Theorem 10.3). Then,

$$\begin{aligned}\mathbf{M}_{\hat{\theta}} &= E_{x,\theta} \left[(\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x})) (\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x}))^T \right] \\ &= E_x E_{\theta|x} \left[(\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x})) (\boldsymbol{\theta} - E(\boldsymbol{\theta}|\mathbf{x}))^T \right] \\ &= E_x (\mathbf{C}_{\theta|x})\end{aligned}$$

where $\mathbf{C}_{\theta|x}$ is the covariance matrix of the posterior PDF $p(\boldsymbol{\theta}|\mathbf{x})$. If $\boldsymbol{\theta}$ and $\mathbf{x}$ are jointly Gaussian, then from (10.25)

$$\mathbf{M}_{\hat{\theta}} = \mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} \quad (11.27)$$

since $\mathbf{C}_{\theta|x}$ does not depend on $\mathbf{x}$. Specializing to the Bayesian linear model, we have from (10.29) (shortening the notation from $\mathbf{C}_{\theta\theta}$ to $\mathbf{C}_{\theta}$)

$$\mathbf{M}_{\hat{\theta}} = \mathbf{C}_{\theta} - \mathbf{C}_{\theta} \mathbf{H}^T (\mathbf{H} \mathbf{C}_{\theta} \mathbf{H}^T + \mathbf{C}_w)^{-1} \mathbf{H} \mathbf{C}_{\theta} \quad (11.28)$$

or from (10.33)

$$\mathbf{M}_{\hat{\theta}} = (\mathbf{C}_{\theta}^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1}. \quad (11.29)$$

Finally, it should be observed that for the Bayesian linear model $\epsilon = \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}$ is Gaussian since from (10.28)

$$\begin{aligned}\epsilon &= \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \\ &= \boldsymbol{\theta} - \boldsymbol{\mu}_{\theta} - \mathbf{C}_{\theta} \mathbf{H}^T (\mathbf{H} \mathbf{C}_{\theta} \mathbf{H}^T + \mathbf{C}_w)^{-1} (\mathbf{x} - \mathbf{H} \boldsymbol{\mu}_{\theta})\end{aligned}$$

and thus ϵ is a linear transformation of $\mathbf{x}, \boldsymbol{\theta}$, which are themselves jointly Gaussian. Hence, the error vector for the MMSE estimator of the parameters of a Bayesian linear model is characterized by

$$\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{M}_{\hat{\theta}}) \quad (11.30)$$

where $\mathbf{M}_{\hat{\theta}}$ is given by (11.28) or by (11.29). An example follows.

Example 11.7 - Bayesian Fourier Analysis (continued)

We now again consider Example 11.1. Recall that

$$\begin{aligned}\mathbf{C}_{\theta} &= \sigma_{\theta}^2 \mathbf{I} \\ \mathbf{C}_w &= \sigma^2 \mathbf{I} \\ \mathbf{H}^T \mathbf{H} &= \frac{N}{2} \mathbf{I}.\end{aligned}$$

Hence, from (11.29)




Figure 11.7 Error ellipses for Bayesian Fourier analysis

$$\begin{aligned}
 \mathbf{M}_{\hat{\theta}} &= (\mathbf{C}_{\theta}^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \\
 &= \left(\frac{1}{\sigma_{\theta}^2} \mathbf{I} + \frac{N}{2\sigma^2} \mathbf{I} \right)^{-1} \\
 &= \left(\frac{1}{\sigma_{\theta}^2} + \frac{N}{2\sigma^2} \right)^{-1} \mathbf{I}.
 \end{aligned}$$

Therefore, the error $\epsilon = [\epsilon_{\hat{a}} \ \epsilon_{\hat{b}}]^T$ has the PDF

$$\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{M}_{\hat{\theta}})$$

where

$$\mathbf{M}_{\hat{\theta}} = \begin{bmatrix} \frac{1}{\frac{1}{\sigma_{\theta}^2} + \frac{1}{2\sigma^2/N}} & 0 \\ 0 & \frac{1}{\frac{1}{\sigma_{\theta}^2} + \frac{1}{2\sigma^2/N}} \end{bmatrix}.$$

The error components are seen to be independent and, furthermore,

$$\begin{aligned}
 \text{Bmse}(\hat{a}) &= [\mathbf{M}_{\hat{\theta}}]_{11} \\
 \text{Bmse}(\hat{b}) &= [\mathbf{M}_{\hat{\theta}}]_{22}
 \end{aligned}$$

in agreement with the results of Example 11.1. Also, for this example the PDF of the error vector has contours of constant probability density that are circular, as shown in Figure 11.7a. One way to summarize the performance of the estimator is by the error

or concentration ellipse (in this case it is a circle). It is in general an ellipse within which the error vector will lie with probability P . Let

$$\epsilon^T \mathbf{M}_\theta^{-1} \epsilon = c^2. \quad (11.31)$$

Then, the probability P that ϵ will lie within the ellipse described by (11.31) is

$$P = \Pr\{\epsilon^T \mathbf{M}_\theta^{-1} \epsilon \leq c^2\}.$$

But $u = \epsilon^T \mathbf{M}_\theta^{-1} \epsilon$ is a χ_2^2 random variable with PDF (see Problem 11.13)

$$p(u) = \begin{cases} \frac{1}{2} \exp(-\frac{u}{2}) & u > 0 \\ 0 & u < 0 \end{cases}$$

so that

$$\begin{aligned} P &= \Pr\{\epsilon^T \mathbf{M}_\theta^{-1} \epsilon \leq c^2\} \\ &= \int_0^{c^2} \frac{1}{2} \exp\left(-\frac{u}{2}\right) du \\ &= 1 - \exp\left(-\frac{c^2}{2}\right) \end{aligned}$$

or

$$\epsilon^T \mathbf{M}_\theta^{-1} \epsilon = 2 \ln \left(\frac{1}{1-P} \right)$$

describes the error ellipse for a probability P . The error vector will lie within this ellipse with probability P . An example is shown in Figure 11.7b for $\sigma_\theta^2 = 1, \sigma^2/N = 1/2$ so that

$$\mathbf{M}_\theta = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}$$

and thus

$$\epsilon^T \epsilon = \ln \left(\frac{1}{1-P} \right).$$

In general, the contours will be elliptical if the minimum MSE matrix is not a scaled identity matrix (see also Problems 11.14 and 11.15). $\diamond$

We now summarize our results in a theorem.

Theorem 11.1 (Performance of the MMSE Estimator for the Bayesian Linear Model) *If the observed data $\mathbf{x}$ can be modeled by the Bayesian linear model of Theorem 10.3, the MMSE estimator is*

$$\hat{\theta} = \mu_\theta + \mathbf{C}_\theta \mathbf{H}^T (\mathbf{H} \mathbf{C}_\theta \mathbf{H}^T + \mathbf{C}_w)^{-1} (\mathbf{x} - \mathbf{H} \mu_\theta) \quad (11.32)$$

$$= \mu_\theta + (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} (\mathbf{x} - \mathbf{H} \mu_\theta). \quad (11.33)$$

The performance of the estimator is measured by the error $\epsilon = \theta - \hat{\theta}$, whose PDF is Gaussian with mean zero and covariance matrix

$$\begin{aligned} \mathbf{C}_\epsilon &= E_{x,\theta}(\epsilon\epsilon^T) \\ &= \mathbf{C}_\theta - \mathbf{C}_\theta \mathbf{H}^T (\mathbf{H} \mathbf{C}_\theta \mathbf{H}^T + \mathbf{C}_w)^{-1} \mathbf{H} \mathbf{C}_\theta \end{aligned} \quad (11.34)$$

$$= (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1}. \quad (11.35)$$

The error covariance matrix is also the minimum MSE matrix $\mathbf{M}_{\hat{\theta}}$, whose diagonal elements yield the minimum Bayesian MSE or

$$\begin{aligned} [\mathbf{M}_{\hat{\theta}}]_{ii} &= [\mathbf{C}_\epsilon]_{ii} \\ &= \text{Bmse}(\hat{\theta}_i). \end{aligned} \quad (11.36)$$


See Problems 11.16 and 11.17 for an application to line fitting.

11.7 Signal Processing Example


There are many signal processing problems that fit the Bayesian linear model. We mention but a few. Consider a communication problem in which we transmit a signal $s(t)$ through a channel with impulse response $h(t)$. At the channel output we observe a noise corrupted waveform as shown in Figure 11.8. The problem is to estimate $s(t)$ over the interval $0 \leq t \leq T_s$. Because the channel will distort and lengthen the signal, we observe $x(t)$ over the longer interval $0 \leq t \leq T$. Such a problem is sometimes referred to as a *deconvolution* problem [Haykin 1991]. We wish to deconvolve $s(t)$ from a noise corrupted version of $s_o(t) = s(t) \star h(t)$, where $\star$ denotes convolution. This problem also arises in seismic processing in which $s(t)$ represents a series of acoustic reflections of an explosively generated signal due to rock inhomogeneities [Robinson and Treitel 1980]. The filter impulse response $h(t)$ models the earth medium. In image processing the same problem arises for two-dimensional signals in which the two-dimensional version of $s(t)$ represents an image, while the two-dimensional version of $h(t)$ models a distortion due to a poor lens, for example [Jain 1989]. To make any progress in this problem we assume that $h(t)$ is *known*. Otherwise, the problem is one of *blind deconvolution*, which is much more difficult [Haykin 1991]. We further assume that $s(t)$ is a *realization of a random process*. This modeling is appropriate for speech, as an example, for which the signal changes with speaker and content. Hence, the Bayesian assumption appears to be a reasonable one. The observed continuous-time data are

$$x(t) = \int_0^{T_s} h(t - \tau) s(\tau) d\tau + w(t) \quad 0 \leq t \leq T. \quad (11.37)$$


It is assumed that $s(t)$ is nonzero over the interval $[0, T_s]$, and $h(t)$ is nonzero over the interval $[0, T_h]$, so that the observation interval is chosen to be $[0, T]$, where $T = T_s + T_h$. We observe the output signal $s_o(t)$ embedded in noise $w(t)$. In converting the problem



(a) System model



(b) Typical signals



(c) Typical data waveform

Figure 11.8 Generic deconvolution problem

to discrete time we will assume that $s(t)$ is essentially bandlimited to B Hz so that the output signal $s_o(t)$ is also bandlimited to B Hz. The continuous-time noise is assumed to be a zero mean WSS Gaussian random process with PSD

$$P_{ww}(F) = \begin{cases} \frac{N_0}{2} & |F| < B \\ 0 & |F| > B. \end{cases}$$

We sample $x(t)$ at times $t = n\Delta$, where $\Delta = 1/(2B)$, to produce the equivalent discrete-time data (see Appendix 11A)

$$x[n] = \sum_{m=0}^{n_s-1} h[n-m]s[m] + w[n] \quad n = 0, 1, \dots, N-1$$

where $x[n] = x(n\Delta)$, $h[n] = \Delta h(n\Delta)$, $s[n] = s(n\Delta)$, and $w[n] = w(n\Delta)$. The discrete-time signal $s[n]$ is nonzero over the interval $[0, n_s - 1]$, and $h[n]$ is nonzero over the interval $[0, n_h - 1]$. The largest integers less than or equal to T_s/Δ and T_h/Δ are denoted by $n_s - 1$ and $n_h - 1$, respectively. Therefore, $N = n_s + n_h - 1$. Finally, it follows that $w[n]$ is WGN with variance $\sigma^2 = N_0B$. This is because the ACF of $w(t)$ is a sinc function and the samples of $w(t)$ correspond to the zeros of the sinc function (see Example 3.13). If we assume that $s(t)$ is a Gaussian process, then $s[n]$ is a discrete-time Gaussian process. In vector-matrix notation we have

$$\underbrace{\begin{bmatrix} x[0] \\ x[1] \\ \vdots \\ x[N-1] \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} h[0] & 0 & \dots & 0 \\ h[1] & h[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h[N-1] & h[N-2] & \dots & h[N-n_s] \end{bmatrix}}_{\mathbf{H}} \underbrace{\begin{bmatrix} s[0] \\ s[1] \\ \vdots \\ s[n_s-1] \end{bmatrix}}_{\boldsymbol{\theta}} + \underbrace{\begin{bmatrix} w[0] \\ w[1] \\ \vdots \\ w[N-1] \end{bmatrix}}_{\mathbf{w}}. \quad (11.38)$$

Note that in $\mathbf{H}$ the elements $h[n] = 0$ for $n > n_h - 1$. This is exactly in the form of the Bayesian linear model with $p = n_s$. Also, $\mathbf{H}$ is lower triangular due to the causal nature of $h(t)$ and thus $h[n]$. We need only specify the mean and covariance of $\boldsymbol{\theta} = \mathbf{s}$ to complete the Bayesian linear model description. If the signal is zero mean (as in speech, for example), we can assume the prior PDF

$$\mathbf{s} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_s).$$

An explicit form for the covariance matrix can be found by assuming the signal is WSS, at least over a short time interval (as in speech, for example). Then,

$$[\mathbf{C}_s]_{ij} = r_{ss}[i-j]$$

where $r_{ss}[k]$ is the ACF. If the PSD is known, then the ACF may be found and therefore $\mathbf{C}_s$ specified. As a result, the MMSE estimator of $\mathbf{s}$ is from (11.32)

$$\hat{\mathbf{s}} = \mathbf{C}_s \mathbf{H}^T (\mathbf{H} \mathbf{C}_s \mathbf{H}^T + \sigma^2 \mathbf{I})^{-1} \mathbf{x}. \quad (11.39)$$

An interesting special case occurs when $\mathbf{H} = \mathbf{I}$, where $\mathbf{I}$ is the $n_s \times n_s$ identity matrix. In this case the channel impulse response is $h[n] = \delta[n]$, so that the Bayesian linear model becomes

$$\mathbf{x} = \boldsymbol{\theta} + \mathbf{w}.$$

In effect the channel is transparent. Note that in the classical linear model we always assumed that $\mathbf{H}$ was $N \times p$ with $N > p$. This was important from a practical viewpoint since for $\mathbf{H} = \mathbf{I}$ the MVU estimator is

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x} \\ &= \mathbf{x}.\end{aligned}$$

In this case there is no averaging because we are estimating as many parameters as observations. In the Bayesian case we can let $N = p$ or even $N < p$ since we have prior knowledge about $\boldsymbol{\theta}$. If $N < p$, (11.39) still exists, while in classical estimation, the MVU estimator does not. We now assume that $\mathbf{H} = \mathbf{I}$. The MMSE estimator of s is

$$\hat{s} = \mathbf{C}_s (\mathbf{C}_s + \sigma^2 \mathbf{I})^{-1} \mathbf{x}. \quad (11.40)$$

Note that the estimator may be written as $\hat{s} = \mathbf{A} \mathbf{x}$, where $\mathbf{A}$ is an $n_s \times n_s$ matrix. The $\mathbf{A}$ matrix is called the *Wiener filter*. We will have more to say about Wiener filters in Chapter 12. As an example, for the scalar case in which we estimate $s[0]$ based on $x[0]$ the Wiener filter becomes

$$\begin{aligned}\hat{s}[0] &= \frac{r_{ss}[0]}{r_{ss}[0] + \sigma^2} x[0] \\ &= \frac{\eta}{\eta + 1} x[0]\end{aligned}$$

where $\eta = r_{ss}[0]/\sigma^2$ is the SNR. For a high SNR we have $\hat{s}[0] \approx x[0]$, while for a low SNR $\hat{s}[0] \approx 0$. A more interesting example can be obtained by assuming that $s[n]$ is a realization of an AR(1) process or

$$s[n] = -a[1]s[n-1] + u[n]$$

where $u[n]$ is WGN with variance σ_u^2 and $a[1]$ is the filter coefficient. The ACF for this process can be shown to be (see Appendix 1)

$$r_{ss}[k] = \frac{\sigma_u^2}{1 - a^2[1]} (-a[1])^{|k|}$$

and the PSD is

$$P_{ss}(f) = \frac{\sigma_u^2}{|1 + a[1] \exp(-j2\pi f)|^2}.$$

For $a[1] < 0$ the PSD is that of a low-pass process, an example of which is shown in Figure 11.9 for $a[1] = -0.95$ and $\sigma_u^2 = 1$. For the same AR parameters a realization of $s[n]$ is shown in Figure 11.10a as the dashed curve. When WGN is added to yield an




Figure 11.9 Power spectral density of AR(1) process




Figure 11.10 Wiener filtering example

SNR of 5 dB, the data $x[n]$ become the solid curve shown in Figure 11.10a. The Wiener filter given by (11.40) will smooth the noise fluctuations, as shown in Figure 11.10b by the solid curve. However, the price paid is that the signal will suffer some smoothing as well. This is a typical tradeoff. Finally, as might be expected, it can be shown that the Wiener filter acts as a low-pass filter (see Problem 11.18).

References

Haykin, S., *Adaptive Filter Theory*, Prentice-Hall, Englewood Cliffs, N.J., 1991.
 Jain, A.K., *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1989.
 Robinson, E.A., S. Treitel, *Geophysical Signal Analysis*, Prentice-Hall, Englewood Cliffs, N.J., 1980.
 Van Trees, H.L., *Detection, Estimation, and Modulation Theory*, Part I, J. Wiley, New York, 1968.

Problems

11.1 The data $x[n]$, $n = 0, 1, \dots, N-1$ are observed having the PDF

$$p(x[n]|\mu) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}(x[n] - \mu)^2\right].$$

The $x[n]$'s are independent when conditioned on μ . The mean μ has the prior PDF

$$\mu \sim \mathcal{N}(\mu_0, \sigma_0^2).$$

Find the MMSE and MAP estimators of μ . What happens as $\sigma_0^2 \rightarrow 0$ and $\sigma_0^2 \rightarrow \infty$?

11.2 For the posterior PDF

$$p(\theta|x) = \frac{\epsilon}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(\theta - x)^2\right] + \frac{1-\epsilon}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(\theta + x)^2\right].$$

Plot the PDF for $\epsilon = 1/2$ and $\epsilon = 3/4$. Next, find the MMSE and MAP estimators for the same values of ϵ .

11.3 For the posterior PDF

$$p(\theta|x) = \begin{cases} \exp[-(\theta - x)] & \theta > x \\ 0 & \theta < x \end{cases}$$

find the MMSE and MAP estimators.

11.4 The data $x[n] = A + w[n]$ for $n = 0, 1, \dots, N-1$ are observed. The unknown parameter A is assumed to have the prior PDF

$$p(A) = \begin{cases} \lambda \exp(-\lambda A) & A > 0 \\ 0 & A < 0 \end{cases}$$

where $\lambda > 0$, and $w[n]$ is WGN with variance σ^2 and is independent of A . Find the MAP estimator of A .

11.5 If $\boldsymbol{\theta}$ and $\mathbf{x}$ are jointly Gaussian so that

$$\begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{x} \end{bmatrix} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C})$$

where

$$\begin{aligned} \boldsymbol{\mu} &= \begin{bmatrix} E(\boldsymbol{\theta}) \\ E(\mathbf{x}) \end{bmatrix} \\ \mathbf{C} &= \begin{bmatrix} \mathbf{C}_{\theta\theta} & \mathbf{C}_{\theta\mathbf{x}} \\ \mathbf{C}_{\mathbf{x}\theta} & \mathbf{C}_{\mathbf{x}\mathbf{x}} \end{bmatrix} \end{aligned}$$

find the MMSE estimator of $\boldsymbol{\theta}$ based on $\mathbf{x}$. Also, determine the minimum Bayesian MSE for each component of $\boldsymbol{\theta}$. Explain what happens if $\mathbf{C}_{\theta\mathbf{x}} = \mathbf{0}$.

11.6 Consider the data model in Example 11.1 for a single sinusoid in WGN. Rewrite the model as

$$x[n] = A \cos(2\pi f_0 n + \phi) + w[n]$$

where

$$\begin{aligned} A &= \sqrt{a^2 + b^2} \\ \phi &= \arctan\left(\frac{-b}{a}\right). \end{aligned}$$

If $\boldsymbol{\theta} = [a \ b]^T \sim \mathcal{N}(\mathbf{0}, \sigma_a^2 \mathbf{I})$, show that the PDF of A is Rayleigh, the PDF of ϕ is $\mathcal{U}[0, 2\pi]$, and that A and ϕ are independent.

11.7 In the classical general linear model the MVU estimator is also efficient and therefore the maximum likelihood method produces it or the MVU estimator is found by maximizing

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}) \right].$$

For the Bayesian linear model the MMSE estimator is identical to the MAP estimator. Argue why in the absence of prior information the MMSE estimator for the Bayesian linear model is identical in form to the MVU estimator for the classical general linear model.

11.8 Consider a parameter $\boldsymbol{\theta}$ which changes with time according to the deterministic relation

$$\boldsymbol{\theta}[n] = \mathbf{A}\boldsymbol{\theta}[n-1] \quad n \geq 1$$

where $\mathbf{A}$ is a known $p \times p$ invertible matrix and $\boldsymbol{\theta}[0]$ is an unknown parameter which is modeled as a random vector. Note that once $\boldsymbol{\theta}[0]$ is specified, so is $\boldsymbol{\theta}[n]$ for $n \geq 1$. Prove that the MMSE estimator of $\boldsymbol{\theta}[n]$ is

$$\hat{\boldsymbol{\theta}}[n] = \mathbf{A}^n \hat{\boldsymbol{\theta}}[0]$$

where $\hat{\theta}[0]$ is the MMSE estimator of $\theta[0]$ or, equivalently,

$$\hat{\theta}[n] = \mathbf{A}\hat{\theta}[n-1].$$

- 11.9** A vehicle starts at an unknown position $(x[0], y[0])$ and moves at a constant velocity according to

$$\begin{aligned}x[n] &= x[0] + v_x n \\ y[n] &= y[0] + v_y n\end{aligned}$$

where v_x, v_y are the velocity components in the x and y directions, respectively, and are both unknown. It is desired to estimate the position of the vehicle for each time n , as well as the velocity. Although the initial position $(x[0], y[0])$ and velocity $[v_x v_y]^T$ are unknown, they can be modeled by a random vector. Show that $\theta[n] = [x[n] y[n] v_x v_y]^T$ can be estimated by the MMSE estimator

$$\hat{\theta}[n] = \mathbf{A}^n \hat{\theta}[0]$$

where $\hat{\theta}[0]$ is the MMSE estimator of $[x[0] y[0] v_x v_y]^T$. Also, determine $\mathbf{A}$.

- 11.10** Show that the MAP estimator in Example 11.2 is consistent in the Bayesian sense or as $N \rightarrow \infty$, $\hat{\theta} \rightarrow \theta$ for any realization θ of the unknown parameter. Give an argument why this may be true in general. Hint: See Example 9.4.

- 11.11** Consider the vector MAP estimator or

$$\hat{\theta} = \arg \max_{\theta} p(\theta | \mathbf{x}).$$

Show that this estimator minimizes the Bayes risk for the cost function

$$C(\epsilon) = \begin{cases} 1 & \|\epsilon\| > \delta \\ 0 & \|\epsilon\| < \delta \end{cases}$$

where $\epsilon = \theta - \hat{\theta}$, $\|\epsilon\|^2 = \sum_{i=1}^p \epsilon_i^2$, and $\delta \rightarrow 0$.

- 11.12** Show that the MAP estimator of

$$\alpha = \mathbf{A}\theta$$

is

$$\hat{\alpha} = \mathbf{A}\hat{\theta}$$

where α is a $p \times 1$ vector and $\mathbf{A}$ is an invertible $p \times p$ matrix. In other words, the MAP estimator commutes over invertible linear transformations.

- 11.13** If $\mathbf{x}$ is a 2×1 random vector with PDF

$$\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}),$$

prove that the PDF of $\mathbf{x}^T \mathbf{C}^{-1} \mathbf{x}$ is that of a χ^2_2 random variable. Hint: Note that $\mathbf{C}^{-1}$ may be factored as $\mathbf{D}^T \mathbf{D}$, where $\mathbf{D}$ is a whitening transformation as described in Section 4.5.

11.14 If ϵ is a 2×1 random vector and $\epsilon \sim \mathcal{N}(0, \mathbf{M}_\epsilon)$, plot the error ellipse for $P = 0.9$ for the following cases:

- a. $\mathbf{M}_\epsilon = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$
 b. $\mathbf{M}_\epsilon = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$
 c. $\mathbf{M}_\epsilon = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}.$

11.15 As a special case of the Bayesian linear model assume that

$$\mathbf{x} = \boldsymbol{\theta} + \mathbf{w}$$

where $\boldsymbol{\theta}$ is a 2×1 random vector and

$$\begin{aligned} \mathbf{C}_w &= \sigma^2 \mathbf{I} \\ \mathbf{C}_\theta &= \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}. \end{aligned}$$

Plot the error ellipse if $\sigma_1^2 = \sigma_2^2$ and $\sigma_2^2 > \sigma_1^2$. What happens if $\sigma_2^2 \gg \sigma_1^2$?

11.16 In fitting a line through experimental data we assume the model

$$x[n] = A + Bn + w[n] \quad -M \leq n \leq M$$

where $w[n]$ is WGN with variance σ^2 . If we have some prior knowledge of the slope B and intercept A such as

$$\begin{bmatrix} A \\ B \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} A_0 \\ B_0 \end{bmatrix}, \begin{bmatrix} \sigma_A^2 & 0 \\ 0 & \sigma_B^2 \end{bmatrix} \right)$$

find the MMSE estimator of A and B as well as the minimum Bayesian MSE. Assume that A, B are independent of $w[n]$. Which parameter will benefit most from the prior knowledge?

11.17 If $s[n] = A + Bn$ for $n = -M, -(M-1), \dots, M$ and $\hat{A}, \hat{B}$ are the MMSE estimators for A and B , respectively, find the MMSE estimator $\hat{s}$ of $s = [s[-M] \ s[-(M-1)] \ \dots \ s[M]]^T$. Next, find the PDF of the error $\epsilon = s - \hat{s}$ if the data model is that given in Problem 11.16. How does the error vary with n and why?

11.18 In this problem we derive the noncausal Wiener filter for estimation of a signal in WGN. We do so by extending the results in Section 11.7. There the MMSE estimator of the signal was given by (11.40). First, show that

$$(\mathbf{C}_s + \sigma^2 \mathbf{I}) \hat{\mathbf{s}} = \mathbf{C}_s \mathbf{x}.$$

Next, assume the signal is WSS, so that

$$\begin{aligned} [\mathbf{C}_s + \sigma^2 \mathbf{I}]_{ij} &= r_{ss}[i-j] + \sigma^2 \delta[i-j] = r_{xx}[i-j] \\ [\mathbf{C}_s]_{ij} &= r_{ss}[i-j]. \end{aligned}$$

As a result, the equations to be solved for $\hat{s}$ can be written as

$$\sum_{j=0}^{N-1} r_{xx}[i-j] \hat{s}[j] = \sum_{j=0}^{N-1} r_{ss}[i-j] x[j]$$

for $i = 0, 1, \dots, N-1$. Next, extend this result to the case where $s[n]$ is to be estimated for $|n| \leq M$ based on $x[n]$ for $|n| \leq M$, and let $M \rightarrow \infty$ to obtain

$$\sum_{j=-\infty}^{\infty} r_{xx}[i-j] \hat{s}[j] = \sum_{j=-\infty}^{\infty} r_{ss}[i-j] x[j]$$

for $-\infty < i < \infty$. Using Fourier transforms, find the frequency response of the Wiener filter $H(f)$ or

$$H(f) = \frac{\hat{S}(f)}{X(f)}.$$

(You may assume that the Fourier transforms of $x[n]$ and $\hat{s}[n]$ exist.) Explain your results.

Appendix 11A

Conversion of Continuous-Time System to Discrete-Time System

A continuous-time signal $s(t)$ is the input to a linear time invariant system with impulse response $h(t)$. The output $s_o(t)$ is

$$s_o(t) = \int_{-\infty}^{\infty} h(t - \tau) s(\tau) d\tau.$$


We assume that $s(t)$ is bandlimited to B Hz. The output signal must therefore also be bandlimited to B Hz and hence we can assume that the frequency response of the system or $\mathcal{F}\{h(t)\}$ is also bandlimited to B Hz. As a result, the original system shown in Figure 11A.1a may be replaced by the one shown in Figure 11A.1b, where the operator shown in the dashed box is a sampler and low-pass filter. The sampling function is

$$p(t) = \sum_{n=-\infty}^{\infty} \delta(t - n\Delta)$$


and the low-pass filter has the frequency response

$$H_{\text{lpf}}(F) = \begin{cases} \Delta & |F| < B \\ 0 & |F| > B. \end{cases}$$


For bandlimited signals with bandwidth B the system within the dashed box does not change the signal. This is because the sampling operation creates replicas of the spectrum of $s(t)$, while the low-pass filter retains only the original spectrum scaled by Δ to cancel the $1/\Delta$ factor introduced by sampling. Next, note that the first low-pass filter can be combined with the system since $h(t)$ is also assumed to be bandlimited to B Hz. The result is shown in Figure 11A.1c. To reconstruct $s_o(t)$ we need only the




(a) Original continuous-time system



(b) Conversion to sampled data system



(c) Sampled data system



$$h[n] = \Delta h(n\Delta)$$

(d) Discrete-time system

Figure 11A.1 Conversion of continuous-time system to discrete-time system

samples $s_o(n\Delta)$. Referring to Figures 11.11b and 11.11c, we have

$$\begin{aligned}s_o(t) &= \int_{-\infty}^{\infty} \Delta h(t-\tau) \sum_{m=-\infty}^{\infty} s(m\Delta) \delta(\tau-m\Delta) d\tau \\ &= \Delta \sum_{m=-\infty}^{\infty} s(m\Delta) \int_{-\infty}^{\infty} h(t-\tau) \delta(\tau-m\Delta) d\tau \\ &= \Delta \sum_{m=-\infty}^{\infty} s(m\Delta) h(t-m\Delta)\end{aligned}$$

so that the output samples are

$$s_o(n\Delta) = \sum_{m=-\infty}^{\infty} \Delta h(n\Delta - m\Delta) s(m\Delta)$$

or letting $s_o[n] = s_o(n\Delta)$, $h[n] = \Delta h(n\Delta)$, and $s[n] = s(n\Delta)$, this becomes

$$s_o[n] = \sum_{m=-\infty}^{\infty} h[n-m] s[m].$$

The continuous-time output signal is obtained by passing the signal

$$\sum_{n=-\infty}^{\infty} s_o[n] \delta(t-n\Delta)$$

through the low-pass filter as shown in Figure 11A.1c. In any practical system we will sample the continuous-time signal $s_o(t)$ to produce our data set $s_o[n]$. Hence, we may replace the continuous-time data $s_o(t)$ by the discrete-time data $s_o[n]$, as shown in Figure 11A.1d, without a loss of information. Note finally that in reality $s(t)$ cannot be bandlimited since it is time-limited. We generally assume, and it is borne out in practice, that the signal is *approximately* bandlimited. The loss of information is usually negligible.

Chapter 12

Linear Bayesian Estimators

12.1 Introduction

The optimal Bayesian estimators discussed in the previous chapter are difficult to determine in closed form, and in practice too computationally intensive to implement. They involve multidimensional integration for the MMSE estimator and multidimensional maximization for the MAP estimator. Although under the jointly Gaussian assumption these estimators are easily found, in general, they are not. When we are unable to make the Gaussian assumption, another approach must be used. To fill this gap we can choose to retain the MMSE criterion but constrain the estimator to be *linear*. Then, an explicit form for the estimator may be determined which depends only on the first two moments of the PDF. In many ways this approach is analogous to the BLUE in classical estimation, and some parallels will become apparent. In practice, this class of estimators, which are generically termed *Wiener filters*, are extensively utilized.

12.2 Summary

The linear estimator is defined by (12.1), and the corresponding Bayesian MSE by (12.2). Minimizing the Bayesian MSE results in the linear MMSE (LMMSE) estimator of (12.6) and the minimum Bayesian MSE of (12.8). The estimator may also be derived using a vector space viewpoint as described in Section 12.4. This approach leads to the important orthogonality principle which says that the error of the linear MMSE estimator must be uncorrelated with the data. The vector LMMSE estimator is given by (12.20), and the minimum Bayesian MSE by (12.21) and (12.22). The estimator commutes over linear transformations (12.23) and has an additivity property (12.24). For a data vector having the Bayesian linear model form (the Bayesian linear model without the Gaussian assumption), the LMMSE estimator and its performance are summarized in Theorem 12.1. If desired, the estimator for the Bayesian linear model form can be implemented sequentially in time using (12.47)–(12.49).

12.3 Linear MMSE Estimation

We begin our discussion by assuming a scalar parameter θ is to be estimated based on the data set $\{x[0], x[1], \dots, x[N-1]\}$ or in vector form $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[N-1]]^T$. The unknown parameter is modeled as the realization of a random variable. We do not assume any specific form for the joint PDF $p(\mathbf{x}, \theta)$, but as we shall see shortly, only a knowledge of the first two moments. That θ may be estimated from $\mathbf{x}$ is due to the assumed statistical dependence of θ on $\mathbf{x}$ as summarized by the joint PDF $p(\mathbf{x}, \theta)$, and in particular, for a linear estimator we rely on the correlation between θ and $\mathbf{x}$. We now consider the class of all linear (actually affine) estimators of the form

$$\hat{\theta} = \sum_{n=0}^{N-1} a_n x[n] + a_N \quad (12.1)$$

and choose the weighting coefficients a_n 's to minimize the Bayesian MSE

$$\text{Bmse}(\hat{\theta}) = E[(\theta - \hat{\theta})^2] \quad (12.2)$$

where the expectation is with respect to the PDF $p(\mathbf{x}, \theta)$. The resultant estimator is termed the *linear minimum mean square error (LMMSE) estimator*. Note that we have included the a_N coefficient to allow for nonzero means of $\mathbf{x}$ and θ . If the means are both zero, then this coefficient may be omitted, as will be shown later.

Before determining the LMMSE estimator we should keep in mind that the estimator will be suboptimal unless the MMSE estimator happens to be linear. Such would be the case, for example, if the Bayesian linear model applied (see Section 10.6). Otherwise, better estimators will exist, although they will be nonlinear (see the introductory example in Section 10.3). Since the LMMSE estimator relies on the correlation between random variables, a parameter uncorrelated with the data cannot be linearly estimated. Consequently, the proposed approach is not always feasible. This is illustrated by the following example. Consider a parameter θ to be estimated based on the single data sample $x[0]$, where $x[0] \sim \mathcal{N}(0, \sigma^2)$. If the parameter to be estimated is the power of the $x[0]$ realization or $\theta = x^2[0]$, then a perfect estimator will be

$$\hat{\theta} = x^2[0]$$

since the minimum Bayesian MSE will be zero. This estimator is clearly nonlinear. If, however, we attempt to use a LMMSE estimator or

$$\hat{\theta} = a_0 x[0] + a_1,$$

then the optimal weighting coefficients a_0 and a_1 can be found by minimizing

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= E[(\theta - \hat{\theta})^2] \\ &= E[(\theta - a_0 x[0] - a_1)^2]. \end{aligned}$$

We differentiate this with respect to a_0 and a_1 and set the results equal to zero to produce

$$\begin{aligned} E[(\theta - a_0 x[0] - a_1)x[0]] &= 0 \\ E(\theta - a_0 x[0] - a_1) &= 0 \end{aligned}$$

or

$$\begin{aligned} a_0 E(x^2[0]) + a_1 E(x[0]) &= E(\theta x[0]) \\ a_0 E(x[0]) + a_1 &= E(\theta). \end{aligned}$$

But $E(x[0]) = 0$ and $E(\theta x[0]) = E(x^3[0]) = 0$, so that

$$\begin{aligned} a_0 &= 0 \\ a_1 &= E(\theta) = E(x^2[0]) = \sigma^2. \end{aligned}$$

Therefore, the LMMSE estimator is $\hat{\theta} = \sigma^2$ and does not depend on the data. This is because θ and $x[0]$ are uncorrelated. The minimum MSE is

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= E[(\theta - \hat{\theta})^2] \\ &= E[(\theta - \sigma^2)^2] \\ &= E[(x^2[0] - \sigma^2)^2] \\ &= E(x^4[0]) - 2\sigma^2 E(x^2[0]) + \sigma^4 \\ &= 3\sigma^4 - 2\sigma^4 + \sigma^4 \\ &= 2\sigma^4 \end{aligned}$$

as opposed to a minimum MSE of zero for the nonlinear estimator $\theta = x^2[0]$. Clearly, the LMMSE estimator is inappropriate for this problem. Problem 12.1 explores how to modify the LMMSE estimator to make it applicable.

We now derive the optimal weighting coefficients for use in (12.1). Substituting (12.1) into (12.2) and differentiating

$$\frac{\partial}{\partial a_N} E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] - a_N \right)^2 \right] = -2E \left[\theta - \sum_{n=0}^{N-1} a_n x[n] - a_N \right].$$

Setting this equal to zero produces

$$a_N = E(\theta) - \sum_{n=0}^{N-1} a_n E(x[n]) \quad (12.3)$$

which as asserted earlier is zero if the means are zero. Continuing, we need to minimize

$$\text{Bmse}(\hat{\theta}) = E \left\{ \left[\sum_{n=0}^{N-1} a_n (x[n] - E(x[n])) - (\theta - E(\theta)) \right]^2 \right\}$$

over the remaining a_n 's, where a_N has been replaced by (12.3). Letting $\mathbf{a} = [a_0 \ a_1 \ \dots \ a_{N-1}]^T$, we have

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= E \left\{ [\mathbf{a}^T (\mathbf{x} - E(\mathbf{x})) - (\theta - E(\theta))]^2 \right\} \\ &= E [\mathbf{a}^T (\mathbf{x} - E(\mathbf{x})) (\mathbf{x} - E(\mathbf{x}))^T \mathbf{a}] - E [\mathbf{a}^T (\mathbf{x} - E(\mathbf{x})) (\theta - E(\theta))] \\ &\quad - E [(\theta - E(\theta)) (\mathbf{x} - E(\mathbf{x}))^T \mathbf{a}] + E [(\theta - E(\theta))^2] \\ &= \mathbf{a}^T \mathbf{C}_{xx} \mathbf{a} - \mathbf{a}^T \mathbf{C}_{x\theta} - \mathbf{C}_{\theta x} \mathbf{a} + C_{\theta\theta} \end{aligned} \quad (12.4)$$

where $\mathbf{C}_{xx}$ is the $N \times N$ covariance matrix of $\mathbf{x}$, and $\mathbf{C}_{\theta x}$ is the $1 \times N$ cross-covariance vector having the property that $\mathbf{C}_{\theta x}^T = \mathbf{C}_{x\theta}$, and $C_{\theta\theta}$ is the variance of θ . Making use of (4.3) we can minimize (12.4) by taking the gradient to yield

$$\frac{\partial \text{Bmse}(\hat{\theta})}{\partial \mathbf{a}} = 2\mathbf{C}_{xx} \mathbf{a} - 2\mathbf{C}_{x\theta}$$

which when set to zero results in

$$\mathbf{a} = \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta}. \quad (12.5)$$

Using (12.3) and (12.5) in (12.1) produces

$$\begin{aligned} \hat{\theta} &= \mathbf{a}^T \mathbf{x} + a_N \\ &= \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} \mathbf{x} + E(\theta) - \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} E(\mathbf{x}) \end{aligned}$$

or finally the *LMMSE estimator* is

$$\hat{\theta} = E(\theta) + \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})). \quad (12.6)$$

Note that it is identical in form to the MMSE estimator for jointly Gaussian $\mathbf{x}$ and θ , as can be verified from (10.24). This is because in the Gaussian case the MMSE estimator happens to be linear, and hence our constraint is automatically satisfied. If the means of θ and $\mathbf{x}$ are zero, then

$$\hat{\theta} = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x}. \quad (12.7)$$

The minimum Bayesian MSE is obtained by substituting (12.5) into (12.4) to yield

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} \mathbf{C}_{xx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} - \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} \\ &\quad - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} + C_{\theta\theta} \\ &= \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} - 2\mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} + C_{\theta\theta} \end{aligned}$$

or finally

$$\text{Bmse}(\hat{\theta}) = C_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta}. \quad (12.8)$$

Again this is identical to that obtained by substituting (10.25) into (11.12). An example follows.

Example 12.1 - DC Level in WGN with Uniform Prior PDF

Consider the introductory example in Chapter 10. The data model is

$$x[n] = A + w[n] \quad n = 0, 1, \dots, N-1$$

where $A \sim \mathcal{U}[-A_0, A_0]$, $w[n]$ is WGN with variance σ^2 , and A and $w[n]$ are independent. We wish to estimate A . The MMSE estimator cannot be obtained in closed form due to the integration required (see (10.9)). Applying the LMMSE estimator, we first note that $E(A) = 0$, and hence $E(x[n]) = 0$. Since $E(\mathbf{x}) = \mathbf{0}$, the covariances are

$$\begin{aligned} \mathbf{C}_{xx} &= E(\mathbf{x}\mathbf{x}^T) \\ &= E[(A\mathbf{1} + \mathbf{w})(A\mathbf{1} + \mathbf{w})^T] \\ &= E(A^2)\mathbf{1}\mathbf{1}^T + \sigma^2\mathbf{I} \\ \mathbf{C}_{\theta x} &= E(A\mathbf{x}^T) \\ &= E[A(A\mathbf{1} + \mathbf{w})^T] \\ &= E(A^2)\mathbf{1}^T \end{aligned}$$

where $\mathbf{1}$ is an $N \times 1$ vector of all ones. Hence, from (12.7)

$$\begin{aligned} \hat{A} &= \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x} \\ &= \sigma_A^2 \mathbf{1}^T (\sigma_A^2 \mathbf{1}\mathbf{1}^T + \sigma^2 \mathbf{I})^{-1} \mathbf{x} \end{aligned}$$

where we have let $\sigma_A^2 = E(A^2)$. But the form of the estimator is identical to that encountered in Example 10.2 if we let $\mu_A = 0$, so that from (10.31)

$$\hat{A} = \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{x}.$$

Since $\sigma_A^2 = E(A^2) = (2A_0^2)/12 = A_0^2/3$, the LMMSE estimator of A is

$$\hat{A} = \frac{\frac{A_0^2}{3}}{\frac{A_0^2}{3} + \frac{\sigma^2}{N}} \bar{x}. \quad (12.9)$$

As opposed to the original MMSE estimator which required integration, we have obtained the LMMSE estimator in closed form. Also, note that we did not really need to know that A was uniformly distributed but only its mean and variance, or that $w[n]$ was Gaussian but only that it is white and its variance. Likewise, independence of A and $\mathbf{w}$ was not required, only that they were uncorrelated. In general, all that is required to determine the LMMSE estimator are the first two moments of $p(\mathbf{x}, \theta)$ or

$$\begin{bmatrix} E(\theta) \\ E(\mathbf{x}) \end{bmatrix}, \begin{bmatrix} C_{\theta\theta} & \mathbf{C}_{\theta x} \\ \mathbf{C}_{x\theta} & \mathbf{C}_{xx} \end{bmatrix}.$$

However, we must realize that the LMMSE of (12.9) will be *suboptimal* since it has been *constrained to be linear*. The optimal estimator for this problem is given by (10.9).

◇




Figure 12.1 Vector space interpretation of random variables

12.4 Geometrical Interpretations

In Chapter 8 we discussed a geometrical interpretation of the LSE based on the concept of a vector space. The LMMSE estimator admits a similar interpretation, although now the “vectors” are random variables. (Actually, both vector spaces are special cases of the more general Hilbert space [Luenberger 1969].) This alternative formulation assumes that θ and $\mathbf{x}$ are zero mean. If they are not, we can always define the zero mean random variables $\theta' = \theta - E(\theta)$ and $\mathbf{x}' = \mathbf{x} - E(\mathbf{x})$, and consider estimation of θ' by a linear function of $\mathbf{x}'$ (see also Problem 12.5). Now we wish to find the a_n 's so that

$$\hat{\theta} = \sum_{n=0}^{N-1} a_n x[n]$$

minimizes

$$\text{Bmse}(\hat{\theta}) = E[(\theta - \hat{\theta})^2].$$

Let us now think of the random variables $\theta, x[0], x[1], \dots, x[N-1]$ as elements in a vector space as shown symbolically in Figure 12.1. The reader may wish to verify that the properties of a vector space are satisfied, such as vector addition and multiplication by a scalar, etc. Since, as is usually the case, θ cannot be perfectly expressed as a linear combination of the $x[n]$'s (if it could, then our estimator would be perfect), we picture θ as only partially lying in the subspace spanned by the $x[n]$'s. We may define the “length” of each vector x as $\|x\| = \sqrt{E(x^2)}$ or the square root of the variance. Longer length vectors are those with larger variances. The zero length vector is the random variable with zero variance or, therefore, the one that is identically zero (actually not random at all). Finally, to complete our description we require the notion of an *inner product* between two vectors. (Recall that if $\mathbf{x}, \mathbf{y}$ are Euclidean vectors in R^3 , then the inner product is $(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \alpha$, where α is the angle between the vectors.) It can be shown that an appropriate definition, i.e., one that satisfies the properties of an inner product between the vectors x and y , is (see Problem 12.4)

$$(x, y) = E(xy). \quad (12.10)$$

With this definition we have that

$$(x, x) = E(x^2) = \|x\|^2, \quad (12.11)$$




Figure 12.2 Orthogonal random variables— y cannot be linearly estimated based on x

consistent with our earlier definition of the length of a vector. Also, we can now define two vectors to be orthogonal if

$$(x, y) = E(xy) = 0. \quad (12.12)$$

Since the vectors are zero mean, this is equivalent to saying that two vectors are *orthogonal* if and only if they are *uncorrelated*. (In R^3 two Euclidean vectors are orthogonal if the angle between them is $\alpha = 90^\circ$, so that $(\mathbf{x}, \mathbf{y}) = \|\mathbf{x}\| \|\mathbf{y}\| \cos \alpha = 0$.) Recalling our discussions from the previous section, this implies that if two vectors are orthogonal, we cannot use one to estimate the other. As shown in Figure 12.2, since there is no component of y along x , we cannot use x to estimate y . Attempting to do so means that we need to find an a so that $\hat{y} = ax$ minimizes the Bayesian MSE, $\text{Bmse}(\hat{y}) = E[(y - \hat{y})^2]$. To find the optimal value of a

$$\begin{aligned} \frac{d \text{Bmse}(\hat{y})}{da} &= \frac{d}{da} E[(y - ax)^2] \\ &= \frac{d}{da} [E(y^2) - 2aE(xy) + a^2E(x^2)] \\ &= -2E(xy) + 2aE(x^2) \\ &= 0, \end{aligned}$$

which yields

$$a = \frac{E(xy)}{E(x^2)} = 0.$$

The LMMSE estimator of y is just $\hat{y} = 0$. Of course, this is a special case of (12.7) where $N = 1$, $\theta = y$, $x[0] = x$, and $C_{\theta x} = 0$.

With these ideas in mind we proceed to determine the LMMSE estimator using the vector space viewpoint. This approach is useful for conceptualization of the LMMSE estimation process and will be used later to derive the sequential LMMSE estimator. As before, we assume that

$$\hat{\theta} = \sum_{n=0}^{N-1} a_n x[n]$$

where $a_N = 0$ due to the zero mean assumption. We wish to estimate θ as a linear combination of $x[0], x[1], \dots, x[N-1]$. The weighting coefficients a_n should be chosen




Figure 12.3 Orthogonality principle for LMMSE estimation

to minimize the MSE

$$\begin{aligned} E[(\theta - \hat{\theta})^2] &= E\left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n]\right)^2\right] \\ &= \left\|\theta - \sum_{n=0}^{N-1} a_n x[n]\right\|^2. \end{aligned}$$

But this means that minimization of the MSE is equivalent to a minimization of the squared length of the *error vector* $\epsilon = \theta - \hat{\theta}$. The error vector is shown in Figure 12.3b for several candidate estimates. Clearly, the length of the error vector is minimized when ϵ is orthogonal to the subspace spanned by $\{x[0], x[1], \dots, x[N-1]\}$. Hence, we require

$$\epsilon \perp x[0], x[1], \dots, x[N-1] \quad (12.13)$$

or by using our definition of orthogonality

$$E[(\theta - \hat{\theta})x[n]] = 0 \quad n = 0, 1, \dots, N-1. \quad (12.14)$$

This is the important *orthogonality principle* or *projection theorem*. It says that in *estimating the realization of a random variable by a linear combination of data samples, the optimal estimator is obtained when the error is orthogonal to each data sample*. Using the orthogonality principle, the weighting coefficients are easily found as

$$E\left[\left(\theta - \sum_{m=0}^{N-1} a_m x[m]\right)x[n]\right] = 0 \quad n = 0, 1, \dots, N-1$$

or

$$\sum_{m=0}^{N-1} a_m E[x[m]x[n]] = E(\theta x[n]) \quad n = 0, 1, \dots, N-1.$$

In matrix form this is

$$\begin{bmatrix} E(x^2[0]) & E(x[0]x[1]) & \dots & E(x[0]x[N-1]) \\ E(x[1]x[0]) & E(x^2[1]) & \dots & E(x[1]x[N-1]) \\ \vdots & \vdots & \ddots & \vdots \\ E(x[N-1]x[0]) & E(x[N-1]x[1]) & \dots & E(x^2[N-1]) \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{N-1} \end{bmatrix} = \begin{bmatrix} E(\theta x[0]) \\ E(\theta x[1]) \\ \vdots \\ E(\theta x[N-1]) \end{bmatrix}. \quad (12.15)$$

These are the *normal equations*. The matrix is recognized as $\mathbf{C}_{xx}$, and the right-hand vector as $\mathbf{C}_{x\theta}$. Therefore,

$$\mathbf{C}_{xx}\mathbf{a} = \mathbf{C}_{x\theta} \quad (12.16)$$

and

$$\mathbf{a} = \mathbf{C}_{xx}^{-1}\mathbf{C}_{x\theta}. \quad (12.17)$$

The LMMSE estimator of θ is

$$\begin{aligned} \hat{\theta} &= \mathbf{a}^T \mathbf{x} \\ &= \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} \mathbf{x} \end{aligned}$$

or finally

$$\hat{\theta} = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x} \quad (12.18)$$

in agreement with (12.7). The minimum Bayesian MSE is the squared length of the error vector or

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= \|\epsilon\|^2 \\ &= \left\| \theta - \sum_{n=0}^{N-1} a_n x[n] \right\|^2 \\ &= E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] \right)^2 \right] \end{aligned}$$

for the a_n 's given by (12.17). But

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] \right) \left(\theta - \sum_{m=0}^{N-1} a_m x[m] \right) \right] \\ &= E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] \right) \theta \right] - E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] \right) \sum_{m=0}^{N-1} a_m x[m] \right] \\ &= E(\theta^2) - \sum_{n=0}^{N-1} a_n E(x[n]\theta) - \sum_{m=0}^{N-1} a_m E \left[\left(\theta - \sum_{n=0}^{N-1} a_n x[n] \right) x[m] \right]. \end{aligned}$$




Figure 12.4 Linear estimation by orthogonal vectors

The last term is zero due to the orthogonality principle, producing

$$\begin{aligned} \text{Bmse}(\hat{\theta}) &= C_{\theta\theta} - \mathbf{a}^T \mathbf{C}_{x\theta} \\ &= C_{\theta\theta} - \mathbf{C}_{x\theta}^T \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} \\ &= C_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta} \end{aligned}$$

in agreement with (12.8).

Many important results can be easily derived when we view the LMMSE estimator in a vector space framework. In Section 12.6 we examine its application to sequential estimation. As a simple example of its utility in conceptualizing the estimation problem, we provide the following illustration.

Example 12.2 - Estimation by Orthogonal Vectors

Assume that $x[0]$ and $x[1]$ are zero mean and uncorrelated with each other. However, they are each correlated with θ . Figure 12.4a illustrates this case. The LMMSE estimator of θ based on $x[0]$ and $x[1]$ is the sum of the projections of θ on $x[0]$ and $x[1]$ as shown in Figure 12.4b or

$$\begin{aligned} \hat{\theta} &= \hat{\theta}_0 + \hat{\theta}_1 \\ &= \left(\theta, \frac{x[0]}{\|x[0]\|} \right) \frac{x[0]}{\|x[0]\|} + \left(\theta, \frac{x[1]}{\|x[1]\|} \right) \frac{x[1]}{\|x[1]\|}. \end{aligned}$$

Each component is the length of the projection $(\theta, x[n]/\|x[n]\|)$ times the unit vector in the $x[n]$ direction. Equivalently, since $\|x[n]\| = \sqrt{\text{var}(x[n])}$ is a constant, this may be written

$$\hat{\theta} = \frac{(\theta, x[0])}{(x[0], x[0])} x[0] + \frac{(\theta, x[1])}{(x[1], x[1])} x[1].$$

By the definition of the inner product (12.10) we have

$$\hat{\theta} = \frac{E(\theta x[0])}{E(x^2[0])} x[0] + \frac{E(\theta x[1])}{E(x^2[1])} x[1]$$

$$\begin{aligned}
&= \begin{bmatrix} E(\theta x[0]) & E(\theta x[1]) \end{bmatrix} \begin{bmatrix} E(x^2[0]) & 0 \\ 0 & E(x^2[1]) \end{bmatrix}^{-1} \begin{bmatrix} x[0] \\ x[1] \end{bmatrix} \\
&= \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x}.
\end{aligned}$$

Clearly, the ease with which this result was obtained is due to the orthogonality of $x[0], x[1]$ or, equivalently, the diagonal nature of $\mathbf{C}_{xx}$. For nonorthogonal data samples the same approach can be used if we first orthogonalize the samples or replace the data by uncorrelated samples that span the same subspace. We will have more to say about this approach in Section 12.6. $\diamond$

12.5 The Vector LMMSE Estimator

The vector LMMSE estimator is a straightforward extension of the scalar one. Now we wish to find the linear estimator that *minimizes the Bayesian MSE for each element*. We assume that

$$\hat{\theta}_i = \sum_{n=0}^{N-1} a_{in} x[n] + a_{iN} \quad (12.19)$$

for $i = 1, 2, \dots, p$ and choose the weighting coefficients to minimize

$$\text{Bmse}(\hat{\theta}_i) = E[(\theta_i - \hat{\theta}_i)^2] \quad i = 1, 2, \dots, p$$

where the expectation is with respect to $p(\mathbf{x}, \theta_i)$. Since we are actually determining p separate estimators, the scalar solution can be applied, and we obtain from (12.6)

$$\hat{\theta}_i = E(\theta_i) + \mathbf{C}_{\theta_i x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \quad i = 1, 2, \dots, p$$

and the minimum Bayesian MSE is from (12.8)

$$\text{Bmse}(\hat{\theta}_i) = C_{\theta_i \theta_i} - \mathbf{C}_{\theta_i x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x \theta_i} \quad i = 1, 2, \dots, p.$$

The scalar LMMSE estimators can be combined into a vector estimator as

$$\begin{aligned}
\hat{\boldsymbol{\theta}} &= \begin{bmatrix} E(\theta_1) \\ E(\theta_2) \\ \vdots \\ E(\theta_p) \end{bmatrix} + \begin{bmatrix} \mathbf{C}_{\theta_1 x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \\ \mathbf{C}_{\theta_2 x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \\ \vdots \\ \mathbf{C}_{\theta_p x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \end{bmatrix} \\
&= \begin{bmatrix} E(\theta_1) \\ E(\theta_2) \\ \vdots \\ E(\theta_p) \end{bmatrix} + \begin{bmatrix} \mathbf{C}_{\theta_1 x} \\ \mathbf{C}_{\theta_2 x} \\ \vdots \\ \mathbf{C}_{\theta_p x} \end{bmatrix} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \\
&= E(\boldsymbol{\theta}) + \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \quad (12.20)
\end{aligned}$$

where now $\mathbf{C}_{\theta x}$ is a $p \times N$ matrix. By a similar approach we find that the Bayesian MSE matrix is (see Problem 12.7)

$$\begin{aligned}\mathbf{M}_{\hat{\theta}} &= E[(\theta - \hat{\theta})(\theta - \hat{\theta})^T] \\ &= \mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta}\end{aligned}\quad (12.21)$$

where $\mathbf{C}_{\theta\theta}$ is the $p \times p$ covariance matrix. Consequently, the minimum Bayesian MSE is (see Problem 12.7)

$$\text{Bmse}(\hat{\theta}_i) = [\mathbf{M}_{\hat{\theta}}]_{ii}. \quad (12.22)$$

Of course, these results are identical to those for the MMSE estimator in the Gaussian case, for which the estimator is linear. Note that to determine the LMMSE estimator we require only the first two moments of the PDF.

Two properties of the LMMSE estimator are particularly useful. The first one states that the LMMSE estimator commutes over linear (actually affine) transformations. This is to say that if

$$\alpha = \mathbf{A}\theta + \mathbf{b}$$

then the LMMSE estimator of α is

$$\hat{\alpha} = \mathbf{A}\hat{\theta} + \mathbf{b} \quad (12.23)$$

with $\hat{\theta}$ given by (12.20). The second property states that the LMMSE estimator of a sum of unknown parameters is the sum of the individual estimators. Specifically, if we wish to estimate $\alpha = \theta_1 + \theta_2$, then

$$\hat{\alpha} = \hat{\theta}_1 + \hat{\theta}_2 \quad (12.24)$$

where

$$\begin{aligned}\hat{\theta}_1 &= E(\theta_1) + \mathbf{C}_{\theta_1 x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})) \\ \hat{\theta}_2 &= E(\theta_2) + \mathbf{C}_{\theta_2 x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})).\end{aligned}$$

The proof of these properties is left to the reader as an exercise (see Problem 12.8).

In analogy with the BLUE there is a corresponding Gauss-Markov theorem for the Bayesian case. It asserts that for data having the Bayesian linear model *form*, i.e., the Bayesian linear model *without the Gaussian assumption*, an optimal linear estimator exists. Optimality is measured by the Bayesian MSE. The theorem is just the application of our LMMSE estimator to the Bayesian linear model. More specifically, the data are assumed to be

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{w}$$

where θ is a random vector to be estimated and has mean $E(\theta)$ and covariance $\mathbf{C}_{\theta\theta}$, $\mathbf{H}$ is a known observation matrix, $\mathbf{w}$ is a random vector with zero mean and covariance $\mathbf{C}_w$,

and θ and $\mathbf{w}$ are uncorrelated. Then, the LMMSE estimator of θ is given by (12.20), where

$$\begin{aligned} E(\mathbf{x}) &= \mathbf{H}E(\theta) \\ \mathbf{C}_{xx} &= \mathbf{H}\mathbf{C}_{\theta\theta}\mathbf{H}^T + \mathbf{C}_w \\ \mathbf{C}_{\theta x} &= \mathbf{C}_{\theta\theta}\mathbf{H}^T. \end{aligned}$$

(See Section 10.6 for details.) This is summarized by the Bayesian Gauss-Markov theorem.

Theorem 12.1 (Bayesian Gauss-Markov Theorem) *If the data are described by the Bayesian linear model form*

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{w} \quad (12.25)$$

where $\mathbf{x}$ is an $N \times 1$ data vector, $\mathbf{H}$ is a known $N \times p$ observation matrix, θ is a $p \times 1$ random vector of parameters whose realization is to be estimated and has mean $E(\theta)$ and covariance matrix $\mathbf{C}_{\theta\theta}$, and $\mathbf{w}$ is an $N \times 1$ random vector with zero mean and covariance matrix $\mathbf{C}_w$ and is uncorrelated with θ (the joint PDF $p(\mathbf{w}, \theta)$ is otherwise arbitrary), then the LMMSE estimator of θ is

$$\hat{\theta} = E(\theta) + \mathbf{C}_{\theta\theta}\mathbf{H}^T(\mathbf{H}\mathbf{C}_{\theta\theta}\mathbf{H}^T + \mathbf{C}_w)^{-1}(\mathbf{x} - \mathbf{H}E(\theta)) \quad (12.26)$$

$$= E(\theta) + (\mathbf{C}_{\theta\theta}^{-1} + \mathbf{H}^T\mathbf{C}_w^{-1}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{C}_w^{-1}(\mathbf{x} - \mathbf{H}E(\theta)). \quad (12.27)$$

The performance of the estimator is measured by the error $\epsilon = \theta - \hat{\theta}$ whose mean is zero and whose covariance matrix is

$$\begin{aligned} \mathbf{C}_\epsilon &= E_{x,\theta}(\epsilon\epsilon^T) \\ &= \mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta\theta}\mathbf{H}^T(\mathbf{H}\mathbf{C}_{\theta\theta}\mathbf{H}^T + \mathbf{C}_w)^{-1}\mathbf{H}\mathbf{C}_{\theta\theta} \end{aligned} \quad (12.28)$$

$$= (\mathbf{C}_{\theta\theta}^{-1} + \mathbf{H}^T\mathbf{C}_w^{-1}\mathbf{H})^{-1}. \quad (12.29)$$

The error covariance matrix is also the minimum MSE matrix $\mathbf{M}_{\hat{\theta}}$ whose diagonal elements yield the minimum Bayesian MSE

$$\begin{aligned} [\mathbf{M}_{\hat{\theta}}]_{ii} &= [\mathbf{C}_\epsilon]_{ii} \\ &= \text{Bmse}(\hat{\theta}_i). \end{aligned} \quad (12.30)$$

These results are identical to those in Theorem 11.1 for the Bayesian linear model except that the error vector is not necessarily Gaussian. An example of the determination of this estimator and its minimum Bayesian MSE has already been given in Section 10.6. The Bayesian Gauss-Markov theorem states that within the class of linear estimators the one that minimizes the Bayesian MSE for each element of θ is given by (12.26) or (12.27). It will not be optimal unless the conditional expectation $E(\theta|\mathbf{x})$ happens to be linear. Such was the case for the jointly Gaussian PDF. Although suboptimal, the LMMSE estimator is in practice quite useful, being available in closed form and depending only on the means and covariances.

12.6 Sequential LMMSE Estimation

In Chapter 8 we discussed the sequential LS procedure as the process of updating in time the LSE estimator as new data become available. An analogous procedure can be used for LMMSE estimators. That this is possible follows from the vector space viewpoint. In Example 12.2 the LMMSE estimator was obtained by adding to the old estimate $\hat{\theta}_0$, that based on $x[0]$, the estimate $\hat{\theta}_1$, that based on the new datum $x[1]$. This was possible because $x[0]$ and $x[1]$ were orthogonal. When they are not, the algebra becomes more tedious, although the approach is similar. Before stating the general results we will illustrate the basic operations involved by considering the DC level in white noise, which has the Bayesian linear model form. The derivation will be purely algebraic. Then we will repeat the derivation but appeal to the vector space approach. The reason for doing so will be to “lay the groundwork” for the Kalman filter in Chapter 13, which will be derived using the same approach. We will assume a zero mean for the DC level A , so that the vector space approach is applicable.

To begin the algebraic derivation we have from Example 10.1 with $\mu_A \approx 0$ (see (10.11))

$$\hat{A}[N-1] = \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{x}$$

where $\hat{A}[N-1]$ denotes the LMMSE estimator based on $\{x[0], x[1], \dots, x[N-1]\}$. This is because the LMMSE estimator is identical in form to the MMSE estimator for the Gaussian case. Also, we note that from (10.14)

$$\text{Bmse}(\hat{A}[N-1]) = \frac{\sigma_A^2 \sigma^2}{N \sigma_A^2 + \sigma^2}. \quad (12.31)$$

To update our estimator as $x[N]$ becomes available

$$\begin{aligned} \hat{A}[N] &= \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N+1}} \frac{1}{N+1} \sum_{n=0}^N x[n] \\ &= \frac{N \sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} \frac{1}{N} \left(\sum_{n=0}^{N-1} x[n] + x[N] \right) \\ &= \frac{N \sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} \frac{\sigma_A^2 + \frac{\sigma^2}{N}}{\sigma_A^2} \hat{A}[N-1] + \frac{\sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} x[N] \\ &= \frac{N \sigma_A^2 + \sigma^2}{(N+1) \sigma_A^2 + \sigma^2} \hat{A}[N-1] + \frac{\sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} x[N] \\ &= \hat{A}[N-1] + \left(\frac{N \sigma_A^2 + \sigma^2}{(N+1) \sigma_A^2 + \sigma^2} - 1 \right) \hat{A}[N-1] + \frac{\sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} x[N] \\ &= \hat{A}[N-1] + \frac{\sigma_A^2}{(N+1) \sigma_A^2 + \sigma^2} (x[N] - \hat{A}[N-1]). \end{aligned} \quad (12.32)$$

Similar to the sequential LS estimator we correct the old estimate $\hat{A}[N-1]$ by a scaled version of the prediction error $x[N] - \hat{A}[N-1]$. The scaling or gain factor is from (12.32) and (12.31)

$$\begin{aligned} K[N] &= \frac{\sigma_A^2}{(N+1)\sigma_A^2 + \sigma^2} \\ &= \frac{\text{Bmse}(\hat{A}[N-1])}{\text{Bmse}(\hat{A}[N-1]) + \sigma^2} \end{aligned} \quad (12.33)$$

and decreases to zero as $N \rightarrow \infty$, reflecting the increased confidence in the old estimator. We can also update the minimum Bayesian MSE since from (12.31)

$$\begin{aligned} \text{Bmse}(\hat{A}[N]) &= \frac{\sigma_A^2 \sigma^2}{(N+1)\sigma_A^2 + \sigma^2} \\ &= \frac{N\sigma_A^2 + \sigma^2}{(N+1)\sigma_A^2 + \sigma^2} \frac{\sigma_A^2 \sigma^2}{N\sigma_A^2 + \sigma^2} \\ &= (1 - K[N])\text{Bmse}(\hat{A}[N-1]). \end{aligned}$$

Summarizing our results we have the following sequential LMMSE estimator.

Estimator Update:

$$\hat{A}[N] = \hat{A}[N-1] + K[N](x[N] - \hat{A}[N-1]) \quad (12.34)$$

where

$$K[N] = \frac{\text{Bmse}(\hat{A}[N-1])}{\text{Bmse}(\hat{A}[N-1]) + \sigma^2}. \quad (12.35)$$

Minimum MSE Update:

$$\text{Bmse}(\hat{A}[N]) = (1 - K[N])\text{Bmse}(\hat{A}[N-1]). \quad (12.36)$$

We now use the vector space viewpoint to derive the same results. (The general sequential LMMSE estimator is derived in Appendix 12A using the vector space approach. It is a straightforward extension of the following derivation.) Assume that the LMMSE estimator $\hat{A}[1]$ is to be found, which is based on the data $\{x[0], x[1]\}$ as shown in Figure 12.5a. Because $x[0]$ and $x[1]$ are not orthogonal, we cannot simply add the estimate based on $x[0]$ to the estimate based on $x[1]$. If we did, we would be adding in an extra component along $x[0]$. However, we can form $\hat{A}[1]$ as the sum of $\hat{A}[0]$ and a component *orthogonal* to $\hat{A}[0]$ as shown in Figure 12.5b. That component is $\Delta\hat{A}[1]$. To find a vector in the direction of this component recall that the LMMSE estimator has the property that the error is orthogonal to the data. Hence, if we find the LMMSE estimator of $x[1]$ based on $x[0]$, call it $\hat{x}[1|0]$, the error $x[1] - \hat{x}[1|0]$ will be orthogonal to $x[0]$. This




Figure 12.5 Sequential estimation using vector space approach

is shown in Figure 12.5c. Since $x[0]$ and $x[1] - \hat{x}[1|0]$ are orthogonal, we can project A onto each vector separately and add the results, so that

$$\hat{A}[1] = \hat{A}[0] + \Delta \hat{A}[1].$$

To find the correction term $\Delta \hat{A}[1]$ first note that the LMMSE estimator of a random variable y based on x , where both are zero mean, is from (12.7)

$$\hat{y} = \frac{E(xy)}{E(x^2)} x. \quad (12.37)$$

Thus, we have that

$$\begin{aligned} \hat{x}[1|0] &= \frac{E(x[0]x[1])}{E(x^2[0])} x[0] \\ &= \frac{E[(A + w[0])(A + w[1])]}{E[(A + w[0])^2]} x[0] \\ &= \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2} x[0] \end{aligned} \quad (12.38)$$

and the error vector $\tilde{x}[1] = x[1] - \hat{x}[1|0]$ represents the new information that $x[1]$ contributes to the estimation of A . As such, it is called the *innovation*. The projection of A along this error vector is the desired correction

$$\begin{aligned}\Delta\hat{A}[1] &= \left(A, \frac{\tilde{x}[1]}{\|\tilde{x}[1]\|} \right) \frac{\tilde{x}[1]}{\|\tilde{x}[1]\|} \\ &= \frac{E(A\tilde{x}[1])\tilde{x}[1]}{E(\tilde{x}^2[1])}.\end{aligned}$$

If we let $K[1] = E(A\tilde{x}[1])/E(\tilde{x}^2[1])$, we have for the LMMSE estimator

$$\hat{A}[1] = \hat{A}[0] + K[1](x[1] - \hat{x}[1|0]).$$

To evaluate $\hat{x}[1|0]$ we note that $x[1] = A + w[1]$. Hence, from the additive property (12.24) $\hat{x}[1|0] = \hat{A}[0] + \hat{w}[1|0]$. Since $w[1]$ is uncorrelated with $w[0]$, we have $\hat{w}[1|0] = 0$. Finally, then $\hat{x}[1|0] = \hat{A}[0]$, and $\hat{A}[0]$ is found as

$$\hat{A}[0] = \frac{E(Ax[0])}{E(x^2[0])}x[0] = \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2}x[0].$$

Thus,

$$\hat{A}[1] = \hat{A}[0] + K[1](x[1] - \hat{A}[0]).$$

It remains only to determine the gain $K[1]$ for the correction. Since

$$\begin{aligned}\tilde{x}[1] &= x[1] - \hat{x}[1|0] \\ &= x[1] - \hat{A}[0] \\ &= x[1] - \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2}x[0]\end{aligned}$$

the gain becomes

$$\begin{aligned}K[1] &= \frac{E\left[A\left(x[1] - \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2}x[0]\right)\right]}{E\left[\left(x[1] - \frac{\sigma_A^2}{\sigma_A^2 + \sigma^2}x[0]\right)^2\right]} \\ &= \frac{\sigma_A^2}{2\sigma_A^2 + \sigma^2}\end{aligned}$$

which agrees with (12.33) for $N = 1$. We can continue this procedure to find $\hat{A}[2]$, $\hat{A}[3]$, ... In general, we

1. find the LMMSE estimator of A based on $x[0]$, yielding $\hat{A}[0]$
2. find the LMMSE estimator of $x[1]$ based on $x[0]$, yielding $\hat{x}[1|0]$

3. determine the innovation of the new datum $x[1]$, yielding $x[1] - \hat{x}[1|0]$
4. add to $\hat{A}[0]$ the LMMSE estimator of A based on the innovation, yielding $\hat{A}[1]$
5. continue the process.

In essence we are generating a set of *uncorrelated or orthogonal* random variables, namely, the innovations $\{x[0], x[1] - \hat{x}[1|0], x[2] - \hat{x}[2|0, 1], \dots\}$. This procedure is termed a Gram-Schmidt orthogonalization (see Problem 12.10). To find the LMMSE estimator of A based on $\{x[0], x[1], \dots, x[N-1]\}$ we simply *add* the individual estimators to yield

$$\hat{A}[N-1] = \sum_{n=0}^{N-1} K[n] (x[n] - \hat{x}[n|0, 1, \dots, n-1]) \quad (12.39)$$

where each gain factor is

$$K[n] = \frac{E[A(x[n] - \hat{x}[n|0, 1, \dots, n-1])]}{E[(x[n] - \hat{x}[n|0, 1, \dots, n-1])^2]} \quad (12.40)$$

This simple form is due to the uncorrelated property of the innovations. In sequential form (12.39) becomes

$$\hat{A}[N] = \hat{A}[N-1] + K[N](x[N] - \hat{x}[N|0, 1, \dots, N-1]). \quad (12.41)$$

To complete the derivation we must find $\hat{x}[N|0, 1, \dots, N-1]$ and the gain $K[N]$. We first use the additive property (12.24). Since A and $w[N]$ are uncorrelated and zero mean, the LMMSE estimator of $x[N] = A + w[N]$ is

$$\hat{x}[N|0, 1, \dots, N-1] = \hat{A}[N|0, 1, \dots, N-1] + \hat{w}[N|0, 1, \dots, N-1].$$

But $\hat{A}[N|0, 1, \dots, N-1]$ is by definition the LMMSE estimator of the realization of A at time N based on $\{x[0], x[1], \dots, x[N-1]\}$ or $\hat{A}[N-1]$ (recalling that A does not change over the observation interval). Also, $\hat{w}[N|0, 1, \dots, N-1]$ is zero since $w[N]$ is uncorrelated with the previous data samples since A and $w[n]$ are uncorrelated and $w[n]$ is white noise. Thus,

$$\hat{x}[N|0, 1, \dots, N-1] = \hat{A}[N-1]. \quad (12.42)$$

To find the gain we use (12.40) and (12.42):

$$K[N] = \frac{E[A(x[N] - \hat{A}[N-1])]}{E[(x[N] - \hat{A}[N-1])^2]}.$$

But

$$E[A(x[N] - \hat{A}[N-1])] = E[(A - \hat{A}[N-1])(x[N] - \hat{A}[N-1])] \quad (12.43)$$

since $x[N] - \hat{A}[N-1]$ is the innovation of $x[N]$ which is orthogonal to $\{x[0], x[1], \dots, x[N-1]\}$ and hence to $\hat{A}[N-1]$ (a linear combination of these data samples). Also, $E[w[N](A - \hat{A}[N-1])] = 0$ as explained previously, so that

$$\begin{aligned} E[A(x[N] - \hat{A}[N-1])] &= E[(A - \hat{A}[N-1])^2] \\ &= \text{Bmse}(\hat{A}[N-1]). \end{aligned} \quad (12.44)$$

Also,

$$\begin{aligned} E[(x[N] - \hat{A}[N-1])^2] &= E[(w[N] + A - \hat{A}[N-1])^2] \\ &= E(w^2[N]) + E[(A - \hat{A}[N-1])^2] \\ &= \sigma^2 + \text{Bmse}(\hat{A}[N-1]) \end{aligned} \quad (12.45)$$

so that finally

$$K[N] = \frac{\text{Bmse}(\hat{A}[N-1])}{\sigma^2 + \text{Bmse}(\hat{A}[N-1])}. \quad (12.46)$$

The minimum MSE update can be obtained as

$$\begin{aligned} \text{Bmse}(\hat{A}[N]) &= E[(A - \hat{A}[N])^2] \\ &= E\left[(A - \hat{A}[N-1] - K[N](x[N] - \hat{A}[N-1]))^2\right] \\ &= E\left[(A - \hat{A}[N-1])^2\right] - 2K[N]E[(A - \hat{A}[N-1])(x[N] - \hat{A}[N-1])] \\ &\quad + K^2[N]E[(x[N] - \hat{A}[N-1])^2] \\ &= \text{Bmse}(\hat{A}[N-1]) - 2K[N]\text{Bmse}(\hat{A}[N-1]) \\ &\quad + K^2[N](\sigma^2 + \text{Bmse}(\hat{A}[N-1])) \end{aligned}$$

where we have used (12.43)–(12.45). Using (12.46), we have

$$\text{Bmse}(\hat{A}[N]) = (1 - K[N])\text{Bmse}(\hat{A}[N-1]).$$

We have now derived the sequential LMMSE estimator equations based on the vector space viewpoint.

In Appendix 12A we generalize the preceding vector space derivation for the sequential vector LMMSE estimator. To do so we must assume the data has the Bayesian linear model form (see Theorem 12.1) and that $\mathbf{C}_w$ is a *diagonal* matrix, so that the $w[n]$'s are *uncorrelated* with variance $E(w^2[n]) = \sigma_n^2$. The latter is a critical assumption and was used in the previous example. Fortunately, the set of equations obtained are valid for the case of nonzero means as well (see Appendix 12A). Hence, the equations to follow are the sequential implementation of the vector LMMSE estimator of (12.26)




Figure 12.6 Sequential linear minimum mean square estimator

or (12.27) when the noise covariance matrix $\mathbf{C}_w$ is diagonal. To define the equations we let $\hat{\boldsymbol{\theta}}[n]$ be the LMMSE estimator based on $\{x[0], x[1], \dots, x[n]\}$, and $\mathbf{M}[n]$ be the corresponding minimum MSE matrix (just the sequential version of (12.28) or (12.29)) or

$$\mathbf{M}[n] = E[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n])(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n])^T].$$

Also, we partition the $(n+1) \times p$ observation matrix as

$$\mathbf{H}[n] = \begin{bmatrix} \mathbf{H}[n-1] \\ \mathbf{h}^T[n] \end{bmatrix} = \begin{bmatrix} n \times p \\ 1 \times p \end{bmatrix}.$$

Then, the sequential LMMSE estimator becomes (see Appendix 12A)

Estimator Update:

$$\hat{\boldsymbol{\theta}}[n] = \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n](x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n-1]) \quad (12.47)$$

where

$$\mathbf{K}[n] = \frac{\mathbf{M}[n-1]\mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n]\mathbf{M}[n-1]\mathbf{h}[n]}. \quad (12.48)$$

Minimum MSE Matrix Update:

$$\mathbf{M}[n] = (\mathbf{I} - \mathbf{K}[n]\mathbf{h}^T[n])\mathbf{M}[n-1]. \quad (12.49)$$

The gain factor $\mathbf{K}[n]$ is a $p \times 1$ vector, and the minimum MSE matrix has dimension $p \times p$. The entire estimator is summarized in Figure 12.6, where the thick arrows indicate vector processing. To start the recursion we need to specify initial values for $\hat{\boldsymbol{\theta}}[n-1]$ and $\mathbf{M}[n-1]$, so that $\mathbf{K}[n]$ can be determined from (12.48) and then $\hat{\boldsymbol{\theta}}[n]$ from (12.47). To do so we specify $\hat{\boldsymbol{\theta}}[-1]$ and $\mathbf{M}[-1]$ so that the recursion can begin at $n = 0$.

Since no data have been observed at $n = -1$, then from (12.19) our estimator becomes the constant $\hat{\theta}_i = a_{i0}$. It is easily shown that the LMMSE estimator is just the mean of θ_i (see also (12.3)) or

$$\hat{\theta}[-1] = E(\theta).$$

As a result, the minimum MSE matrix is

$$\begin{aligned} \mathbf{M}[-1] &= E[(\theta - \hat{\theta}[-1])(\theta - \hat{\theta}[-1])^T] \\ &= \mathbf{C}_{\theta\theta}. \end{aligned}$$

Several observations are in order.

1. For no prior knowledge about θ we can let $\mathbf{C}_{\theta\theta} \rightarrow \infty$. Then, we have the same *form* as the sequential LSE (see Section 8.7), although the approaches are fundamentally different.
2. No matrix inversions are required.
3. The gain factor $\mathbf{K}[n]$ depends on our confidence in the new data sample as measured by σ_n^2 versus that in the previous data as summarized by $\mathbf{M}[n-1]$.

An example follows.

Example 12.3 - Bayesian Fourier Analysis

We now use sequential LS to compute the LMMSE estimator in Example 11.1. The data model is

$$x[n] = a \cos 2\pi f_0 n + b \sin 2\pi f_0 n + w[n] \quad n = 0, 1, \dots, N-1$$

where f_0 is any frequency (in Example 11.1 we assumed f_0 to be a multiple of $1/N$), excepting 0 and $1/2$ for which $\sin 2\pi f_0 n$ is identically zero, and $w[n]$ is white noise with variance σ^2 . It is desired to estimate $\theta = [a \ b]^T$ sequentially. We further assume that the Bayesian linear model form applies and that θ has mean $E(\theta)$ and covariance $\sigma_\theta^2 \mathbf{I}$. The sequential estimator is initialized by

$$\begin{aligned} \hat{\theta}[-1] &= E(\theta) \\ \mathbf{M}[-1] &= \mathbf{C}_{\theta\theta} = \sigma_\theta^2 \mathbf{I}. \end{aligned}$$

Then, from (12.47)

$$\hat{\theta}[0] = \hat{\theta}[-1] + \mathbf{K}[0](x[0] - \mathbf{h}^T[0]\hat{\theta}[-1]).$$

The $\mathbf{h}^T[0]$ vector is the first row of the observation matrix or

$$\mathbf{h}^T[0] = \begin{bmatrix} 1 & 0 \end{bmatrix}.$$

Subsequent rows are

$$\mathbf{h}^T[n] = \begin{bmatrix} \cos 2\pi f_0 n & \sin 2\pi f_0 n \end{bmatrix}.$$

The 2×1 gain vector is, from (12.48),

$$\mathbf{K}[0] = \frac{\mathbf{M}[-1]\mathbf{h}[0]}{\sigma^2 + \mathbf{h}^T[0]\mathbf{M}[-1]\mathbf{h}[0]}$$

where $\mathbf{M}[-1] = \sigma_s^2 \mathbf{I}$. Once the gain vector has been found, $\hat{\theta}[0]$ can be computed. Finally, the MSE matrix update is, from (12.49),

$$\mathbf{M}[0] = (\mathbf{I} - \mathbf{K}[0]\mathbf{h}^T[0])\mathbf{M}[-1]$$

and the procedure continues in like manner for $n \geq 1$. ◇

12.7 Signal Processing Examples - Wiener Filtering

We now examine in detail some of the important applications of the LMMSE estimator. In doing so we will assume that the data $\{x[0], x[1], \dots, x[N-1]\}$ is WSS with zero mean. As such, the $N \times N$ covariance matrix $\mathbf{C}_{xx}$ takes the symmetric Toeplitz form

$$\begin{aligned} \mathbf{C}_{xx} &= \begin{bmatrix} r_{xx}[0] & r_{xx}[1] & \dots & r_{xx}[N-1] \\ r_{xx}[1] & r_{xx}[0] & \dots & r_{xx}[N-2] \\ \vdots & \vdots & \ddots & \vdots \\ r_{xx}[N-1] & r_{xx}[N-2] & \dots & r_{xx}[0] \end{bmatrix} \\ &= \mathbf{R}_{xx} \end{aligned} \quad (12.50)$$

where $r_{xx}[k]$ is the ACF of the $x[n]$ process and $\mathbf{R}_{xx}$ denotes the autocorrelation matrix. Furthermore, the parameter θ to be estimated is also assumed to be zero mean. The group of applications that we will describe are generically called *Wiener filters*.

There are three main problems that we will study. They are (see Figure 12.7)

1. *Filtering*, where $\theta = s[n]$ is to be estimated based on $x[m] = s[m] + w[m]$ for $m = 0, 1, \dots, n$. The sequences $s[n]$ and $w[n]$ represent signal and noise processes. The problem is to filter the signal from the noise. Note that the signal sample is estimated based on the *present and past data only*, so that as n increases, we view the estimation process as the application of a causal filter to the data.
2. *Smoothing*, where $\theta = s[n]$ is to be estimated for $n = 0, 1, \dots, N-1$ based on the data set $\{x[0], x[1], \dots, x[N-1]\}$, where $x[n] = s[n] + w[n]$. In contrast to filtering we are allowed to use future data. For instance, to estimate $s[1]$ we can use the entire data set $\{x[0], x[1], \dots, x[N-1]\}$, while in filtering we are constrained to use only $\{x[0], x[1]\}$. Clearly, in smoothing an estimate cannot be obtained until all the data has been collected.
3. *Prediction*, where $\theta = x[N-1+l]$ for l a positive integer is to be estimated based on $\{x[0], x[1], \dots, x[N-1]\}$. This is referred to as the l -step prediction problem.




Figure 12.7 Wiener filtering problem definitions

A related problem that is explored in Problem 12.14 is to estimate $x[n]$ based on $\{x[0], \dots, x[n-1], x[n+1], \dots, x[N-1]\}$ and is termed the *interpolation* problem. To solve all three problems we use (12.20) with $E(\boldsymbol{\theta}) = E(\mathbf{x}) = \mathbf{0}$ or

$$\hat{\boldsymbol{\theta}} = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x} \quad (12.51)$$

and the minimum Bayesian MSE matrix given by (12.21)

$$\mathbf{M}_{\hat{\theta}} = \mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta}. \quad (12.52)$$

Consider first the smoothing problem. We wish to estimate $\boldsymbol{\theta} = \mathbf{s} = [s[0] s[1] \dots s[N-1]]^T$ based on $\mathbf{x} = [x[0] x[1] \dots x[N-1]]^T$. We will make the reasonable assumption that the signal and noise processes are uncorrelated. Hence,

$$r_{xx}[k] = r_{ss}[k] + r_{ww}[k].$$

Then, we have

$$\mathbf{C}_{xx} = \mathbf{R}_{xx} = \mathbf{R}_{ss} + \mathbf{R}_{ww}.$$

Also,

$$\mathbf{C}_{\theta x} = E(\mathbf{s}\mathbf{x}^T) = E(\mathbf{s}(\mathbf{s} + \mathbf{w})^T) = \mathbf{R}_{ss}.$$

Therefore, the Wiener estimator of the signal is, from (12.51),

$$\hat{\mathbf{s}} = \mathbf{R}_{ss}(\mathbf{R}_{ss} + \mathbf{R}_{ww})^{-1}\mathbf{x}. \quad (12.53)$$

The $N \times N$ matrix

$$\mathbf{W} = \mathbf{R}_{ss}(\mathbf{R}_{ss} + \mathbf{R}_{ww})^{-1} \quad (12.54)$$

is referred to as the *Wiener smoothing matrix*. The corresponding minimum MSE matrix is, from (12.52),

$$\begin{aligned} \mathbf{M}_{\hat{\mathbf{s}}} &= \mathbf{R}_{ss} - \mathbf{R}_{ss}(\mathbf{R}_{ss} + \mathbf{R}_{ww})^{-1}\mathbf{R}_{ss} \\ &= (\mathbf{I} - \mathbf{W})\mathbf{R}_{ss}. \end{aligned} \quad (12.55)$$

As an example, if $N = 1$ we would wish to estimate $s[0]$ based on $x[0] = s[0] + w[0]$. Then, the Wiener smoother $\mathbf{W}$ is just a scalar W given by

$$W = \frac{r_{ss}[0]}{r_{ss}[0] + r_{ww}[0]} = \frac{\eta}{\eta + 1} \quad (12.56)$$

where $\eta = r_{ss}[0]/r_{ww}[0]$ is the SNR. For a high SNR so that $W \rightarrow 1$, we have $\hat{s}[0] \rightarrow x[0]$, while for a low SNR so that $W \rightarrow 0$, we have $\hat{s}[0] \rightarrow 0$. The corresponding minimum MSE is

$$M_{\hat{\mathbf{s}}} = (1 - W)r_{ss}[0] = \left(1 - \frac{\eta}{\eta + 1}\right)r_{ss}[0]$$

which for these two extremes is either 0 for a high SNR or $r_{ss}[0]$ for a low SNR. A numerical example of a Wiener smoother has been given in Section 11.7. See also Problem 12.15.

We next consider the filtering problem in which we wish to estimate $\theta = s[n]$ based on $\mathbf{x} = [x[0] x[1] \dots x[n]]^T$. The problem is repeated for each value of n until the entire signal $s[n]$ for $n = 0, 1, \dots, N-1$ has been estimated. As before, $x[n] = s[n] + w[n]$, where $s[n]$ and $w[n]$ are signal and noise processes that are uncorrelated with each other. Thus,

$$\mathbf{C}_{xx} = \mathbf{R}_{ss} + \mathbf{R}_{ww}$$

where $\mathbf{R}_{ss}$, $\mathbf{R}_{ww}$ are $(n+1) \times (n+1)$ autocorrelation matrices. Also,

$$\begin{aligned} \mathbf{C}_{\theta x} &= E(s[n] [x[0] \ x[1] \ \dots \ x[n]]) \\ &= E(s[n] [s[0] \ s[1] \ \dots \ s[n]]) \\ &= [r_{ss}[n] \ r_{ss}[n-1] \ \dots \ r_{ss}[0]]. \end{aligned}$$

Letting the latter row vector be denoted as $\mathbf{r}_{ss}^T$, we have from (12.51)

$$\hat{s}[n] = \mathbf{r}_{ss}^T (\mathbf{R}_{ss} + \mathbf{R}_{ww})^{-1} \mathbf{x}. \quad (12.57)$$

The $(n+1) \times 1$ vector of weights is seen to be

$$\mathbf{a} = (\mathbf{R}_{ss} + \mathbf{R}_{ww})^{-1} \mathbf{r}'_{ss}$$

recalling that

$$\hat{s}[n] = \mathbf{a}^T \mathbf{x}$$

where $\mathbf{a} = [a_0 \ a_1 \ \dots \ a_n]^T$ as in our original formulation of the scalar LMMSE estimator of (12.1). We may interpret the process of forming the estimator as n increases as a filtering operation if we define a *time varying* impulse response $h^{(n)}[k]$. Specifically, we let $h^{(n)}[k]$ be the response of a filter at time n to an impulse applied k samples before. To make the filtering correspondence we let

$$h^{(n)}[k] = a_{n-k} \quad k = 0, 1, \dots, n.$$

Note for future reference that the vector $\mathbf{h} = [h^{(n)}[0] \ h^{(n)}[1] \ \dots \ h^{(n)}[n]]^T$ is just the vector $\mathbf{a}$ when flipped upside down. Then,

$$\begin{aligned} \hat{s}[n] &= \sum_{k=0}^n a_k x[k] \\ &= \sum_{k=0}^n h^{(n)}[n-k] x[k] \end{aligned}$$

or

$$\hat{s}[n] = \sum_{k=0}^n h^{(n)}[k] x[n-k] \quad (12.58)$$

which is recognized as a time varying FIR filter. To explicitly find the impulse response $\mathbf{h}$ we note that since

$$(\mathbf{R}_{ss} + \mathbf{R}_{ww}) \mathbf{a} = \mathbf{r}'_{ss}$$

it follows that

$$(\mathbf{R}_{ss} + \mathbf{R}_{ww}) \mathbf{h} = \mathbf{r}_{ss}$$

where $\mathbf{r}_{ss} = [r_{ss}[0] \ r_{ss}[1] \ \dots \ r_{ss}[n]]^T$. This result depends on the symmetric Toeplitz nature of $\mathbf{R}_{ss} + \mathbf{R}_{ww}$ and the fact that $\mathbf{h}$ is just $\mathbf{a}$ when flipped upside down. Written out, the set of linear equations becomes

$$\begin{bmatrix} r_{xx}[0] & r_{xx}[1] & \dots & r_{xx}[n] \\ r_{xx}[1] & r_{xx}[0] & \dots & r_{xx}[n-1] \\ \vdots & \vdots & \ddots & \vdots \\ r_{xx}[n] & r_{xx}[n-1] & \dots & r_{xx}[0] \end{bmatrix} \begin{bmatrix} h^{(n)}[0] \\ h^{(n)}[1] \\ \vdots \\ h^{(n)}[n] \end{bmatrix} = \begin{bmatrix} r_{ss}[0] \\ r_{ss}[1] \\ \vdots \\ r_{ss}[n] \end{bmatrix} \quad (12.59)$$

where $r_{xx}[k] = r_{ss}[k] + r_{ww}[k]$. These are the *Wiener-Hopf filtering equations*. It appears that they must be solved for each value of n . However, a computationally efficient algorithm for doing so is the Levinson recursion, which solves the equations recursively to avoid resolving them for each value of n [Marple 1987]. For large enough

n it can be shown that the filter becomes time invariant so that only a single solution is necessary. In this case an analytical solution may be possible. The Wiener-Hopf filtering equations can be written as

$$\sum_{k=0}^n h^{(n)}[k] r_{xx}[l-k] = r_{ss}[l] \quad l = 0, 1, \dots, n$$

where we have used the property $r_{xx}[-k] = r_{xx}[k]$. As $n \rightarrow \infty$, we have upon replacing the time varying impulse response $h^{(n)}[k]$ by its time invariant version $h[k]$

$$\sum_{k=0}^{\infty} h[k] r_{xx}[l-k] = r_{ss}[l] \quad l = 0, 1, \dots \quad (12.60)$$

The same set of equations result if we attempt to estimate $s[n]$ based on the present and infinite past or based on $x[m]$ for $m \leq n$. This is termed the *infinite Wiener filter*. To see why this is so we let

$$\hat{s}[n] = \sum_{k=0}^{\infty} h[k] x[n-k]$$

and use the orthogonality principle (12.13). Then, we have

$$s[n] - \hat{s}[n] \perp \dots, x[n-1], x[n]$$

or by using the definition of orthogonality

$$E[(s[n] - \hat{s}[n])x[n-l]] = 0 \quad l = 0, 1, \dots$$

Hence,

$$E\left(\sum_{k=0}^{\infty} h[k] x[n-k] x[n-l]\right) = E(s[n] x[n-l])$$

and therefore, the equations to be solved for the infinite Wiener filter impulse response are

$$\sum_{k=0}^{\infty} h[k] r_{xx}[l-k] = r_{ss}[l] \quad l = 0, 1, \dots$$

A little thought will convince the reader that the solutions must be identical since the problem of estimating $s[n]$ based on $x[m]$ for $0 \leq m \leq n$ as $n \rightarrow \infty$ is really just that of using the present and infinite past to estimate the current sample. The time invariance of the filter makes the solution independent of which sample is to be estimated or independent of n . The solution of (12.60) utilizes spectral factorization and is explored in Problem 12.16. At first glance it might appear that (12.60) could be solved by using Fourier transform techniques since the left-hand side of the equation is a convolution of two sequences. This is not the case, however, since the equations hold only for $l \geq 0$. If, indeed, they were valid for $l < 0$ as well, then the Fourier transform approach would be viable. Such a set of equations arises in the smoothing problem in which $s[n]$ is

to be estimated based on $\{\dots, x[-1], x[0], x[1], \dots\}$ or $x[k]$ for *all* k . In this case the smoothing estimator takes the form

$$\hat{s}[n] = \sum_{k=-\infty}^{\infty} a_k x[k]$$

and by letting $h[k] = a_{n-k}$ we have the convolution sum

$$\hat{s}[n] = \sum_{k=-\infty}^{\infty} h[k] x[n-k]$$

where $h[k]$ is the impulse response of an infinite two-sided time invariant filter. The Wiener-Hopf equations become

$$\sum_{k=-\infty}^{\infty} h[k] r_{xx}[l-k] = r_{ss}[k] \quad -\infty < l < \infty \quad (12.61)$$

(see Problem 12.17). The difference from the filtering case is that now the equations must be satisfied for all l , and there is no constraint that $h[k]$ must be causal. Hence, we can use Fourier transform techniques to solve for the impulse response since (12.61) becomes

$$h[n] \star r_{xx}[n] = r_{ss}[n] \quad (12.62)$$

where $\star$ denotes convolution. Letting $H(f)$ be the frequency response of the infinite Wiener smoother, we have upon taking the Fourier transform of (12.62):

$$\begin{aligned} H(f) &= \frac{P_{ss}(f)}{P_{xx}(f)} \\ &= \frac{P_{ss}(f)}{P_{ss}(f) + P_{ww}(f)}. \end{aligned}$$

The frequency response is a real even function of frequency, so that the impulse response must also be real and even. This means that the filter is not causal. Of course, this is to be expected since we sought to estimate $s[n]$ using *future* as well as present and past data. Before leaving this topic it is worthwhile to point out that the Wiener smoother emphasizes portions of the frequency spectrum of the data where the SNR is high and attenuates those where it is low. This is evident if we define the “local SNR” as the SNR in a narrow band of frequencies centered about f as

$$\eta(f) = \frac{P_{ss}(f)}{P_{ww}(f)}.$$

Then, the optimal filter frequency response becomes

$$H(f) = \frac{\eta(f)}{\eta(f) + 1}.$$

Clearly, the filter response satisfies $0 < H(f) < 1$, and thus the Wiener smoother response is $H(f) \approx 0$ when $\eta(f) \approx 0$ (low local SNR) and $H(f) \approx 1$ when $\eta(f) \rightarrow \infty$ (high local SNR). The reader may wish to compare these results to those of the Wiener filter given by (12.56). See also Problem 12.18 for further results on the infinite Wiener smoother.

Finally, we examine the prediction problem in which we wish to estimate $\theta = x[N-1+l]$ for $l \geq 1$ based on $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[N-1]]^T$. The resulting estimator is termed the l -step linear predictor. We use (12.51) for which

$$\mathbf{C}_{xx} = \mathbf{R}_{xx}$$

where $\mathbf{R}_{xx}$ is of dimension $N \times N$ and

$$\begin{aligned} \mathbf{C}_{\theta x} &= E [x[N-1+l] [x[0] \ x[1] \ \dots \ x[N-1]]^T] \\ &= [r_{xx}[N-1+l] \ r_{xx}[N-2+l] \ \dots \ r_{xx}[l]]^T. \end{aligned}$$

Let the latter vector be denoted by $\mathbf{r}_{xx}'^T$. Then,

$$\hat{x}[N-1+l] = \mathbf{r}_{xx}'^T \mathbf{R}_{xx}^{-1} \mathbf{x}.$$

Recalling that

$$\mathbf{a} = \mathbf{R}_{xx}^{-1} \mathbf{r}_{xx}'^T \quad (12.63)$$

we have

$$\hat{x}[N-1+l] = \sum_{k=0}^{N-1} a_k x[k].$$

If we let $h[N-k] = a_k$ to allow a “filtering” interpretation, then

$$\begin{aligned} \hat{x}[N-1+l] &= \sum_{k=0}^{N-1} h[N-k] x[k] \\ &= \sum_{k=1}^N h[k] x[N-k] \end{aligned} \quad (12.64)$$

and it is observed that the predicted sample is the output of a filter with impulse response $h[n]$. The equations to be solved are from (12.63) (noting once again that $\mathbf{h}$ is just $\mathbf{a}$ when flipped upside down)

$$\mathbf{R}_{xx} \mathbf{h} = \mathbf{r}_{xx}$$

where $\mathbf{r}_{xx} = [r_{xx}[l] \ r_{xx}[l+1] \ \dots \ r_{xx}[N-1+l]]^T$. In explicit form they become

$$\begin{bmatrix} r_{xx}[0] & r_{xx}[1] & \dots & r_{xx}[N-1] \\ r_{xx}[1] & r_{xx}[0] & \dots & r_{xx}[N-2] \\ \vdots & \vdots & \ddots & \vdots \\ r_{xx}[N-1] & r_{xx}[N-2] & \dots & r_{xx}[0] \end{bmatrix} \begin{bmatrix} h[1] \\ h[2] \\ \vdots \\ h[N] \end{bmatrix} = \begin{bmatrix} r_{xx}[l] \\ r_{xx}[l+1] \\ \vdots \\ r_{xx}[N-1+l] \end{bmatrix}$$

$$= \begin{bmatrix} r_{xx}[l] \\ r_{xx}[l+1] \\ \vdots \\ r_{xx}[N-1+l] \end{bmatrix}. \quad (12.65)$$

These are the Wiener-Hopf prediction equations for the l -step linear predictor based on N past samples. A computationally efficient method for solving these equations is the Levinson recursion [Marple 1987]. For the specific case where $l = 1$, the one-step linear predictor, the values of $-h[n]$ are termed the *linear prediction coefficients* which are used extensively in speech modeling [Makhoul 1975]. Also, for $l = 1$ the resulting equations are identical to the Yule-Walker equations used to solve for the AR filter parameters of an AR(N) process (see Example 7.18 and Appendix 1).

The minimum MSE for the l -step linear predictor is, from (12.52),

$$M_{\hat{x}} = r_{xx}[0] - \mathbf{r}_{xx}^T \mathbf{R}_{xx}^{-1} \mathbf{r}_{xx}'$$

or, equivalently,

$$\begin{aligned} M_{\hat{x}} &= r_{xx}[0] - \mathbf{r}_{xx}^T \mathbf{a} \\ &= r_{xx}[0] - \sum_{k=0}^{N-1} a_k r_{xx}[N-1+l-k] \\ &= r_{xx}[0] - \sum_{k=0}^{N-1} h[N-k] r_{xx}[N-1+l-k] \\ &= r_{xx}[0] - \sum_{k=1}^N h[k] r_{xx}[k+(l-1)]. \end{aligned} \quad (12.66)$$

As an example, assume that $x[n]$ is an AR(1) process with ACF (see Appendix 1)

$$r_{xx}[k] = \frac{\sigma_u^2}{1-a^2[1]} (-a[1])^{|k|}$$

and we wish to find the one-step predictor $\hat{x}[N]$ as (see (12.64))

$$\hat{x}[N] = \sum_{k=1}^N h[k] x[N-k].$$

To solve for the $h[k]$'s we use (12.65) for $l = 1$, which is

$$\sum_{k=1}^N h[k] r_{xx}[m-k] = r_{xx}[m] \quad m = 1, 2, \dots, N.$$

Substituting the ACF, we have

$$\sum_{k=1}^N h[k] (-a[1])^{|m-k|} = (-a[1])^{|m|}.$$

It is easily verified that these equations are solved for

$$h[k] = \begin{cases} -a[1] & k = 1 \\ 0 & k = 2, 3, \dots, N. \end{cases}$$

Hence, the one-step linear predictor is

$$\hat{x}[N] = -a[1]x[N-1]$$

and depends only on the previous sample. This is not surprising in light of the fact that the AR(1) process satisfies

$$x[n] = -a[1]x[n-1] + u[n]$$

where $u[n]$ is white noise, so that the sample to be predicted satisfies

$$x[N] = -a[1]x[N-1] + u[N].$$

The predictor cannot predict $u[N]$ since it is uncorrelated with the past data samples (recall that since the AR(1) filter is causal, $x[n]$ is a linear combination of $\{u[n], u[n-1], \dots\}$ and thus $x[n]$ is uncorrelated with all future samples of $u[n]$). The prediction error is $x[N] - \hat{x}[N] = u[N]$, and the minimum MSE is just the driving noise power σ_u^2 . To verify this we have from (12.66)

$$\begin{aligned} M_{\hat{x}} &= r_{xx}[0] - \sum_{k=1}^N h[k]r_{xx}[k] \\ &= r_{xx}[0] + a[1]r_{xx}[1] \\ &= \frac{\sigma_u^2}{1 - a^2[1]} + a[1] \frac{\sigma_u^2}{1 - a^2[1]} (-a[1]) \\ &= \sigma_u^2. \end{aligned}$$

We can extend these results to the l -step predictor by solving

$$\sum_{k=1}^N h[k]r_{xx}[m-k] = r_{xx}[m+l-1] \quad m = 1, 2, \dots, N.$$

Substituting the ACF for an AR(1) process this becomes

$$\sum_{k=1}^N h[k] (-a[1])^{|m-k|} = (-a[1])^{|m+l-1|}.$$

The solution, which can be easily verified, is

$$h[k] = \begin{cases} (-a[1])^l & k = 1 \\ 0 & k = 2, 3, \dots, N \end{cases}$$




Figure 12.8 Linear prediction for realization of AR(1) process

and therefore the l -step predictor is

$$\hat{x}[(N-1)+l] = (-a[1])^l x[N-1]. \quad (12.67)$$

The minimum MSE is, from (12.66),

$$\begin{aligned} M_{\hat{x}} &= r_{xx}[0] - h[1]r_{xx}[l] \\ &= \frac{\sigma_u^2}{1-a^2[1]} - (-a[1])^l \frac{\sigma_u^2}{1-a^2[1]} (-a[1])^l \\ &= \frac{\sigma_u^2}{1-a^2[1]} (1-a^{2l}[1]). \end{aligned}$$

It is interesting to note that the predictor decays to zero with l (since $|a[1]| < 1$). This is reasonable since the correlation between $x[(N-1)+l]$, the sample to be predicted, and $x[N-1]$, the data sample on which the prediction is based, is $r_{xx}[l]$. As l increases, $r_{xx}[l]$ decays to zero and thus so does $\hat{x}[(N-1)+l]$. This is also reflected in the minimum MSE, which is smallest for $l=1$ and increases for larger l .

A numerical example illustrates the behavior of the l -step predictor. If $a[1] = -0.95$ and $\sigma_u^2 = 0.1$, so that the process has a low-pass PSD, then for a given realization of $x[n]$ we obtain the results shown in Figure 12.8. The true data are displayed as a solid line, while the predicted data for $n \geq 11$ is shown as a dashed line. The predictions are given by (12.67), where $N = 11$ and $l = 1, 2, \dots, 40$, and thus decay to zero with increasing l . As can be observed, the predictions are generally poor except for small l . See also Problems 12.19 and 12.20.

References

- Kay, S., "Some Results in Linear Interpolation Theory," *IEEE Trans. Acoust., Speech, Signal Process.*, Vol. 31, pp. 746-749, June 1983.
- Luenberger, D.G., *Optimization by Vector Space Methods*, J. Wiley, New York, 1969.
- Makhoul, J., "Linear Prediction: A Tutorial Review," *Proc. IEEE*, Vol. 63, pp. 561-580, April 1975.
- Marple, S.L., Jr., *Digital Spectral Analysis*, Prentice-Hall, Englewood Cliffs, N.J., 1987.
- Orfanidis, S.J., *Optimum Signal Processing*, Macmillan, New York, 1985.

Problems

12.1 Consider the quadratic estimator

$$\hat{\theta} = ax^2[0] + bx[0] + c$$

of a scalar parameter θ based on the single data sample $x[0]$. Find the coefficients a, b, c that minimize the Bayesian MSE. If $x[0] \sim \mathcal{U}[-\frac{1}{2}, \frac{1}{2}]$, find the LMMSE estimator and the quadratic MMSE estimator if $\theta = \cos 2\pi x[0]$. Also, compare the minimum MSEs.

12.2 Consider the data

$$x[n] = Ar^n + w[n] \quad n = 0, 1, \dots, N-1$$

where A is a parameter to be estimated, r is a known constant, and $w[n]$ is zero mean white noise with variance σ^2 . The parameter A is modeled as a random variable with mean μ_A and variance σ_A^2 and is independent of $w[n]$. Find the LMMSE estimator of A and the minimum Bayesian MSE.

- 12.3** A Gaussian random vector $\mathbf{x} = [x_1 \ x_2]^T$ has zero mean and covariance matrix $\mathbf{C}_{xx}$. If x_2 is to be linearly estimated based on x_1 , find the estimator that minimizes the Bayesian MSE. Also, find the minimum MSE and prove that it is zero if and only if $\mathbf{C}_{xx}$ is singular. Extend your results to show that if the covariance matrix of an $N \times 1$ zero mean Gaussian random vector is not positive definite, then any random variable may be perfectly estimated by a linear combination of the others. Hint: Note that

$$E \left[\left(\sum_{i=1}^N a_i x_i \right)^2 \right] = \mathbf{a}^T \mathbf{C}_{xx} \mathbf{a}.$$

- 12.4** An inner product (x, y) between two vectors x and y of a vector space must satisfy the following properties

- $(x, x) \geq 0$ and $(x, x) = 0$ if and only if $x = 0$.
- $(x, y) = (y, x)$.
- $(c_1 x_1 + c_2 x_2, y) = c_1 (x_1, y) + c_2 (x_2, y)$.

Prove that the definition $(x, y) = E(xy)$ for x and y zero mean random variables satisfies these properties.

- 12.5** If we assume nonzero mean random variables, then a reasonable approach is to define the inner product between x and y as $(x, y) = \text{cov}(x, y)$. With this definition x and y are orthogonal if and only if they are uncorrelated. For this “inner product” which of the properties given in Problem 12.4 is violated?
- 12.6** We observe the data $x[n] = s[n] + w[n]$ for $n = 0, 1, \dots, N-1$, where $s[n]$ and $w[n]$ are zero mean WSS random processes which are uncorrelated with each other. The ACFs are

$$\begin{aligned} r_{ss}[k] &= \sigma_s^2 \delta[k] \\ r_{ww}[k] &= \sigma_w^2 \delta[k]. \end{aligned}$$

Determine the LMMSE estimator of $s = [s[0] s[1] \dots s[N-1]]^T$ based on $\mathbf{x} = [x[0] x[1] \dots x[N-1]]^T$ and the corresponding minimum MSE matrix.

- 12.7** Derive the Bayesian MSE matrix for the vector LMMSE estimator as given by (12.21), as well as the minimum Bayesian MSE of (12.22). Keep in mind the averaging PDFs implied by the expectation operator.
- 12.8** Prove the commutative and additive properties of the LMMSE estimator as given by (12.23) and (12.24).
- 12.9** Derive a sequential LMMSE estimator analogous to (12.34)–(12.36) but for the case where $\mu_A \neq 0$. Hint: Use an algebraic approach based on the results in Example 10.1.
- 12.10** Given a set of vectors $\{x_1, x_2, \dots, x_n\}$, the Gram-Schmidt orthogonalization procedure finds a new set of vectors $\{e_1, e_2, \dots, e_n\}$ which are orthonormal (orthogonal and having unit length) or $(e_i, e_j) = \delta_{ij}$. The procedure is

$$\begin{aligned} \text{a.} \quad e_1 &= \frac{x_1}{\|x_1\|} \\ \text{b.} \quad z_2 &= x_2 - (x_2, e_1)e_1 \\ e_2 &= \frac{z_2}{\|z_2\|} \end{aligned}$$

c. and so on,

or in general for $n \geq 2$

$$\begin{aligned} z_n &= x_n - \sum_{i=1}^{n-1} (x_n, e_i)e_i \\ e_n &= \frac{z_n}{\|z_n\|}. \end{aligned}$$

Give a geometrical interpretation of the procedure. For the Euclidean vectors

$$\begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}, \quad \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \quad \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

find three orthonormal vectors using the Gram-Schmidt procedure. The inner product is defined as $(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y}$ for $\mathbf{x}$ and $\mathbf{y}$ Euclidean vectors.

- 12.11** Let $\mathbf{x}$ denote the vector composed of three zero mean random variables with a covariance matrix

$$\mathbf{C}_{xx} = \begin{bmatrix} 1 & \rho & \rho^2 \\ \rho & 1 & \rho \\ \rho^2 & \rho & 1 \end{bmatrix}.$$

If $\mathbf{y} = \mathbf{A}\mathbf{x}$, determine the 3×3 matrix $\mathbf{A}$ so that the covariance matrix of $\mathbf{y}$ is $\mathbf{I}$ or, equivalently, so that the random variables $\{y_1, y_2, y_3\}$ are uncorrelated and have unit variance. Use the Gram-Schmidt orthogonalization procedure in Problem 12.10 to do so. What is interesting about $\mathbf{A}$? Finally, relate $\mathbf{C}_{xx}^{-1}$ to $\mathbf{A}$. Note that $\mathbf{A}$ may be viewed as a *whitening* transformation (see Section 4.5).

- 12.12** For the sequential LMMSE estimator of a scalar parameter explain what would happen if $\sigma_n^2 \rightarrow 0$ for some n . Do you obtain the same results for the vector case?
- 12.13** Find the sequential LMMSE estimator for the problem described in Example 12.1. Verify that the sequential estimator is identical to the batch estimator (12.9). Hint: Solve for the $1/K[n]$ sequence.
- 12.14** In this problem we examine the interpolation of a data sample. We assume that the data set $\{x[n-M], \dots, x[n-1], x[n+1], \dots, x[n+M]\}$ is available and that we wish to estimate or interpolate $x[n]$. The data and $x[n]$ are assumed to be a realization of a zero mean WSS random process. Let the LMMSE estimator of $x[n]$ be

$$\hat{x}[n] = \sum_{\substack{k=-M \\ k \neq 0}}^M a_k x[n-k].$$

Find the set of linear equations to be solved for the weighting coefficients by using the orthogonality principle. Next, prove that $a_{-k} = a_k$ and explain why this must be true. See also [Kay 1983] for a further discussion of interpolation.

- 12.15** Consider the Wiener smoother for a single data sample as given by (12.56). Rewrite W and M_s as a function of the correlation coefficient

$$\rho = \frac{\text{cov}(s[0], x[0])}{\sqrt{\text{var}(s[0])\text{var}(x[0])}}$$

between $s[0]$ and $x[0]$. Explain your results.

12.16 In this problem we explore the solution of the Wiener-Hopf filter equations based on the present and infinite past or the solution of

$$\sum_{k=0}^{\infty} h[k] r_{xx}[l-k] = r_{ss}[l] \quad l \geq 0$$

by the use of z transforms. First we define the one-sided z transform as

$$\begin{aligned} [\mathcal{X}(z)]_+ &= \left[\sum_{n=-\infty}^{\infty} x[n] z^{-n} \right]_+ \\ &= \sum_{n=0}^{\infty} x[n] z^{-n}. \end{aligned}$$

This is seen to be the usual z transform but with the positive powers of z omitted. Next, we write the Wiener-Hopf equation as

$$h[n] \star r_{xx}[n] - r_{ss}[n] = 0 \quad n \geq 0 \quad (12.68)$$

where $h[n]$ is constrained to be causal. The z transform (two-sided) of the left-hand side is

$$\mathcal{H}(z) \mathcal{P}_{xx}(z) - \mathcal{P}_{ss}(z)$$

so that to satisfy (12.68) we must have

$$[\mathcal{H}(z) \mathcal{P}_{xx}(z) - \mathcal{P}_{ss}(z)]_+ = 0.$$

By the spectral factorization theorem [Orfanidis 1985], if $\mathcal{P}_{xx}(z)$ has no zeros on the unit circle, then it may be factored as

$$\mathcal{P}_{xx}(z) = \mathcal{B}(z) \mathcal{B}(z^{-1})$$

where $\mathcal{B}(z)$ is the z transform of a causal sequence and $\mathcal{B}(z^{-1})$ is the z transform of an anticausal sequence. Hence,

$$[\mathcal{H}(z) \mathcal{B}(z) \mathcal{B}(z^{-1}) - \mathcal{P}_{ss}(z)]_+ = 0.$$

Let $\mathcal{G}(z) = \mathcal{H}(z) \mathcal{B}(z)$, so that

$$\left[\mathcal{B}(z^{-1}) \left(\mathcal{G}(z) - \frac{\mathcal{P}_{ss}(z)}{\mathcal{B}(z^{-1})} \right) \right]_+ = 0.$$

Noting that $\mathcal{G}(z)$ is the z transform of a causal sequence, show that

$$\mathcal{H}(z) = \frac{1}{\mathcal{B}(z)} \left[\frac{\mathcal{P}_{ss}(z)}{\mathcal{B}(z^{-1})} \right]_+.$$

For an example of the computations involved in determining the Wiener filter, see [Orfanidis 1985].

12.17 Rederive the infinite Wiener smoother (12.61) by first assuming that

$$\hat{s}[n] = \sum_{k=-\infty}^{\infty} h[k]x[n-k]$$

and then using the orthogonality principle.

12.18 For the infinite Wiener smoother show that the minimum Bayesian MSE is

$$M_s = r_{ss}[0] - \sum_{k=-\infty}^{\infty} h[k]r_{ss}[k].$$

Then, using Fourier transform techniques, show that this can be rewritten as

$$M_s = \int_{-\frac{1}{2}}^{\frac{1}{2}} P_{ss}(f)(1 - H(f)) df$$

where

$$H(f) = \frac{P_{ss}(f)}{P_{ss}(f) + P_{ww}(f)}.$$

Evaluate the Wiener smoother and minimum MSE if

$$\begin{aligned} P_{ss}(f) &= \begin{cases} P_0 & |f| \leq \frac{1}{4} \\ 0 & \frac{1}{4} < |f| \leq \frac{1}{2} \end{cases} \\ P_{ww}(f) &= \sigma^2 \end{aligned}$$

and explain your results.

12.19 Assume that $x[n]$ is a zero mean WSS random process. We wish to predict $x[n]$ based on $\{x[n-1], x[n-2], \dots, x[n-N]\}$. Use the orthogonality principle to derive the LMMSE estimator or predictor. Explain why the equations to be solved are the same as (12.65) for $l = 1$, i.e., they are independent of n . Also, rederive the minimum MSE (12.66) by again invoking the orthogonality principle.

12.20 Consider an AR(N) process

$$x[n] = - \sum_{k=1}^N a[k]x[n-k] + u[n]$$

where $u[n]$ is white noise with variance σ_u^2 . Prove that the optimal one-step linear predictor of $x[n]$ is

$$\hat{x}[n] = - \sum_{k=1}^N a[k]x[n-k].$$

Also, find the minimum MSE. Hint: Compare the equations to be solved to the Yule-Walker equations (see Appendix 1).

Appendix 12A

Derivation of Sequential LMMSE Estimator for the Bayesian Linear Model

We will use a vector space approach to update the i th component of $\hat{\boldsymbol{\theta}}[n-1]$ as follows.

$$\hat{\theta}_i[n] = \hat{\theta}_i[n-1] + K_i[n] (x[n] - \hat{x}[n|n-1]) \quad (12A.1)$$

where $\hat{x}[n|n-1]$ is the LMMSE estimator of $x[n]$ based on $\{x[0], x[1], \dots, x[n-1]\}$. We now shorten the notation from that used in Section 12.6 where we denoted the same LMMSE estimator as $\hat{x}[n|0, 1, \dots, n-1]$. The motivation for (12A.1) follows by the same rationale as depicted in Figure 12.5 (see also (12.41)). Also, we let $\tilde{x}[n] = x[n] - \hat{x}[n|n-1]$, which is the innovation sequence. We can then form the update for the entire parameter as

$$\hat{\boldsymbol{\theta}}[n] = \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n](x[n] - \hat{x}[n|n-1]).$$

Before beginning the derivation we state some properties which are needed. The reader should recall that all random variables are assumed to be zero mean.

1. The LMMSE estimator of $\mathbf{A}\boldsymbol{\theta}$ is $\mathbf{A}\hat{\boldsymbol{\theta}}$, the commutative property of (12.23).
2. The LMMSE estimator of $\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2$ is $\hat{\boldsymbol{\theta}}_1 + \hat{\boldsymbol{\theta}}_2$, the additive property of (12.24).
3. Since $\hat{\theta}_i[n-1]$ is a linear combination of the samples $\{x[0], x[1], \dots, x[n-1]\}$, and the innovation $x[n] - \hat{x}[n|n-1]$ is uncorrelated with the past data samples, it follows that

$$E \left[\hat{\theta}_i[n-1] (x[n] - \hat{x}[n|n-1]) \right] = 0.$$

4. Since $\boldsymbol{\theta}$ and $w[n]$ are uncorrelated as assumed in the Bayesian linear model and $\hat{\boldsymbol{\theta}}[n-1]$ and $w[n]$ are also uncorrelated, it follows that

$$E \left[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1]) w[n] \right] = 0.$$

To see why $\hat{\theta}[n-1]$ and $w[n]$ are uncorrelated note first that $\hat{\theta}[n-1]$ depends linearly on the past data samples or on θ and $\{w[0], w[1], \dots, w[n-1]\}$. But $w[n]$ is uncorrelated with θ and by the assumption that $\mathbf{C}_w$ is *diagonal* or that $w[n]$ is a sequence of uncorrelated random variables, $w[n]$ is uncorrelated with the past noise samples.

With these results the derivation simplifies. To begin we note that

$$x[n] = \mathbf{h}^T[n]\theta + w[n].$$

Using properties 1 and 2 we have that

$$\hat{x}[n|n-1] = \mathbf{h}^T[n]\hat{\theta}[n-1] + \hat{w}[n|n-1].$$

But $\hat{w}[n|n-1] = 0$ since $w[n]$ is uncorrelated with the past data samples, as explained in property 4. Thus,

$$\hat{x}[n|n-1] = \mathbf{h}^T[n]\hat{\theta}[n-1].$$

We next find the gain factor $K_i[n]$. Using (12.37), we have

$$K_i[n] = \frac{E[\theta_i(x[n] - \hat{x}[n|n-1])]}{E[(x[n] - \hat{x}[n|n-1])^2]}. \quad (12A.2)$$

Evaluating the denominator

$$\begin{aligned} E[(x[n] - \hat{x}[n|n-1])^2] &= E\left[\left(\mathbf{h}^T[n]\theta + w[n] - \mathbf{h}^T[n]\hat{\theta}[n-1]\right)^2\right] \\ &= E\left[\left(\mathbf{h}^T[n](\theta - \hat{\theta}[n-1]) + w[n]\right)^2\right] \\ &= \mathbf{h}^T[n]E[(\theta - \hat{\theta}[n-1])(\theta - \hat{\theta}[n-1])^T]\mathbf{h}[n] \\ &\quad + E(w^2[n]) \\ &= \mathbf{h}^T[n]\mathbf{M}[n-1]\mathbf{h}[n] + \sigma_n^2 \end{aligned}$$

where we have used property 4. Now, using properties 3 and 4 to evaluate the numerator

$$\begin{aligned} E[\theta_i(x[n] - \hat{x}[n|n-1])] &= E[(\theta_i - \hat{\theta}_i[n-1])(x[n] - \hat{x}[n|n-1])] \\ &= E[(\theta_i - \hat{\theta}_i[n-1])(\mathbf{h}^T[n]\theta + w[n] - \mathbf{h}^T[n]\hat{\theta}[n-1])] \\ &= E[(\theta_i - \hat{\theta}_i[n-1])(\mathbf{h}^T[n](\theta - \hat{\theta}[n-1]))] \\ &= E[(\theta_i - \hat{\theta}_i[n-1])(\theta - \hat{\theta}[n-1])^T]\mathbf{h}[n] \end{aligned} \quad (12A.3)$$

so that

$$K_i[n] = \frac{E[(\theta_i - \hat{\theta}_i[n-1])(\theta - \hat{\theta}[n-1])^T]\mathbf{h}[n]}{\mathbf{h}^T[n]\mathbf{M}[n-1]\mathbf{h}[n] + \sigma_n^2}$$

and therefore

$$\mathbf{K}[n] = \frac{\mathbf{M}[n-1]\mathbf{h}[n]}{\mathbf{h}^T[n]\mathbf{M}[n-1]\mathbf{h}[n] + \sigma_n^2}. \quad (12A.4)$$

To determine the update for the MSE matrix

$$\begin{aligned} \mathbf{M}[n] &= E[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n])(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n])^T] \\ &= E[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1] - \mathbf{K}[n](x[n] - \hat{x}[n|n-1])) \\ &\quad \cdot (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1] - \mathbf{K}[n](x[n] - \hat{x}[n|n-1]))^T] \\ &= \mathbf{M}[n-1] - E[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1])(x[n] - \hat{x}[n|n-1])\mathbf{K}^T[n]] \\ &\quad - \mathbf{K}[n]E[(x[n] - \hat{x}[n|n-1])(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1])^T] \\ &\quad + \mathbf{K}[n]E[(x[n] - \hat{x}[n|n-1])^2] \mathbf{K}^T[n]. \end{aligned}$$

But from (12A.2) and (12A.4)

$$\mathbf{K}[n]E[(x[n] - \hat{x}[n|n-1])^2] = \mathbf{M}[n-1]\mathbf{h}[n]$$

and from property 3 and (12A.3)

$$\begin{aligned} E[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}[n-1])(x[n] - \hat{x}[n|n-1])] &= E[\boldsymbol{\theta}(x[n] - \hat{x}[n|n-1])] \\ &= \mathbf{M}[n-1]\mathbf{h}[n] \end{aligned}$$

so that

$$\begin{aligned} \mathbf{M}[n] &= \mathbf{M}[n-1] - \mathbf{M}[n-1]\mathbf{h}[n]\mathbf{K}^T[n] - \mathbf{K}[n]\mathbf{h}^T[n]\mathbf{M}[n-1] \\ &\quad + \mathbf{M}[n-1]\mathbf{h}[n]\mathbf{K}^T[n] \\ &= (\mathbf{I} - \mathbf{K}[n]\mathbf{h}^T[n])\mathbf{M}[n-1]. \end{aligned}$$

We next show that the same equations result if $\boldsymbol{\theta}$ and $x[n]$ are not zero mean. Since the sequential implementation of the minimum MSE matrix (12.49) must be identical to (12.28), which does not depend on the means, $\mathbf{M}[n]$ likewise must be independent of the means. Also, since $\mathbf{M}[n]$ depends on $\mathbf{K}[n]$, the gain vector must also be independent of the means. Finally, the estimator update equation (12.47) is valid for zero means, as we have already shown in this appendix. If the means are not zero, we may still apply (12.47) to the data set $x[n] - E(x[n])$ and view the estimate as that of $\boldsymbol{\theta} - E(\boldsymbol{\theta})$. By the commutative property, however, the LMMSE estimator of $\boldsymbol{\theta} + \mathbf{b}$ for $\mathbf{b}$ a constant vector is $\hat{\boldsymbol{\theta}} + \mathbf{b}$ (see (12.23)). Thus, (12.47) becomes

$$\hat{\boldsymbol{\theta}}[n] - E(\boldsymbol{\theta}) = \hat{\boldsymbol{\theta}}[n-1] - E(\boldsymbol{\theta}) + \mathbf{K}[n][x[n] - E(x[n]) - \mathbf{h}^T[n](\hat{\boldsymbol{\theta}}[n-1] - E(\boldsymbol{\theta}))]$$

where $\hat{\boldsymbol{\theta}}[n]$ is the LMMSE estimator for *nonzero means*. Rearranging and canceling terms we have

$$\hat{\boldsymbol{\theta}}[n] = \hat{\boldsymbol{\theta}}[n-1] + \mathbf{K}[n][x[n] - \mathbf{h}^T[n]\hat{\boldsymbol{\theta}}[n-1] - (E(x[n]) - \mathbf{h}^T[n]E(\boldsymbol{\theta}))]$$

and since

$$\begin{aligned}E(x[n]) &= E(\mathbf{h}^T[n]\boldsymbol{\theta} + w[n]) \\ &= \mathbf{h}^T[n]E(\boldsymbol{\theta})\end{aligned}$$

we arrive at the identical equation for the estimator update.

Chapter 13

Kalman Filters

13.1 Introduction

We now discuss an important generalization of the Wiener filter. The significance of the extension is in its ability to accommodate vector signals and noises which additionally may be nonstationary. This is in contrast to the Wiener filter, which is restricted to stationary scalar signals and noises. This generalization is termed the *Kalman filter*. It may be thought of as a sequential MMSE estimator of a signal embedded in noise, where the signal is characterized by a dynamical or state model. It generalizes the sequential MMSE estimator in Section 12.6, to allow the unknown parameters to evolve in time according to a dynamical model. If the signal and noise are jointly Gaussian, then the Kalman filter is an optimal MMSE estimator, and if not, it is the optimal LMMSE estimator.

13.2 Summary

The scalar Gauss-Markov signal model is given in recursive form by (13.1) and explicitly by (13.2). Its mean, covariance, and variance are given by (13.4), (13.5), and (13.6), respectively. Generalizing the model to a vector signal results in (13.12) with the statistical assumptions summarized just below. Also, the explicit representation is given in (13.13). The corresponding mean and covariances are (13.14), (13.15), and (13.16) or in recursive form by (13.17) and (13.18). A summary of the vector Gauss-Markov signal model is given in Theorem 13.1. The sequential MMSE estimator or Kalman filter for a scalar signal and scalar observations is given by (13.38)–(13.42). If the filter attains steady-state, then it becomes the infinite length Wiener filter as described in Section 13.5. Two generalizations of the Kalman filter are to a vector signal, whose equations are (13.50)–(13.54), and to a vector signal and vector observations, whose assumptions and implementation are described in Theorem 13.2. When the signal model and/or observation model is nonlinear, the preceding Kalman filter cannot be applied directly. However, using a linearization approach, one obtains the suboptimal extended Kalman filter, whose equations are (13.67)–(13.71).

13.3 Dynamical Signal Models

We begin our discussion of signal modeling by recalling our usual example of a DC level in WGN or

$$x[n] = A + w[n]$$

where A is the parameter to be estimated and $w[n]$ is WGN with variance σ^2 . The signal model A might represent the voltage output of a DC power supply, and the noise $w[n]$ might model the error introduced by an inaccurate voltmeter as successive measurements are taken. Hence, $x[n]$ represents the noise corrupted observations of the power supply output. Now, even though the power supply should generate a constant voltage of A , in practice the true voltage will vary slightly as time progresses. This is due to the effects of temperature, component aging, etc., on the circuitry. Hence, a more accurate measurement model would be

$$x[n] = A[n] + w[n]$$

where $A[n]$ is the true voltage at time n . However, with this model our estimation problem becomes considerably more complicated since we will need to estimate $A[n]$ for $n = 0, 1, \dots, N-1$ instead of just the single parameter A . To underscore the difficulty assume that we model the voltage $A[n]$ as a sequence of unknown deterministic parameters. Then, the MVU estimator of $A[n]$ is easily shown to be

$$\hat{A}[n] = x[n].$$

The estimates will be inaccurate due to a lack of averaging, and in fact the variability will be identical to that of the noise or $\text{var}(\hat{A}[n]) = \sigma^2$. This estimator is undesirable in that it allows estimates such as those shown in Figure 13.1. We can expect that if the power supply is set for $A = 10$ volts, then the true voltage will be near this and the variation over time will be slow (otherwise it's time to buy a new power supply!). An example of the true voltage is given in Figure 13.1 and is seen to vary about 10 volts. Successive samples of $A[n]$ will not be too different, leading us to conclude that they display a high degree of "correlation." This reasoning naturally leads us to consider $A[n]$ as a realization of a random process with a mean of 10 and some correlation between samples. The imposition of a correlation constraint will prevent the estimate of $A[n]$ from fluctuating too wildly in time. Thus, we will consider $A[n]$ to be a realization of a random process to be estimated, for which Bayesian approaches are appropriate. This type of modeling was used in Chapter 12 in our discussion of Wiener filtering. There the signal to be estimated was termed $s[n]$, and it was assumed to be zero mean. In keeping with this standard notation, we now adopt $s[n]$ as our notation as opposed to $\theta[n]$. Also, because of the zero mean assumption, $s[n]$ will represent the signal model for $A[n] - 10$. Once we specify the signal model for zero mean $s[n]$, it is easily modified to accommodate nonzero mean processes by adding $E(s[n])$ to it. We will always assume that the mean is *known*.




Figure 13.1 True voltage and MVU estimator

A simple model for $s[n]$ which allows us to specify the correlation between samples is the *first-order Gauss-Markov process*

$$s[n] = as[n - 1] + u[n] \quad n \geq 0 \quad (13.1)$$

where $u[n]$ is WGN with variance σ_u^2 , $s[-1] \sim \mathcal{N}(\mu_s, \sigma_s^2)$, and $s[-1]$ is independent of $u[n]$ for all $n \geq 0$. (The reader should not confuse the Gauss-Markov *process* with the model considered in the Gauss-Markov *theorem* since they are different.) The noise $u[n]$ is termed the *driving or excitation noise* since $s[n]$ may be viewed as the output of a linear time invariant system driven by $u[n]$. In the control literature the system is referred to as the *plant*, and $u[n]$ is termed the *plant noise* [Jaswinski 1970]. The model of (13.1) is also called the *dynamical* or *state model*. The current output $s[n]$ depends only on the state of the system at the previous time, or $s[n-1]$, and the current input $u[n]$. The state of a system at time n_0 is generally defined to be the amount of information, which together with the input for $n \geq n_0$ determines the output for $n \geq n_0$ [Chen 1970]. Clearly, the state is $s[n-1]$, and it summarizes the effect of all past inputs to the system. We will henceforth refer to (13.1) as the Gauss-Markov model, where it is understood to be first order.

The signal model of (13.1) resembles an AR(1) process except that the signal starts at $n = 0$, and hence may not be WSS. We shall see shortly that as $n \rightarrow \infty$, so that the effect of the initial condition is negligible, the process is actually WSS and may be regarded as an AR(1) process with filter parameter $a[1] = -a$. A typical realization of $s[n]$ is shown in Figure 13.2 for $a = 0.98$, $\sigma_u^2 = 0.1$, $\mu_s = 5$, $\sigma_s^2 = 1$. Note that the mean starts off at about 5 but quickly decreases to zero. This behavior is just the transient

response of the system to the large initial sample $s[-1] \approx \mu_s = 5$. Also, the samples are heavily correlated. We may quantify these results by determining the mean and covariance of $s[n]$. Then, a complete statistical description will have been specified since $s[n]$ is a Gaussian process, as we will now show.

First, we express $s[n]$ as a function of the initial condition, and the inputs as

$$\begin{aligned} s[0] &= as[-1] + u[0] \\ s[1] &= as[0] + u[1] \\ &= a^2s[-1] + au[0] + u[1] \\ &\text{etc.} \end{aligned}$$

In general, we have

$$s[n] = a^{n+1}s[-1] + \sum_{k=0}^n a^k u[n-k] \quad (13.2)$$

and we see that $s[n]$ is a linear function of the initial condition and the driving noise inputs from the initial time to the present time. Since these random variables are all independent and Gaussian, $s[n]$ is Gaussian, and it may further be shown that $s[n]$ is a Gaussian random process (see Problem 13.1). Now, the mean follows from (13.2) as

$$E(s[n]) = a^{n+1}E(s[-1]) \quad (13.3)$$

$$= a^{n+1}\mu_s. \quad (13.4)$$

The covariance between samples $s[m]$ and $s[n]$ is, from (13.2) and (13.4),

$$\begin{aligned} c_s[m, n] &= E[(s[m] - E(s[m]))(s[n] - E(s[n]))] \\ &= E \left[\left(a^{m+1}(s[-1] - \mu_s) + \sum_{k=0}^m a^k u[m-k] \right) \right. \\ &\quad \cdot \left. \left(a^{n+1}(s[-1] - \mu_s) + \sum_{l=0}^n a^l u[n-l] \right) \right] \\ &= a^{m+n+2}\sigma_s^2 + \sum_{k=0}^m \sum_{l=0}^n a^{k+l} E(u[m-k]u[n-l]). \end{aligned}$$

But

$$E(u[m-k]u[n-l]) = \sigma_u^2 \delta[l - (n-m+k)]$$

or

$$E(u[m-k]u[n-l]) = \begin{cases} \sigma_u^2 & l = n-m+k \\ 0 & \text{otherwise.} \end{cases}$$

Thus, for $m \geq n$

$$\begin{aligned} c_s[m, n] &= a^{m+n+2}\sigma_s^2 + \sigma_u^2 \sum_{k=m-n}^m a^{2k+n-m} \\ &= a^{m+n+2}\sigma_s^2 + \sigma_u^2 a^{m-n} \sum_{k=0}^n a^{2k} \quad (13.5) \end{aligned}$$




Figure 13.2 Typical realization of first order Gauss-Markov process

and of course $c_s[m, n] = c_s[n, m]$ for $m < n$. Note that the variance is

$$\begin{aligned} \text{var}(s[n]) &= c_s[n, n] \\ &= a^{2n+2} \sigma_s^2 + \sigma_u^2 \sum_{k=0}^n a^{2k}. \end{aligned} \quad (13.6)$$

Clearly, $s[n]$ is not WSS since the mean depends on n and the covariance depends on m and n , not the difference. However, as $n \rightarrow \infty$, we have from (13.4) and (13.5)

$$\begin{aligned} E(s[n]) &\rightarrow 0 \\ c_s[m, n] &\rightarrow \frac{\sigma_u^2 a^{m-n}}{1 - a^2} \end{aligned}$$

since $|a| < 1$, which is necessary for stability of the process. Otherwise, the mean and variance would increase exponentially with n . Thus, as $n \rightarrow \infty$, the mean is zero and the covariance becomes an ACF with

$$r_{ss}[k] = c_s[m, n]|_{m-n=k} = \frac{\sigma_u^2}{1 - a^2} a^k \quad k \geq 0$$

which is recognized as the ACF of an AR(1) process. It is even possible that $s[n]$ will be WSS for $n \geq 0$ if the initial conditions are suitably chosen (see Problem 13.2). In any event it is seen from (13.5) that by properly choosing a the process can be made heavily correlated ($|a| \rightarrow 1$) or not correlated ($|a| \rightarrow 0$). This is also apparent from (13.1). As an illustration, for the process described earlier we plot the mean, variance, and steady-state covariance or ACF in Figure 13.3.

Because of the special form of the Gauss-Markov process, the mean and variance can also be expressed recursively. This is useful for conceptualization purposes as well as for extending the results to the vector Gauss-Markov process. The mean and variance are obtained directly from (13.1) as

$$E(s[n]) = aE(s[n-1]) + E(u[n])$$

or

$$E(s[n]) = aE(s[n-1]) \quad (13.7)$$

and

$$\begin{aligned} \text{var}(s[n]) &= E[(s[n] - E(s[n]))^2] \\ &= E[(as[n-1] + u[n] - aE(s[n-1]))^2] \\ &= a^2 \text{var}(s[n-1]) + \sigma_u^2 \end{aligned} \quad (13.8)$$


where we have used $E(u[n]s[n-1]) = 0$. This follows from (13.2) because $s[n-1]$ depends only on $\{s[-1], u[0], \dots, u[n-1]\}$ and they are independent of $u[n]$ by assumption. The equations (13.7) and (13.8) are termed the *mean and variance propagation equations*. A covariance propagation equation is explored in Problem 13.4. Note from (13.8) that the variance is decreased due to the effect of $as[n-1]$ but increased due to $u[n]$. In steady-state or as $n \rightarrow \infty$, these effects balance each other to yield the variance $\sigma_u^2/(1-a^2)$ (just let $\text{var}(s[n-1]) = \text{var}(s[n])$ in (13.8) and solve).

The ease with which we were able to obtain the mean and variance propagation equations resulted from the dependence of $s[n]$ on only the previous sample $s[n-1]$ and the input $u[n]$. Consider now a p th-order Gauss-Markov process expressed in a notation reminiscent of an AR(p) process or

$$s[n] = - \sum_{k=1}^p a[k]s[n-k] + u[n]. \quad (13.9)$$

Because $s[n]$ now depends on the p previous samples, the mean and variance propagation equations become more complicated. To extend our previous results we first note that the state of the system at time n is $\{s[n-1], s[n-2], \dots, s[n-p]\}$ since the previous p samples together with $u[n]$ determine the output. We thus define the *state vector* as

$$s[n-1] = \begin{bmatrix} s[n-p] \\ s[n-p+1] \\ \vdots \\ s[n-1] \end{bmatrix}. \quad (13.10)$$

**Figure 13.3** Statistics of first-order Gauss-Markov process

With this definition we can rewrite (13.9) in the form

$$\begin{bmatrix} s[n-p+1] \\ s[n-p+2] \\ \vdots \\ s[n-1] \\ s[n] \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ -a[p] & -a[p-1] & -a[p-2] & \dots & -a[1] \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} s[n-p] \\ s[n-p+1] \\ \vdots \\ s[n-2] \\ s[n-1] \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}}_{\mathbf{B}} u[n]$$

where the additional $(p-1)$ equations are just identities. Hence, using the definition of the state vector, we have

$$\mathbf{s}[n] = \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}u[n] \quad (13.11)$$

where $\mathbf{A}$ is a $p \times p$ nonsingular matrix (termed the *state transition matrix*) and $\mathbf{B}$ is a $p \times 1$ vector. Now we have the desired form (compare this to (13.1)) in which the vector signal $\mathbf{s}[n]$ is easily computed based on its value at the previous time instant, the state vector, and the input. This is termed the *vector Gauss-Markov model*. A final level of generality allows the input $u[n]$ to be an $r \times 1$ vector so that, as shown in Figure 13.4, we have a model for a vector signal as the output of a linear time invariant system ($\mathbf{A}$ and $\mathbf{B}$ are constant matrices) excited by a vector input. In Figure 13.4a $\mathcal{H}(z)$ is the $p \times r$ matrix system function. Summarizing, our general vector Gauss-Markov model takes the form

$$\mathbf{s}[n] = \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] \quad n \geq 0 \quad (13.12)$$

where $\mathbf{A}$, $\mathbf{B}$ are constant matrices with dimension $p \times p$ and $p \times r$, respectively, $\mathbf{s}[n]$ is the $p \times 1$ signal vector, and $\mathbf{u}[n]$ is the $r \times 1$ driving noise vector. We will from time to time refer to (13.12) as the *state model*. The statistical assumptions are that


1. The input $\mathbf{u}[n]$ is a vector WGN sequence, i.e., $\mathbf{u}[n]$ is a sequence of uncorrelated jointly Gaussian vectors with $E(\mathbf{u}[n]) = \mathbf{0}$. As a result, we have that

$$E(\mathbf{u}[m]\mathbf{u}^T[n]) = \mathbf{0} \quad m \neq n$$

and the covariance of $\mathbf{u}[n]$ is

$$E(\mathbf{u}[n]\mathbf{u}^T[n]) = \mathbf{Q},$$

where $\mathbf{Q}$ is an $r \times r$ positive definite matrix. Note that the vector samples are independent due to the jointly Gaussian assumption.

**Figure 13.4** Vector Gauss-Markov signal system model

2. The initial state or $\mathbf{s}[-1]$ is a random vector with

$$\mathbf{s}[-1] \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$$

and is independent of $\mathbf{u}[n]$ for all $n \geq 0$.

We now illustrate with an example.

Example 13.1 - Two DC Power Supplies

Recalling the introductory example, consider now the model for the outputs of two DC power supplies that vary with time. If we assume that the outputs are independent (in a functional sense) of each other, then a reasonable model would be the scalar model of (13.1) for each output or

$$\begin{aligned} s_1[n] &= a_1 s_1[n-1] + u_1[n] \\ s_2[n] &= a_2 s_2[n-1] + u_2[n] \end{aligned}$$

where $s_1[-1] \sim \mathcal{N}(\mu_{s_1}, \sigma_{s_1}^2)$, $s_2[-1] \sim \mathcal{N}(\mu_{s_2}, \sigma_{s_2}^2)$, $u_1[n]$ is WGN with variance $\sigma_{u_1}^2$, $u_2[n]$ is WGN with variance $\sigma_{u_2}^2$, and all random variables are independent of each other. Considering $\mathbf{s}[n] = [s_1[n] \ s_2[n]]^T$ as the vector parameter to be estimated, we have the model

$$\underbrace{\begin{bmatrix} s_1[n] \\ s_2[n] \end{bmatrix}}_{\mathbf{s}[n]} = \underbrace{\begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} s_1[n-1] \\ s_2[n-1] \end{bmatrix}}_{\mathbf{s}[n-1]} + \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} u_1[n] \\ u_2[n] \end{bmatrix}}_{\mathbf{u}[n]}$$

where $\mathbf{u}[n]$ is vector WGN with zero mean and covariance

$$\begin{aligned} E(\mathbf{u}[m]\mathbf{u}^T[n]) &= \begin{bmatrix} E(u_1[m]u_1[n]) & E(u_1[m]u_2[n]) \\ E(u_2[m]u_1[n]) & E(u_2[m]u_2[n]) \end{bmatrix} \\ &= \begin{bmatrix} \sigma_{u_1}^2 & 0 \\ 0 & \sigma_{u_2}^2 \end{bmatrix} \delta[m-n] \end{aligned}$$

so that

$$\mathbf{Q} = \begin{bmatrix} \sigma_{u_1}^2 & 0 \\ 0 & \sigma_{u_2}^2 \end{bmatrix}$$

and

$$\begin{aligned} \mathbf{s}[-1] &= \begin{bmatrix} s_1[-1] \\ s_2[-1] \end{bmatrix} \\ &\sim \mathcal{N}\left(\begin{bmatrix} \mu_{s_1} \\ \mu_{s_2} \end{bmatrix}, \begin{bmatrix} \sigma_{s_1}^2 & 0 \\ 0 & \sigma_{s_2}^2 \end{bmatrix}\right). \end{aligned}$$

If, on the other hand, the two outputs were generated from the same source (maybe some of the circuitry was shared), then they would undoubtedly be correlated. As we will see shortly, we could model this by letting any of the matrices $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{Q}$ be nondiagonal (see (13.26)). $\diamond$

We now complete our discussion of the vector Gauss-Markov model by deriving its statistical properties. The computations are simple extensions of those for the scalar model. First, we determine an explicit expression for $\mathbf{s}[n]$. From (13.12) we have

$$\begin{aligned} \mathbf{s}[0] &= \mathbf{A}\mathbf{s}[-1] + \mathbf{B}\mathbf{u}[0] \\ \mathbf{s}[1] &= \mathbf{A}\mathbf{s}[0] + \mathbf{B}\mathbf{u}[1] \\ &= \mathbf{A}^2\mathbf{s}[-1] + \mathbf{A}\mathbf{B}\mathbf{u}[0] + \mathbf{B}\mathbf{u}[1] \\ &\text{etc.} \end{aligned}$$

In general, we have

$$\mathbf{s}[n] = \mathbf{A}^{n+1}\mathbf{s}[-1] + \sum_{k=0}^n \mathbf{A}^k \mathbf{B} \mathbf{u}[n-k] \quad (13.13)$$

where $\mathbf{A}^0 = \mathbf{I}$. It is seen that $\mathbf{s}[n]$ is a linear function of the initial condition and the driving noise inputs. As a result, $\mathbf{s}[n]$ is a Gaussian random process. It remains only to determine the mean and covariance. From (13.13)

$$\begin{aligned} E(\mathbf{s}[n]) &= \mathbf{A}^{n+1}E(\mathbf{s}[-1]) \\ &= \mathbf{A}^{n+1}\boldsymbol{\mu}_s. \end{aligned} \quad (13.14)$$

The covariance is

$$\begin{aligned}
\mathbf{C}_s[m, n] &= E[(\mathbf{s}[m] - E(\mathbf{s}[m]))(\mathbf{s}[n] - E(\mathbf{s}[n]))^T] \\
&= E \left[\left(\mathbf{A}^{m+1}(\mathbf{s}[-1] - \boldsymbol{\mu}_s) + \sum_{k=0}^m \mathbf{A}^k \mathbf{B} \mathbf{u}[m-k] \right) \right. \\
&\quad \cdot \left. \left(\mathbf{A}^{n+1}(\mathbf{s}[-1] - \boldsymbol{\mu}_s) + \sum_{l=0}^n \mathbf{A}^l \mathbf{B} \mathbf{u}[n-l] \right)^T \right] \\
&= \mathbf{A}^{m+1} \mathbf{C}_s \mathbf{A}^{n+1T} + \sum_{k=0}^m \sum_{l=0}^n \mathbf{A}^k \mathbf{B} E(\mathbf{u}[m-k] \mathbf{u}^T[n-l]) \mathbf{B}^T \mathbf{A}^{lT}.
\end{aligned}$$

But

$$\begin{aligned}
E(\mathbf{u}[m-k] \mathbf{u}^T[n-l]) &= \mathbf{Q} \delta[l - (n - m + k)] \\
&= \begin{cases} \mathbf{Q} & l = n - m + k \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

Thus, for $m \geq n$

$$\mathbf{C}_s[m, n] = \mathbf{A}^{m+1} \mathbf{C}_s \mathbf{A}^{n+1T} + \sum_{k=m-n}^m \mathbf{A}^k \mathbf{B} \mathbf{Q} \mathbf{B}^T \mathbf{A}^{n-m+kT} \quad (13.15)$$

and for $m < n$

$$\mathbf{C}_s[m, n] = \mathbf{C}_s^T[n, m].$$

The covariance matrix for $\mathbf{s}[n]$ is

$$\begin{aligned}
\mathbf{C}[n] &= \mathbf{C}_s[n, n] \\
&= \mathbf{A}^{n+1} \mathbf{C}_s \mathbf{A}^{n+1T} + \sum_{k=0}^n \mathbf{A}^k \mathbf{B} \mathbf{Q} \mathbf{B}^T \mathbf{A}^{kT}.
\end{aligned} \quad (13.16)$$

Also, note that the mean and covariance propagation equations can be written as

$$E(\mathbf{s}[n]) = \mathbf{A} E(\mathbf{s}[n-1]) \quad (13.17)$$

$$\mathbf{C}[n] = \mathbf{A} \mathbf{C}[n-1] \mathbf{A}^T + \mathbf{B} \mathbf{Q} \mathbf{B}^T \quad n \geq 0 \quad (13.18)$$

which follow from (13.12) (see Problem 13.5). As in the scalar case, the covariance matrix decreases due to the $\mathbf{A} \mathbf{C}[n-1] \mathbf{A}^T$ term but increases due to the $\mathbf{B} \mathbf{Q} \mathbf{B}^T$ term. (It can be shown that for a stable process the eigenvalues of $\mathbf{A}$ must be less than 1 in magnitude. See Problem 13.6.) Steady-state properties similar to those of the scalar Gauss-Markov model are evident from (13.14) and (13.16). As $n \rightarrow \infty$, the mean will converge to zero or

$$E(\mathbf{s}[n]) = \mathbf{A}^{n+1} \boldsymbol{\mu}_s \rightarrow \mathbf{0}$$

since the eigenvalues of $\mathbf{A}$ are all less than 1 in magnitude. Also, from (13.16) it can be shown that as $n \rightarrow \infty$ (see Problem 13.7)

$$\mathbf{A}^{n+1} \mathbf{C}_s \mathbf{A}^{n+1^T} \rightarrow \mathbf{0}$$

so that

$$\mathbf{C}[n] \rightarrow \mathbf{C} = \sum_{k=0}^{\infty} \mathbf{A}^k \mathbf{B} \mathbf{Q} \mathbf{B}^T \mathbf{A}^{k^T}. \quad (13.19)$$

It is interesting to note that the steady-state covariance is also the solution of (13.18) if we set $\mathbf{C}[n-1] = \mathbf{C}[n] = \mathbf{C}$ in (13.18). Then, the steady-state covariance satisfies

$$\mathbf{C} = \mathbf{A} \mathbf{C} \mathbf{A}^T + \mathbf{B} \mathbf{Q} \mathbf{B}^T \quad (13.20)$$

and (13.19) is the solution, as can be verified by direct substitution. This is known as the *Lyapunov equation*.

Although in our definition of the Gauss-Markov model we assumed that the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{Q}$ did not depend on n , it is perfectly permissible to do so, and in some cases quite useful. Similar expressions for the mean and covariance can be developed. One major difference though is that the process may not attain a statistical steady-state.

We now summarize the model and its properties.

Theorem 13.1 (Vector Gauss-Markov Model) *The Gauss-Markov model for a $p \times 1$ vector signal $\mathbf{s}[n]$ is*

$$\mathbf{s}[n] = \mathbf{A} \mathbf{s}[n-1] + \mathbf{B} \mathbf{u}[n] \quad n \geq 0. \quad (13.21)$$

The $\mathbf{A}$, $\mathbf{B}$ are known matrices having dimensions $p \times p$ and $p \times r$, respectively, and it is assumed that the eigenvalues of $\mathbf{A}$ are less than 1 in magnitude. The driving noise vector $\mathbf{u}[n]$ has dimension $r \times 1$ and is vector WGN or $\mathbf{u}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ with the $\mathbf{u}[n]$'s independent. The initial condition $\mathbf{s}[-1]$ is a $p \times 1$ random vector distributed according to $\mathbf{s}[-1] \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$ and is independent of the $\mathbf{u}[n]$'s. Then, the signal process is Gaussian with mean

$$E(\mathbf{s}[n]) = \mathbf{A}^{n+1} \boldsymbol{\mu}_s \quad (13.22)$$

and covariance for $m \geq n$

$$\begin{aligned} \mathbf{C}_s[m, n] &= E[(\mathbf{s}[m] - E(\mathbf{s}[m]))(\mathbf{s}[n] - E(\mathbf{s}[n]))^T] \\ &= \mathbf{A}^{m+1} \mathbf{C}_s \mathbf{A}^{n+1^T} + \sum_{k=m-n}^m \mathbf{A}^k \mathbf{B} \mathbf{Q} \mathbf{B}^T \mathbf{A}^{n-m+k^T} \end{aligned} \quad (13.23)$$

and for $m < n$

$$\mathbf{C}_s[m, n] = \mathbf{C}_s^T[n, m]$$

and covariance matrix

$$\mathbf{C}[n] = \mathbf{C}_s[n, n] = \mathbf{A}^{n+1} \mathbf{C}_s \mathbf{A}^{n+1^T} + \sum_{k=0}^n \mathbf{A}^k \mathbf{B} \mathbf{Q} \mathbf{B}^T \mathbf{A}^{k^T}. \quad (13.24)$$

The mean and covariance propagation equations are

$$E(s[n]) = \mathbf{A}E(s[n-1]) \quad (13.25)$$

$$\mathbf{C}[n] = \mathbf{A}\mathbf{C}[n-1]\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T. \quad (13.26)$$

13.4 Scalar Kalman Filter

The scalar Gauss-Markov signal model discussed in the previous section had the form

$$s[n] = as[n-1] + u[n] \quad n \geq 0.$$

We now describe a sequential MMSE estimator which will allow us to estimate $s[n]$ based on the data $\{x[0], x[1], \dots, x[n]\}$ as n increases. Such an operation is referred to as *filtering*. The approach computes the estimator $\hat{s}[n]$ based on the estimator for the previous time sample $\hat{s}[n-1]$ and so is recursive in nature. This is the so-called *Kalman filter*. As explained in the introduction, the versatility of the Kalman filter accounts for its widespread use. It can be applied to estimation of a scalar Gauss-Markov signal as well as to its vector extension. Furthermore, the data, which previously in all our discussions consisted of a *scalar* sequence such as $\{x[0], x[1], \dots, x[n]\}$, can be extended to vector observations or $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n]\}$. A common example occurs in array processing in which at each time instant we sample the outputs of a group of sensors. If we have M sensors, then each data sample $\mathbf{x}[n]$ or observation will be a $M \times 1$ vector. Three different levels of generality are now summarized in hierarchical order.

1. scalar state – scalar observation ($s[n-1], x[n]$)
2. vector state – scalar observation ($s[n-1], x[n]$)
3. vector state – vector observation ($s[n-1], \mathbf{x}[n]$).

In this section we discuss the first case, leaving the remaining ones to Section 13.6.

Consider the scalar state equation and the scalar observation equation

$$\begin{aligned} s[n] &= as[n-1] + u[n] \\ x[n] &= s[n] + w[n] \end{aligned} \quad (13.27)$$

where $u[n]$ is zero mean Gaussian noise with independent samples and $E(u^2[n]) = \sigma_u^2$, $w[n]$ is zero mean Gaussian noise with independent samples and $E(w^2[n]) = \sigma_w^2$. We further assume that $s[-1]$, $u[n]$, and $w[n]$ are all independent. Finally, we assume that $s[-1] \sim \mathcal{N}(\mu_s, \sigma_s^2)$. The noise process $w[n]$ differs from WGN only in that its variance is allowed to change with time. To simplify the derivation we will assume that $\mu_s = 0$, so that according to (13.4) $E(s[n]) = 0$ for $n \geq 0$. Later we will account for a nonzero initial signal mean. We wish to estimate $s[n]$ based on the observations $\{x[0], x[1], \dots, x[n]\}$ or to filter $x[n]$ to produce $\hat{s}[n]$. More generally, the estimator of $s[n]$ based on the observations $\{x[0], x[1], \dots, x[m]\}$ will be denoted by $\hat{s}[n|m]$. Our criterion of optimality will be the minimum Bayesian MSE or

$$E[(s[n] - \hat{s}[n|m])^2]$$

where the expectation is with respect to $p(x[0], x[1], \dots, x[n], s[n])$. But the MMSE estimator is just the mean of the posterior PDF or

$$\hat{s}[n|n] = E(s[n]|x[0], x[1], \dots, x[n]). \quad (13.28)$$

Using Theorem 10.2 with zero means this becomes

$$\hat{s}[n|n] = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x} \quad (13.29)$$

since $\theta = s[n]$ and $\mathbf{x} = [x[0] x[1] \dots x[n]]^T$ are jointly Gaussian. Because we are assuming Gaussian statistics for the signal and noise, the MMSE estimator is *linear* and is *identical* in algebraic form to the LMMSE estimator. The algebraic properties allow us to utilize the vector space approach to find the estimator. The implicit linear constraint does not detract from the generality since we already know that the optimal estimator is linear. Furthermore, if the Gaussian assumption is not valid, then the resulting estimator is still valid but can only be said to be the optimal LMMSE estimator. Returning to the sequential computation of (13.29), we note that if $x[n]$ is uncorrelated with $\{x[0], x[1], \dots, x[n-1]\}$, then from (13.28) and the orthogonality principle we will have (see Example 12.2)

$$\begin{aligned} \hat{s}[n|n] &= E(s[n]|x[0], x[1], \dots, x[n-1]) + E(s[n]|x[n]) \\ &= \hat{s}[n|n-1] + E(s[n]|x[n]) \end{aligned}$$

which has the desired sequential form. Unfortunately, the $x[n]$'s are correlated due to their dependence on $s[n]$, which is correlated from sample to sample. From our discussions in Chapter 12 of the sequential LMMSE estimator, we can use our vector space interpretation to determine the correction of the old estimator $\hat{s}[n|n-1]$ due to the observation of $x[n]$. Before doing so we will summarize some properties of the MMSE estimator that will be used.

1. The MMSE estimator of θ based on two uncorrelated data vectors, assuming jointly Gaussian statistics, is (see Section 11.4)

$$\begin{aligned} \hat{\theta} &= E(\theta|\mathbf{x}_1, \mathbf{x}_2) \\ &= E(\theta|\mathbf{x}_1) + E(\theta|\mathbf{x}_2) \end{aligned}$$

if θ is zero mean.

2. The MMSE estimator is additive in that if $\theta = \theta_1 + \theta_2$, then

$$\begin{aligned} \hat{\theta} &= E(\theta|\mathbf{x}) \\ &= E(\theta_1 + \theta_2|\mathbf{x}) \\ &= E(\theta_1|\mathbf{x}) + E(\theta_2|\mathbf{x}). \end{aligned}$$

With these properties we begin the derivation of (13.38)–(13.42). Let $\mathbf{X}[n] = [x[0] x[1] \dots x[n]]^T$ (we now use $\mathbf{X}$ as our notation to avoid confusion with our previous notation

$\mathbf{x}[n]$ which will subsequently be used for vector observations) and $\tilde{x}[n]$ denote the innovation. Recall that the innovation is the part of $x[n]$ that is uncorrelated with the previous samples $\{x[0], x[1], \dots, x[n-1]\}$ or

$$\tilde{x}[n] = x[n] - \hat{x}[n|n-1]. \quad (13.30)$$

This is because by the orthogonality principle $\hat{x}[n|n-1]$ is the MMSE estimator of $x[n]$ based on the data $\{x[0], x[1], \dots, x[n-1]\}$, the error or $\tilde{x}[n]$ being orthogonal (uncorrelated) with the data. The data $\mathbf{X}[n-1], \tilde{x}[n]$ are equivalent to the original data set since $x[n]$ may be recovered from

$$\begin{aligned} x[n] &= \tilde{x}[n] + \hat{x}[n|n-1] \\ &= \tilde{x}[n] + \sum_{k=0}^{n-1} a_k x[k] \end{aligned}$$

where the a_k 's are the optimal weighting coefficients of the MMSE estimator of $x[n]$ based on $\{x[0], x[1], \dots, x[n-1]\}$. Now we can rewrite (13.28) as

$$\hat{s}[n|n] = E(s[n]|\mathbf{X}[n-1], \tilde{x}[n])$$

and because $\mathbf{X}[n-1]$ and $\tilde{x}[n]$ are uncorrelated, we have from property 1 that

$$\hat{s}[n|n] = E(s[n]|\mathbf{X}[n-1]) + E(s[n]|\tilde{x}[n]).$$

But $E(s[n]|\mathbf{X}[n-1])$ is the prediction of $s[n]$ based on the previous data, and we denote it by $\hat{s}[n|n-1]$. Explicitly the prediction is, from (13.1) and property 2,

$$\begin{aligned} \hat{s}[n|n-1] &= E(s[n]|\mathbf{X}[n-1]) \\ &= E(as[n-1] + u[n]|\mathbf{X}[n-1]) \\ &= aE(s[n-1]|\mathbf{X}[n-1]) \\ &= a\hat{s}[n-1|n-1] \end{aligned}$$

since $E(u[n]|\mathbf{X}[n-1]) = 0$. This is because

$$E(u[n]|\mathbf{X}[n-1]) = E(u[n]) = 0$$

since $u[n]$ is independent of $\{x[0], x[1], \dots, x[n-1]\}$. This follows by noting that $u[n]$ is independent of all $w[n]$, and from (13.2) $s[0], s[1], \dots, s[n-1]$ are linear combinations of the random variables $\{u[0], u[1], \dots, u[n-1], s[-1]\}$ and these are also independent of $u[n]$. We now have that

$$\hat{s}[n|n] = \hat{s}[n|n-1] + E(s[n]|\tilde{x}[n]) \quad (13.31)$$

where

$$\hat{s}[n|n-1] = a\hat{s}[n-1|n-1].$$

To determine $E(s[n]|\hat{x}[n])$ we note that it is the MMSE estimator of $s[n]$ based on $\hat{x}[n]$. As such, it is linear, and because of the zero mean assumption of $s[n]$, it takes the form

$$\begin{aligned} E(s[n]|\hat{x}[n]) &= K[n]\hat{x}[n] \\ &= K[n](x[n] - \hat{x}[n|n-1]) \end{aligned}$$

where

$$K[n] = \frac{E(s[n]\hat{x}[n])}{E(\hat{x}^2[n])}. \quad (13.32)$$

This follows from the general MMSE estimator for jointly Gaussian θ and x

$$\hat{\theta} = C_{\theta x} C_{xx}^{-1} x = \frac{E(\theta x)}{E(x^2)} x.$$

But $x[n] = s[n] + w[n]$, so that by property 2

$$\begin{aligned} \hat{x}[n|n-1] &= \hat{s}[n|n-1] + \hat{w}[n|n-1] \\ &= \hat{s}[n|n-1] \end{aligned}$$

since $\hat{w}[n|n-1] = 0$ due to $w[n]$ being independent of $\{x[0], x[1], \dots, x[n-1]\}$. Thus,

$$E(s[n]|\hat{x}[n]) = K[n](x[n] - \hat{s}[n|n-1])$$

and from (13.31) we now have

$$\hat{s}[n] = \hat{s}[n|n-1] + K[n](x[n] - \hat{s}[n|n-1]) \quad (13.33)$$

where

$$\hat{s}[n|n-1] = a\hat{s}[n-1|n-1]. \quad (13.34)$$

It remains only to determine the gain factor $K[n]$. From (13.32) the gain factor is

$$K[n] = \frac{E[s[n](x[n] - \hat{s}[n|n-1])]}{E[(x[n] - \hat{s}[n|n-1])^2]}.$$

To evaluate this we need the results

1.

$$E[s[n](x[n] - \hat{s}[n|n-1])] = E[(s[n] - \hat{s}[n|n-1])(x[n] - \hat{s}[n|n-1])]$$

2.

$$E[w[n](s[n] - \hat{s}[n|n-1])] = 0.$$

The first result is a consequence of the fact that the innovation

$$\begin{aligned} \tilde{x}[n] &= x[n] - \hat{x}[n|n-1] \\ &= x[n] - \hat{s}[n|n-1] \end{aligned} \quad (13.35)$$

is uncorrelated with the past data and hence with $\hat{s}[n|n-1]$, which is a linear combination of $x[0], x[1], \dots, x[n-1]$. The second result follows from $s[n]$ and $w[n]$ being uncorrelated and $w[n]$ being uncorrelated with the past data (since $w[n]$ is an uncorrelated process). Using these properties, the gain becomes

$$\begin{aligned} K[n] &= \frac{E[(s[n] - \hat{s}[n|n-1])(x[n] - \hat{s}[n|n-1])]}{E[(s[n] - \hat{s}[n|n-1] + w[n])^2]} \\ &= \frac{E[(s[n] - \hat{s}[n|n-1])^2]}{\sigma_u^2 + E[(s[n] - \hat{s}[n|n-1])^2]}. \end{aligned} \quad (13.36)$$

But the numerator is just the minimum MSE incurred when $s[n]$ is estimated based on the previous data or the minimum one-step prediction error. We will denote this by $M[n|n-1]$, so that

$$K[n] = \frac{M[n|n-1]}{\sigma_u^2 + M[n|n-1]}. \quad (13.37)$$

To evaluate the gain we need an expression for the minimum prediction error. Using (13.34)

$$\begin{aligned} M[n|n-1] &= E[(s[n] - \hat{s}[n|n-1])^2] \\ &= E[(as[n-1] + u[n] - \hat{s}[n|n-1])^2] \\ &= E[(a(s[n-1] - \hat{s}[n-1|n-1]) + u[n])^2]. \end{aligned}$$

We note that

$$E[(s[n-1] - \hat{s}[n-1|n-1])u[n]] = 0$$

since $s[n-1]$ depends on $\{u[0], u[1], \dots, u[n-1], s[-1]\}$, which are independent of $u[n]$, and $\hat{s}[n-1|n-1]$ depends on past data samples or $\{s[0] + w[0], s[1] + w[1], \dots, s[n-1] + w[n-1]\}$, which also are independent of $u[n]$. Thus,

$$M[n|n-1] = a^2 M[n-1|n-1] + \sigma_u^2.$$

Finally, we require a recursion for $M[n|n]$. Using (13.33), we have

$$\begin{aligned} M[n|n] &= E[(s[n] - \hat{s}[n|n])^2] \\ &= E[(s[n] - \hat{s}[n|n-1] - K[n](x[n] - \hat{s}[n|n-1]))^2] \\ &= E[(s[n] - \hat{s}[n|n-1])^2] - 2K[n]E[(s[n] - \hat{s}[n|n-1])(x[n] - \hat{s}[n|n-1])] \\ &\quad + K^2[n]E[(x[n] - \hat{s}[n|n-1])^2]. \end{aligned}$$

But the first term is $M[n|n-1]$, the second expectation is the numerator of $K[n]$, and the last expectation is the denominator of $K[n]$, as seen from (13.36). Hence, from (13.37)

$$\begin{aligned} M[n|n] &= M[n|n-1] - 2K^2[n](M[n|n-1] + \sigma_u^2) + K[n]M[n|n-1] \\ &= M[n|n-1] - 2K[n]M[n|n-1] + K[n]M[n|n-1] \\ &= (1 - K[n])M[n|n-1]. \end{aligned}$$

This completes the derivation of the scalar state–scalar observation Kalman filter. Although tedious, the final equations are actually quite simple and intuitive. We summarize them below. For $n \geq 0$

Prediction:

$$\hat{s}[n|n-1] = a\hat{s}[n-1|n-1]. \quad (13.38)$$

Minimum Prediction MSE:

$$M[n|n-1] = a^2 M[n-1|n-1] + \sigma_u^2. \quad (13.39)$$

Kalman Gain:

$$K[n] = \frac{M[n|n-1]}{\sigma_n^2 + M[n|n-1]}. \quad (13.40)$$

Correction:

$$\hat{s}[n|n] = \hat{s}[n|n-1] + K[n](x[n] - \hat{s}[n|n-1]). \quad (13.41)$$

Minimum MSE:

$$M[n|n] = (1 - K[n])M[n|n-1]. \quad (13.42)$$

Although derived for $\mu_s = 0$ so that $E(s[n]) = 0$, the same equations result for $\mu_s \neq 0$ (see Appendix 13A). Hence, to initialize the equations we use $\hat{s}[-1|-1] = E(s[-1]) = \mu_s$ and $M[-1|-1] = \sigma_s^2$ since this amounts to the estimation of $s[-1]$ without any data. A block diagram of the Kalman filter is given in Figure 13.5. It is interesting to note that the dynamical model for the signal is an integral part of the estimator. Furthermore, we may view the output of the gain block as an estimator $\hat{u}[n]$ of $u[n]$. From (13.38) and (13.41) the signal estimate is

$$\hat{s}[n|n] = a\hat{s}[n-1|n-1] + \hat{u}[n]$$

where $\hat{u}[n] = K[n](x[n] - \hat{s}[n|n-1])$. To the extent that this estimate is approximately $u[n]$, we will have $\hat{s}[n|n] \approx s[n]$, as desired. We now consider an example to illustrate the flow of the Kalman filter.

Example 13.2 - Scalar State–Scalar Observation Kalman Filter

The signal model is a first-order Gauss-Markov process

$$s[n] = \frac{1}{2}s[n-1] + u[n] \quad n \geq 0$$

where $s[-1] \sim \mathcal{N}(0, 1)$ and $\sigma_u^2 = 2$. We assume that the signal is observed in noise $w[n]$, so that the data are $x[n] = s[n] + w[n]$, where $w[n]$ is zero mean Gaussian noise




Figure 13.5 Scalar state-scalar observation Kalman filter and relationship to dynamic model

with independent samples, a variance of $\sigma_n^2 = (1/2)^n$, and independent of $s[-1]$ and $u[n]$ for $n \geq 0$. We initialize the filter with

$$\begin{aligned}\hat{s}[-1|-1] &= E(s[-1]) = 0 \\ M[-1|-1] &= E[(s[-1] - \hat{s}[-1|-1])^2] \\ &= E(s^2[-1]) = 1.\end{aligned}$$

According to (13.38) and (13.39), we first predict the $s[0]$ sample to obtain

$$\begin{aligned}\hat{s}[0|-1] &= a\hat{s}[-1|-1] \\ &= \frac{1}{2}(0) = 0 \\ M[0|-1] &= a^2 M[-1|-1] + \sigma_u^2 \\ &= \frac{1}{4}(1) + 2 = \frac{9}{4}.\end{aligned}$$

Next, as $x[0]$ is observed, we correct our predicted estimate by using (13.40) and (13.41)

to yield

$$\begin{aligned} K[0] &= \frac{M[0|-1]}{\sigma_0^2 + M[0|-1]} \\ &= \frac{\frac{9}{4}}{1 + \frac{9}{4}} = \frac{9}{13} \end{aligned}$$

$$\begin{aligned} \hat{s}[0|0] &= \hat{s}[0|-1] + K[0](x[0] - \hat{s}[0|-1]) \\ &= 0 + \frac{9}{13}(x[0] - 0) = \frac{9}{13}x[0]. \end{aligned}$$

Next, the minimum MSE is updated using (13.42).

$$\begin{aligned} M[0|0] &= (1 - K[0])M[0|-1] \\ &= \left(1 - \frac{9}{13}\right)\frac{9}{4} = \frac{9}{13}. \end{aligned}$$


For $n = 1$ we obtain the results

$$\begin{aligned} \hat{s}[1|0] &= \frac{9}{26}x[0] \\ M[1|0] &= \frac{113}{52} \\ K[1] &= \frac{113}{129} \\ \hat{s}[1|1] &= \frac{9}{26}x[0] + \frac{113}{129}\left(x[1] - \frac{9}{26}x[0]\right) \\ M[1|1] &= \frac{452}{1677} \end{aligned}$$

which the reader should verify. $\diamond$

We should note that *the same set of equations result if $E(s[-1]) \neq 0$* . In this case the mean of $s[n]$ will be nonzero since $E(s[n]) = a^{n+1}E(s[-1])$ (see (13.3)). This extension is explored in Problem 13.13. Recall that the reason for the zero mean assumption was to allow us to use the orthogonality principle. We now discuss some important properties of the Kalman filter. They are:

1. The Kalman filter extends the sequential MMSE estimator in Chapter 12 to the case where the unknown parameter evolves in time according to the dynamical model. In Chapter 12 we derived the equations for the sequential LMMSE estimator (see (12.47)–(12.49)), which are identical in form to those for the sequential MMSE estimator which assumes Gaussian statistics. In particular, for the scalar

**Figure 13.6** Kalman filter as time varying filter

case with $h[n] = 1$ and $\theta[n]$ replaced by $s[n]$, those equations become

$$\hat{s}[n] = \hat{s}[n-1] + K[n](x[n] - \hat{s}[n-1]) \quad (13.43)$$

$$K[n] = \frac{M[n-1]}{\sigma_n^2 + M[n-1]} \quad (13.44)$$

$$M[n] = (1 - K[n])M[n-1]. \quad (13.45)$$

The Kalman filter will reduce to these equations when the parameter to be estimated does not evolve in time. This follows by assuming the driving noise to be zero or $\sigma_u^2 = 0$ and also $a = 1$. Then, the state equation becomes from (13.1) $s[n] = s[n-1]$ or explicitly $s[n] = s[-1] = \theta$. The parameter to be estimated is a constant, which is modeled as the realization of a random variable. Then, the prediction is $\hat{s}[n|n-1] = \hat{s}[n-1|n-1]$ with minimum MSE $M[n|n-1] = M[n-1|n-1]$, so that we can omit the prediction stage of the Kalman filter. This says that the prediction is just the last estimate of $s[n]$. The correction stage reduces to (13.43)–(13.45) since we can let $\hat{s}[n|n] = \hat{s}[n]$, $\hat{s}[n|n-1] = \hat{s}[n-1]$, $M[n|n] = M[n]$, and $M[n|n-1] = M[n-1]$.

2. No matrix inversions are required. This should be compared to the *batch* method of estimating $\theta = s[n]$ as

$$\hat{\theta} = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x}$$

where $\mathbf{x} = [x[0] \ x[1] \ \dots \ x[n]]^T$. The use of this formula requires us to invert $\mathbf{C}_{xx}$ for each sample of $s[n]$ to be estimated. And in fact, the dimension of the matrix is $(n+1) \times (n+1)$, becoming larger with n .

3. The Kalman filter is a time varying linear filter. Note from (13.38) and (13.41) that

$$\begin{aligned} \hat{s}[n|n] &= a\hat{s}[n-1|n-1] + K[n](x[n] - a\hat{s}[n-1|n-1]) \\ &= a(1 - K[n])\hat{s}[n-1|n-1] + K[n]x[n]. \end{aligned}$$

This is a first-order recursive filter with time varying coefficients as shown in Figure 13.6.




Figure 13.7 Prediction and correction minimum MSE

4. The Kalman filter provides its own performance measure. From (13.42) the minimum Bayesian MSE is computed as an integral part of the estimator. Also, the error measure may be computed off-line, i.e., before any data are collected. This is because $M[n|n]$ depends only on (13.39) and (13.40), which are independent of the data. We will see later that for the extended Kalman filter the minimum MSE sequence must be computed on-line (see Section 13.7).
5. The prediction stage increases the error, while the correction stage decreases it. There is an interesting interplay between the minimum prediction MSE and the minimum MSE. If as $n \rightarrow \infty$ a steady-state condition is achieved, then $M[n|n]$ and $M[n|n-1]$ each become constant with $M[n|n-1] > M[n-1|n-1]$ (see Problem 13.14). Hence, the error will increase after the prediction stage. When the new data sample is obtained, we correct the estimate, which decreases the MSE according to (13.42) (since $K[n] < 1$). An example of this is shown in Figure 13.7, in which $a = 0.99$, $\sigma_u^2 = 0.1$, $\sigma_n^2 = 0.9^{n+1}$, and $M[-1|-1] = 1$. We will say more about this in the next section when we discuss the relationship of the Kalman filter to the Wiener filter.
6. Prediction is an integral part of the Kalman filter. It is seen from (13.38) that to determine the best filtered estimate of $s[n]$ we employ predictions. We can find the best one-step prediction of $s[n]$ based on $\{x[0], x[1], \dots, x[n-1]\}$ from (13.38). If we desire the best two-step prediction, we can obtain it easily by noting that this is the best estimate of $s[n+1]$ based on $\{x[0], x[1], \dots, x[n-1]\}$. To find this we let $\sigma_n^2 \rightarrow \infty$, implying that $x[n]$ is so noisy that the Kalman filter will not use it. Then, $\hat{s}[n+1|n-1]$ is just the optimal two-step prediction. To evaluate this we have from (13.38)

$$\hat{s}[n+1|n] = a\hat{s}[n|n]$$





Figure 13.8 Innovation-driven Kalman filter

Figure 13.9 Orthogonality (uncorrelated property) of innovation sequence $\tilde{x}(n)$

and since $K[n] \rightarrow 0$, we have from (13.41) that $\hat{s}[n|n] = \hat{s}[n|n-1]$, where $\hat{s}[n|n-1] = a\hat{s}[n-1|n-1]$. Thus, the optimal two-step prediction is

$$\begin{aligned}\hat{s}[n+1|n-1] &= \hat{s}[n+1|n] \\ &= a\hat{s}[n|n] \\ &= a\hat{s}[n|n-1] \\ &= a^2\hat{s}[n-1|n-1].\end{aligned}$$

This can be generalized to the l -step predictor, as shown in Problem 13.15.

7. The Kalman filter is driven by the uncorrelated innovation sequence and in steady-state can also be viewed as a whitening filter. Note from (13.38) and (13.41) that

$$\hat{s}[n|n] = a\hat{s}[n-1|n-1] + K[n](x[n] - \hat{s}[n|n-1])$$

so that the input to the Kalman filter is the innovation sequence $\tilde{x}[n] = x[n] - \hat{s}[n|n-1]$ (see (13.35)) as shown in Figure 13.8. We know from our discussion of the vector space approach that $\tilde{x}[n]$ is uncorrelated with $\{x[0], x[1], \dots, x[n-1]\}$, which translates into a sequence of uncorrelated random variables as shown in Figure 13.9. Alternatively, if we view $\tilde{x}[n]$ as the Kalman filter output and if the filter attains steady-state, then it becomes a linear time invariant whitening filter as shown in Figure 13.10. This is discussed in detail in the next section.




Figure 13.10 Whitening filter interpretation of Kalman filter

8. The Kalman filter is optimal in that it minimizes the Bayesian MSE for each estimator $\hat{s}[n]$. If the Gaussian assumption is not valid, then it is still the optimal *linear* MMSE estimator as described in Chapter 12.

All these properties of the Kalman filter carry over to the vector state case, except for property 2, if the observations are also vectors.

13.5 Kalman Versus Wiener Filters

The causal infinite length Wiener filter described in Chapter 12 produced an estimate of $s[n]$ as the output of a linear time invariant filter or

$$\hat{s}[n] = \sum_{k=0}^{\infty} h[k]x[n-k].$$

The estimator of $s[n]$ is based on the present data sample and the infinite past. To determine the filter impulse response $h[k]$ analytically we needed to assume that the signal $s[n]$ and noise $w[n]$ were WSS processes, so that the Wiener-Hopf equations could be solved (see Problem 12.16). In the Kalman filter formulation the signal and noise need not be WSS. The variance of $w[n]$ may change with n , and furthermore, $s[n]$ will only be WSS as $n \rightarrow \infty$. Additionally, the Kalman filter produces estimates based on only the data samples from 0 to n , not the infinite past as assumed by the infinite length Wiener filter. The two filters will, however, be the same as $n \rightarrow \infty$ if $\sigma_n^2 = \sigma^2$. This is because $s[n]$ will approach statistical steady-state as shown in Section 13.3, i.e., it will become an AR(1) process, and the estimator will then be based on the present

and the infinite past. Since the Kalman filter will approach a linear time invariant filter, we can let $K[n] \rightarrow K[\infty]$, $M[n|n] \rightarrow M[\infty]$, and $M[n|n-1] \rightarrow M_p[\infty]$, where $M_p[\infty]$ is the steady-state one-step prediction error. To find the steady-state Kalman filter we need to first find $M[\infty]$. From (13.39), (13.40), and (13.42) we have

$$\begin{aligned} M[\infty] &= \left(1 - \frac{M_p[\infty]}{\sigma^2 + M_p[\infty]}\right) M_p[\infty] \\ &= \frac{\sigma^2 M_p[\infty]}{M_p[\infty] + \sigma^2} \\ &= \frac{\sigma^2 (a^2 M[\infty] + \sigma_u^2)}{a^2 M[\infty] + \sigma_u^2 + \sigma^2} \end{aligned} \quad (13.46)$$

which must be solved for $M[\infty]$. The resulting equation is termed the steady-state *Ricatti equation* and is seen to be quadratic. Once $M[\infty]$ has been found, $M_p[\infty]$ can be determined and finally $K[\infty]$. Then, the steady-state Kalman filter takes the form of the first-order recursive filter shown in Figure 13.6 with $K[n]$ replaced by the constant $K[\infty]$. Note that a simple way of solving the Ricatti equation numerically is to run the Kalman filter until it converges. This will produce the desired time invariant filter. The steady-state filter will have the form

$$\begin{aligned} \hat{s}[n|n] &= a\hat{s}[n-1|n-1] + K[\infty](x[n] - a\hat{s}[n-1|n-1]) \\ &= a(1 - K[\infty])\hat{s}[n-1|n-1] + K[\infty]x[n] \end{aligned}$$

so that its steady-state transfer function is

$$\mathcal{H}_\infty(z) = \frac{K[\infty]}{1 - a(1 - K[\infty])z^{-1}}. \quad (13.47)$$

As an example, if $a = 0.9$, $\sigma_u^2 = 1$, $\sigma^2 = 1$, we will find from (13.46) that $M[\infty] = 0.5974$ (the other solution is negative). Hence, from (13.39) $M_p[\infty] = 1.4839$, and from (13.40) $K[\infty] = 0.5974$. The steady-state frequency response is

$$\begin{aligned} H_\infty(f) &= \mathcal{H}_\infty(\exp(j2\pi f)) \\ &= \frac{0.5974}{1 - 0.3623 \exp(-j2\pi f)} \end{aligned}$$

whose magnitude is shown in Figure 13.11 as a solid line versus the PSD of the steady-state signal

$$\begin{aligned} P_{ss}(f) &= \frac{\sigma_u^2}{|1 - a \exp(-j2\pi f)|^2} \\ &= \frac{1}{|1 - 0.9 \exp(-j2\pi f)|^2} \end{aligned}$$

shown as a dashed line. The same results would be obtained if the Wiener-Hopf equation had been solved for the causal infinite length Wiener filter. Hence, the steady-state




Figure 13.11 Signal PSD and steady-state Kalman filter magnitude response

Kalman filter is equivalent to the causal infinite length Wiener filter if the signal and the noise become WSS as $n \rightarrow \infty$.

Finally, the whitening filter property discussed in property 7 in the previous section may be verified for this example. The innovation is

$$\begin{aligned}\bar{x}[n] &= x[n] - \hat{s}[n|n-1] \\ &= x[n] - a\hat{s}[n-1|n-1].\end{aligned}\quad (13.48)$$

But in steady-state $\hat{s}[n|n]$ is the output of the filter with system function $\mathcal{H}_\infty(z)$ driven by $x[n]$. Thus, the system function relating the input $x[n]$ to the output $\bar{x}[n]$ is, from (13.48) and (13.47),

$$\begin{aligned}\mathcal{H}_w(z) &= 1 - az^{-1}\mathcal{H}_\infty(z) \\ &= 1 - \frac{az^{-1}K[\infty]}{1 - a(1 - K[\infty])z^{-1}} \\ &= \frac{1 - az^{-1}}{1 - a(1 - K[\infty])z^{-1}}.\end{aligned}$$

For this example we have the whitening filter frequency response

$$H_w(f) = \frac{1 - 0.9 \exp(-j2\pi f)}{1 - 0.3623 \exp(-j2\pi f)}$$

whose magnitude is plotted in Figure 13.12 as a solid line along with the PSD of $x[n]$,




Figure 13.12 Whitening filter property of steady-state Kalman filter

shown as a dashed line. Note that the PSD of $x[n]$ is

$$\begin{aligned}
 P_{xx}(f) &= P_{ss}(f) + \sigma^2 \\
 &= \frac{\sigma_u^2}{|1 - a \exp(-j2\pi f)|^2} + \sigma^2 \\
 &= \frac{\sigma_u^2 + \sigma^2 |1 - a \exp(-j2\pi f)|^2}{|1 - a \exp(-j2\pi f)|^2}
 \end{aligned}$$

which for this example is

$$P_{xx}(f) = \frac{1 + |1 - 0.9 \exp(-j2\pi f)|^2}{|1 - 0.9 \exp(-j2\pi f)|^2}.$$

The PSD at the output of the whitening filter is therefore

$$|H_w(f)|^2 P_{xx}(f) = \frac{1 + |1 - 0.9 \exp(-j2\pi f)|^2}{|1 - 0.3623 \exp(-j2\pi f)|^2}$$

which can be verified to be the constant PSD $P_{\hat{x}\hat{x}}(f) = 2.48$. In general, we have

$$P_{\hat{x}\hat{x}}(f) = \frac{\sigma_{\hat{x}}^2}{|H_w(f)|^2}$$

where $\sigma_{\hat{x}}^2$ is the variance of the innovation. Thus, as shown in Figure 13.13, the output PSD of a filter with frequency response $H_w(f)$ is a flat PSD with height $\sigma_{\hat{x}}^2$.




Figure 13.13 Input and output PSDs of steady-state Kalman whitening filter

13.6 Vector Kalman Filter

The scalar state–scalar observation Kalman filter is easily generalized. The two generalizations are to replace $s[n]$ by $\mathbf{s}[n]$, where $\mathbf{s}[n]$ obeys the Gauss–Markov model described in Theorem 13.1, and to replace the scalar observation $x[n]$ by the vector observation $\mathbf{x}[n]$. The first generalization will produce the vector state–scalar observation Kalman filter, while the second leads to the most general form, the vector state–vector observation Kalman filter. In either case the state model is, from Theorem 13.1,

$$\mathbf{s}[n] = \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] \quad n \geq 0$$

where $\mathbf{A}$, $\mathbf{B}$ are known $p \times p$ and $p \times r$ matrices, $\mathbf{u}[n]$ is vector WGN with $\mathbf{u}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$, $\mathbf{s}[-1] \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$, and $\mathbf{s}[-1]$ is independent of the $\mathbf{u}[n]$'s. The vector state–scalar observation Kalman filter assumes that the observations follow the Bayesian linear model (see Section 10.6) with the added assumption that the noise covariance matrix is *diagonal*. Thus, for the n th data sample we have

$$x[n] = \mathbf{h}^T[n]\mathbf{s}[n] + w[n] \quad (13.49)$$

where $\mathbf{h}[n]$ is a known $p \times 1$ vector and $w[n]$ is zero mean Gaussian noise with uncorrelated samples, with variance σ_w^2 , and also independent of $\mathbf{s}[-1]$ and $\mathbf{u}[n]$. The data model of (13.49) is called the *observation or measurement equation*. An example is given in Section 13.8, where we wish to track the coefficients of a random time varying FIR filter. For that example, the state is comprised of the impulse response values. The Kalman filter for this setup is derived in exactly the same manner as for the scalar state case. The derivation is included in Appendix 13A. We now summarize the results. The reader should note the similarity to (13.38)–(13.42).

Prediction:

$$\hat{\mathbf{s}}[n|n-1] = \mathbf{A}\hat{\mathbf{s}}[n-1|n-1]. \quad (13.50)$$

Minimum Prediction MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n-1] = \mathbf{A}\mathbf{M}[n-1|n-1]\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T. \quad (13.51)$$

Kalman Gain Vector ($p \times 1$):

$$\mathbf{K}[n] = \frac{\mathbf{M}[n|n-1]\mathbf{h}[n]}{\sigma_n^2 + \mathbf{h}^T[n]\mathbf{M}[n|n-1]\mathbf{h}[n]}. \quad (13.52)$$

Correction:

$$\hat{\mathbf{s}}[n|n] = \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](x[n] - \mathbf{h}^T[n]\hat{\mathbf{s}}[n|n-1]). \quad (13.53)$$

Minimum MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n] = (\mathbf{I} - \mathbf{K}[n]\mathbf{h}^T[n])\mathbf{M}[n|n-1] \quad (13.54)$$

where the mean square error matrices are defined as

$$\mathbf{M}[n|n] = E[(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n])(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n])^T] \quad (13.55)$$

$$\mathbf{M}[n|n-1] = E[(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1])(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1])^T]. \quad (13.56)$$

To initialize the equations we use $\hat{\mathbf{s}}[-1|-1] = E(\mathbf{s}[-1]) = \boldsymbol{\mu}_s$ and $\mathbf{M}[-1|-1] = \mathbf{C}_s$. The order of the recursion is identical to the scalar state-scalar observation case. Also, as before, no matrix inversions are required, but this is no longer true when we consider the vector observation case. If $\mathbf{A} = \mathbf{I}$ and $\mathbf{B} = \mathbf{0}$, then the equations are identical to the sequential LMMSE estimator (see (12.47)–(12.49)). This is because the signal models are identical, the signal being constant in time. An application example is given in Section 13.8.

Finally, we generalize the Kalman filter to the case of *vector observations*, which is quite common in practice. An example is in array processing, where at each time instant we sample the output of an array of M sensors. Then, the observations are $\mathbf{x}[n] = [x_1[n] x_2[n] \dots x_M[n]]^T$, where $x_i[n]$ is the output of the i th sensor at time n . The observations are modeled using the Bayesian linear model

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where *each* vector observation has this form. Indexing the quantities by n and replacing $\boldsymbol{\theta}$ by $\mathbf{s}$, we have as our observation model

$$\mathbf{x}[n] = \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n] \quad (13.57)$$

where $\mathbf{H}[n]$ is a known $M \times p$ matrix, $\mathbf{x}[n]$ is an $M \times 1$ observation vector, and $\mathbf{w}[n]$ is a $M \times 1$ *observation* noise sequence. The $\mathbf{w}[n]$'s are independent of each other and of $\mathbf{u}[n]$ and $\mathbf{s}[-1]$, and $\mathbf{w}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{C}[n])$. Except for the dependence of the covariance matrix on n , $\mathbf{w}[n]$ can be thought of as vector WGN. For the array processing problem $\mathbf{s}[n]$ represents a vector of p transmitted signals, which are modeled as random,

and $\mathbf{H}[n]$ models the linear transformation due to the medium. The medium may be modeled as time varying since $\mathbf{H}[n]$ depends on n . Hence, $\mathbf{H}[n]\mathbf{s}[n]$ is the signal output at the M sensors. Also, the sensor outputs are corrupted by noise $\mathbf{w}[n]$. The statistical assumptions on $\mathbf{w}[n]$ indicate the noise is correlated from sensor to sensor at the same time instant with covariance $\mathbf{C}[n]$, and this correlation varies with time. From time instant to time instant, however, the noise samples are independent since $E(\mathbf{w}[i]\mathbf{w}^T[j]) = \mathbf{0}$ for $i \neq j$. With this data model we can derive the vector state–vector observation Kalman filter, the most general estimator. Because of the large number of assumptions required, we summarize the results as a theorem.

Theorem 13.2 (Vector Kalman Filter) *The $p \times 1$ signal vector $\mathbf{s}[n]$ evolves in time according to the Gauss-Markov model*

$$\mathbf{s}[n] = \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] \quad n \geq 0$$

where $\mathbf{A}, \mathbf{B}$ are known matrices of dimension $p \times p$ and $p \times r$, respectively. The driving noise vector $\mathbf{u}[n]$ has the PDF $\mathbf{u}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ and is independent from sample to sample, so that $E(\mathbf{u}[m]\mathbf{u}^T[n]) = \mathbf{0}$ for $m \neq n$ ($\mathbf{u}[n]$ is vector WGN). The initial state vector $\mathbf{s}[-1]$ has the PDF $\mathbf{s}[-1] \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$ and is independent of $\mathbf{u}[n]$.

The $M \times 1$ observation vectors $\mathbf{x}[n]$ are modeled by the Bayesian linear model

$$\mathbf{x}[n] = \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n] \quad n \geq 0$$

where $\mathbf{H}[n]$ is a known $M \times p$ observation matrix (which may be time varying) and $\mathbf{w}[n]$ is an $M \times 1$ observation noise vector with PDF $\mathbf{w}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{C}[n])$ and is independent from sample to sample, so that $E(\mathbf{w}[m]\mathbf{w}^T[n]) = \mathbf{0}$ for $m \neq n$. (If $\mathbf{C}[n]$ did not depend on n , then $\mathbf{w}[n]$ would be vector WGN.)

The MMSE estimator of $\mathbf{s}[n]$ based on $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n]\}$ or

$$\hat{\mathbf{s}}[n|n] = E(\mathbf{s}[n]|\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n])$$

can be computed sequentially in time using the following recursion:

Prediction:

$$\hat{\mathbf{s}}[n|n-1] = \mathbf{A}\hat{\mathbf{s}}[n-1|n-1]. \quad (13.58)$$

Minimum Prediction MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n-1] = \mathbf{A}\mathbf{M}[n-1|n-1]\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T. \quad (13.59)$$

Kalman Gain Matrix ($p \times M$):

$$\mathbf{K}[n] = \mathbf{M}[n|n-1]\mathbf{H}^T[n] (\mathbf{C}[n] + \mathbf{H}[n]\mathbf{M}[n|n-1]\mathbf{H}^T[n])^{-1}. \quad (13.60)$$

Correction:

$$\hat{\mathbf{s}}[n|n] = \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]). \quad (13.61)$$

Minimum MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n] = (\mathbf{I} - \mathbf{K}[n]\mathbf{H}[n])\mathbf{M}[n|n-1]. \quad (13.62)$$

The recursion is initialized by $\hat{\mathbf{s}}[-1|-1] = \boldsymbol{\mu}_s$, and $\mathbf{M}[-1|-1] = \mathbf{C}_s$.

All the comments of the scalar state-scalar observation Kalman filter apply here as well, with the exception of the need for matrix inversions. We now require the inversion of an $M \times M$ matrix to find the Kalman gain. If the dimension of the state vector p is less than the dimension of the observation vector M , a more efficient implementation of the Kalman filter can be obtained. Referred to as the *information form*, this alternative Kalman filter is described in [Anderson and Moore 1979].

Before concluding the discussion of linear Kalman filters, it is worthwhile to note that the Kalman filter summarized in the previous theorem is still not the most general one. It is, however, adequate for many practical problems. Extensions can be made by letting the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{Q}$ be time varying. Fortunately, the equations that result are identical to those of the previous theorem when we replace $\mathbf{A}$ by $\mathbf{A}[n]$, $\mathbf{B}$ by $\mathbf{B}[n]$, and $\mathbf{Q}$ by $\mathbf{Q}[n]$. Also, it is possible to extend the results to colored observation noise and to signal models with deterministic inputs (in addition to the driving noise). Finally, smoothing equations have also been derived based on the Kalman philosophy. These extensions are described in [Anderson and Moore 1979, Gelb 1974, Mendel 1987].

13.7 Extended Kalman Filter

In practice we are often faced with a state equation and/or an observation equation which is nonlinear. The previous approaches then are no longer valid. A simple example that will be explored in some detail in Section 13.8 is vehicle tracking. For this problem the observations or measurements are range estimates $\hat{R}[n]$ and bearing estimates $\hat{\beta}[n]$. If the vehicle state is the position (r_x, r_y) in Cartesian coordinates (it is assumed to travel in the x - y plane), then the noiseless measurements are related to the unknown parameters by

$$\begin{aligned} R[n] &= \sqrt{r_x^2[n] + r_y^2[n]} \\ \beta[n] &= \arctan \frac{r_y[n]}{r_x[n]}. \end{aligned}$$

Due to measurement errors, however, we obtain the estimates $\hat{R}[n]$ and $\hat{\beta}[n]$, which are assumed to be the true range and bearing plus measurement noise. Hence, we have for our measurements

$$\begin{aligned} \hat{R}[n] &= R[n] + w_R[n] \\ \hat{\beta}[n] &= \beta[n] + w_\beta[n] \end{aligned} \quad (13.63)$$

or

$$\begin{aligned}\hat{R}[n] &= \sqrt{r_x^2[n] + r_y^2[n]} + w_R[n] \\ \hat{\beta}[n] &= \arctan \frac{r_y[n]}{r_x[n]} + w_\beta[n].\end{aligned}$$

Clearly, we *cannot* express these in the linear model form or as

$$\mathbf{x}[n] = \mathbf{H}[n]\boldsymbol{\theta}[n] + \mathbf{w}[n]$$

where $\boldsymbol{\theta}[n] = [r_x[n] \ r_y[n]]^T$. The observation equation is *nonlinear*.

An example of a nonlinear state equation occurs if we assume that the vehicle is traveling in a given direction at a known fixed speed and we choose polar coordinates, range and bearing, to describe the state. (Note that this choice would render the measurement equation linear as described by (13.63)). Then, ignoring the driving noise, the state equation becomes

$$\begin{aligned}r_x[n] &= v_x n \Delta + r_x[0] \\ r_y[n] &= v_y n \Delta + r_y[0]\end{aligned}\tag{13.64}$$

where (v_x, v_y) is the known velocity, Δ is the time interval between samples, and $(r_x[0], r_y[0])$ is the initial position. This can be expressed alternatively as

$$\begin{aligned}r_x[n] &= r_x[n-1] + v_x \Delta \\ r_y[n] &= r_y[n-1] + v_y \Delta\end{aligned}$$

so that the range becomes

$$\begin{aligned}R[n] &= \sqrt{r_x^2[n-1] + r_y^2[n-1] + 2v_x \Delta r_x[n-1] + 2v_y \Delta r_y[n-1] + (v_x^2 + v_y^2) \Delta^2} \\ &= \sqrt{R^2[n-1] + 2R[n-1] \Delta (v_x \cos \beta[n-1] + v_y \sin \beta[n-1]) + (v_x^2 + v_y^2) \Delta^2}.\end{aligned}$$

This is clearly very nonlinear in range and bearing. In general, we may be faced with sequential state estimation where the state and/or observation equations are nonlinear. Then, instead of our linear Kalman filter models

$$\begin{aligned}\mathbf{s}[n] &= \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] \\ \mathbf{x}[n] &= \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n]\end{aligned}$$

we would have

$$\mathbf{s}[n] = \mathbf{a}(\mathbf{s}[n-1]) + \mathbf{B}\mathbf{u}[n]\tag{13.65}$$

$$\mathbf{x}[n] = \mathbf{h}(\mathbf{s}[n]) + \mathbf{w}[n]\tag{13.66}$$

where $\mathbf{a}$ is a p -dimensional function and $\mathbf{h}$ is an M -dimensional function. The dimensions of the remaining matrices and vectors are the same as before. Now $\mathbf{a}(\mathbf{s}[n-1])$

represents the true physical model for the evolution of the state, while $\mathbf{u}[n]$ accounts for the modeling errors, unforeseen inputs, etc. Likewise, $\mathbf{h}(\mathbf{s}[n])$ represents the transformation from the state variables to the ideal observations (without noise). For this case the MMSE estimator is intractable. The only hope is an approximate solution based on linearizing $\mathbf{a}$ and $\mathbf{h}$, much the same as was done for nonlinear LS, where the data were nonlinearly related to the unknown parameters. The result of this linearization and the subsequent application of the *linear* Kalman filter of (13.58)–(13.62) results in the *extended Kalman filter*. It has *no optimality properties*, and its performance will depend on the accuracy of the linearization. Being a *dynamic* linearization there is no way to determine its performance beforehand.

Proceeding with the derivation, we linearize $\mathbf{a}(\mathbf{s}[n-1])$ about the estimate of $\mathbf{s}[n-1]$ or about $\hat{\mathbf{s}}[n-1|n-1]$. Likewise, we linearize $\mathbf{h}(\mathbf{s}[n])$ about the estimate of $\mathbf{s}[n]$ based on the previous data or $\hat{\mathbf{s}}[n|n-1]$ since from (13.61) we will need the linearized observation equation to determine $\hat{\mathbf{s}}[n|n]$. Hence, a first-order Taylor expansion yields

$$\begin{aligned}\mathbf{a}(\mathbf{s}[n-1]) &\approx \mathbf{a}(\hat{\mathbf{s}}[n-1|n-1]) \\ &\quad + \left. \frac{\partial \mathbf{a}}{\partial \mathbf{s}[n-1]} \right|_{\mathbf{s}[n-1]=\hat{\mathbf{s}}[n-1|n-1]} (\mathbf{s}[n-1] - \hat{\mathbf{s}}[n-1|n-1]) \\ \mathbf{h}(\mathbf{s}[n]) &\approx \mathbf{h}(\hat{\mathbf{s}}[n|n-1]) + \left. \frac{\partial \mathbf{h}}{\partial \mathbf{s}[n]} \right|_{\mathbf{s}[n]=\hat{\mathbf{s}}[n|n-1]} (\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1]).\end{aligned}$$

We let the Jacobians be denoted by

$$\begin{aligned}\mathbf{A}[n-1] &= \left. \frac{\partial \mathbf{a}}{\partial \mathbf{s}[n-1]} \right|_{\mathbf{s}[n-1]=\hat{\mathbf{s}}[n-1|n-1]} \\ \mathbf{H}[n] &= \left. \frac{\partial \mathbf{h}}{\partial \mathbf{s}[n]} \right|_{\mathbf{s}[n]=\hat{\mathbf{s}}[n|n-1]}\end{aligned}$$

so that the linearized state and observation equations become from (13.65) and (13.66)

$$\begin{aligned}\mathbf{s}[n] &= \mathbf{A}[n-1]\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] + (\mathbf{a}(\hat{\mathbf{s}}[n-1|n-1]) - \mathbf{A}[n-1]\hat{\mathbf{s}}[n-1|n-1]) \\ \mathbf{x}[n] &= \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n] + (\mathbf{h}(\hat{\mathbf{s}}[n|n-1]) - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]).\end{aligned}$$

The equations differ from our standard ones in that $\mathbf{A}$ is now time varying and both equations have *known* terms added to them. It is shown in Appendix 13B that the linear Kalman filter for this model, which is the *extended* Kalman filter, is

Prediction:

$$\hat{\mathbf{s}}[n|n-1] = \mathbf{A}(\hat{\mathbf{s}}[n-1|n-1]). \quad (13.67)$$

Minimum Prediction MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n-1] = \mathbf{A}[n-1]\mathbf{M}[n-1|n-1]\mathbf{A}^T[n-1] + \mathbf{BQB}^T. \quad (13.68)$$

Kalman Gain Matrix ($p \times M$):

$$\mathbf{K}[n] = \mathbf{M}[n|n-1]\mathbf{H}^T[n] (\mathbf{C}[n] + \mathbf{H}[n]\mathbf{M}[n|n-1]\mathbf{H}^T[n])^{-1}. \quad (13.69)$$

Correction:

$$\hat{\mathbf{s}}[n|n] = \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{h}(\hat{\mathbf{s}}[n|n-1])). \quad (13.70)$$

Minimum MSE Matrix ($p \times p$):

$$\mathbf{M}[n|n] = (\mathbf{I} - \mathbf{K}[n]\mathbf{H}[n])\mathbf{M}[n|n-1] \quad (13.71)$$

where

$$\begin{aligned} \mathbf{A}[n-1] &= \left. \frac{\partial \mathbf{a}}{\partial \mathbf{s}[n-1]} \right|_{\mathbf{s}[n-1] = \hat{\mathbf{s}}[n-1|n-1]} \\ \mathbf{H}[n] &= \left. \frac{\partial \mathbf{h}}{\partial \mathbf{s}[n]} \right|_{\mathbf{s}[n] = \hat{\mathbf{s}}[n|n-1]} \end{aligned}$$

Note that in contrast to the linear Kalman filter the gain and MSE matrices must be computed on-line, as they depend upon the state estimates via $\mathbf{A}[n-1]$ and $\mathbf{H}[n]$. Also, the use of the term MSE matrix is itself a misnomer since the MMSE estimator has not been implemented but only an approximation to it. In the next section we will apply the extended Kalman filter to vehicle tracking.


13.8 Signal Processing Examples

We now examine some common signal processing applications of the linear and extended Kalman filters.


Example 13.3 - Time Varying Channel Estimation

Many transmission channels can be characterized as being linear but *not* time invariant. These are referred to by various names such as *fading dispersive* channels or *fading multipath* channels. They arise in communication problems in which the troposphere is used as a medium or in sonar in which the ocean is used [Kennedy 1969]. In either case, the medium acts as a linear filter, causing an impulse at the input to appear as a continuous waveform at the output (the dispersive or multipath nature), as shown in Figure 13.14a. This effect is the result of a continuum of propagation paths, i.e., multipath, each of which delays and attenuates the input signal. Additionally, however, a sinusoid at the input will appear as a narrowband signal at the output or one whose amplitude is modulated (the fading nature), as shown in Figure 13.14b. This effect is due to the changing character of the medium, for example, the movement of the scatterers. A little thought will convince the reader that the channel is acting as a linear time varying filter. If we sample the output of the channel, then it can be shown that a good model is the low-pass tapped delay line model as shown in Figure 13.15 [Van Trees 1971]. The input-output description of this system is


$$y[n] = \sum_{k=0}^{p-1} h_n[k]v[n-k]. \quad (13.72)$$



(a) Multipath channel



(b) Fading channel

Figure 13.14 Input-output waveforms for fading and multipath channels**Figure 13.15** Tapped delay line channel model

This is really nothing more than an FIR filter with time-varying coefficients. To design effective communication or sonar systems it is necessary to have knowledge of these coefficients. Hence, the problem becomes one of estimating $h_n[k]$ based on the noise corrupted output of the channel

$$x[n] = \sum_{k=0}^{p-1} h_n[k]v[n-k] + w[n] \quad (13.73)$$

where $w[n]$ is observation noise. A similar problem was addressed in Example 4.3, except that there we assumed the filter coefficients to be time invariant. Consequently, the linear model could be applied to estimate the deterministic parameters. It is not possible to extend that approach to our current problem since there are too many parameters to estimate. To see why this is so we let $p = 2$ and assume that $v[n] = 0$ for $n < 0$. The observations are, from (13.73),

$$\begin{aligned} x[0] &= h_0[0]v[0] + h_0[1]v[-1] + w[0] = h_0[0]v[0] + w[0] \\ x[1] &= h_1[0]v[1] + h_1[1]v[0] + w[1] \\ x[2] &= h_2[0]v[2] + h_2[1]v[1] + w[2] \\ &\text{etc.} \end{aligned}$$

It is seen that for $n \geq 1$ we have two new parameters for each new data sample. Even without corrupting noise we cannot determine the tapped delay line weights. A way out of this problem is to realize that the weights will not change rapidly from sample to sample, as for example, in a slow-fading channel. For example, in Figure 13.14b this would correspond to an amplitude modulation (caused by the time variation of the weights) which is slow. Statistically, we may interpret the slow variation as a high degree of correlation between samples of the same tap weight. This observation naturally leads us to model the tap weights as random variables whose time variation is described by a Gauss-Markov model. The use of such a signal model allows us to fix the correlation between the successive values of a given tap weight in time. Hence, we suppose that the state vector is

$$\mathbf{h}[n] = \mathbf{A}\mathbf{h}[n-1] + \mathbf{u}[n]$$

where $\mathbf{h}[n] = [h_n[0] \ h_n[1] \ \dots \ h_n[p-1]]^T$, $\mathbf{A}$ is a known $p \times p$ matrix, and $\mathbf{u}[n]$ is vector WGN with covariance matrix $\mathbf{Q}$. (The reader should note that $\mathbf{h}[n]$ no longer refers to the observation vector as in (13.49) but is now the signal $\mathbf{s}[n]$.) A standard assumption that is made to simplify the modeling is that of *uncorrelated scattering* [Van Trees 1971]. It assumes that the tap weights are uncorrelated with each other and hence independent due to the jointly Gaussian assumption. As a result, we can let $\mathbf{A}$, $\mathbf{Q}$, and $\mathbf{C}_h$, the covariance matrix of $\mathbf{h}[-1]$, be diagonal matrices. The vector Gauss-Markov model then becomes p independent scalar models. The measurement model is, from (13.73),

$$x[n] = \underbrace{\begin{bmatrix} v[n] & v[n-1] & \dots & v[n-p+1] \end{bmatrix}}_{\mathbf{v}^T[n]} \mathbf{h}[n] + w[n]$$

where $w[n]$ is assumed to be WGN with variance σ^2 and the $v[n]$ sequence is assumed known (since we provide the input to the channel). We can now form the MMSE estimator for the tapped delay line weights recursively in time using the Kalman filter equations for a vector state and scalar observations. With obvious changes in notation we have from (13.50)–(13.54)

$$\begin{aligned}\hat{\mathbf{h}}[n|n-1] &= \mathbf{A}\hat{\mathbf{h}}[n-1|n-1] \\ \mathbf{M}[n|n-1] &= \mathbf{A}\mathbf{M}[n-1|n-1]\mathbf{A}^T + \mathbf{Q} \\ \mathbf{K}[n] &= \frac{\mathbf{M}[n|n-1]\mathbf{v}[n]}{\sigma^2 + \mathbf{v}^T[n]\mathbf{M}[n|n-1]\mathbf{v}[n]} \\ \hat{\mathbf{h}}[n|n] &= \hat{\mathbf{h}}[n|n-1] + \mathbf{K}[n](x[n] - \mathbf{v}^T[n]\hat{\mathbf{h}}[n|n-1]) \\ \mathbf{M}[n|n] &= (\mathbf{I} - \mathbf{K}[n]\mathbf{v}^T[n])\mathbf{M}[n|n-1]\end{aligned}$$

and is initialized by $\hat{\mathbf{h}}[-1|-1] = \boldsymbol{\mu}_h$, $\mathbf{M}[-1|-1] = \mathbf{C}_h$. As an example, we now implement the Kalman filter estimator for a tapped delay line having $p = 2$ weights. We assume a state model with

$$\begin{aligned}\mathbf{A} &= \begin{bmatrix} 0.99 & 0 \\ 0 & 0.999 \end{bmatrix} \\ \mathbf{Q} &= \begin{bmatrix} 0.0001 & 0 \\ 0 & 0.0001 \end{bmatrix}.\end{aligned}$$

A particular realization is shown in Figure 13.16, in which $h_n[0]$ is decaying to zero while $h_n[1]$ is fairly constant. This is because the mean of the weights will be zero in steady-state (see (13.4)). Due to the smaller value of $[\mathbf{A}]_{11}$, $h_n[0]$ will decay more rapidly. Also, note that the eigenvalues of $\mathbf{A}$ are just the diagonal elements and they are less than 1 in magnitude. For this tap weight realization and the input shown in Figure 13.17a, the output is shown in Figure 13.17b as determined from (13.72). When observation noise is added with $\sigma^2 = 0.1$, we have the channel output shown in Figure 13.17c. We next apply the Kalman filter with $\hat{\mathbf{h}}[-1|-1] = \mathbf{0}$ and $\mathbf{M}[-1|-1] = 100\mathbf{I}$, which were chosen to reflect little knowledge about the initial state. In the theoretical development of the Kalman filter the initial state estimate is given by the mean of $\mathbf{s}[-1]$. In practice this is seldom known, so that we usually just choose an arbitrary initial state estimate with a large initial MSE matrix to avoid “biasing” the Kalman filter towards that assumed state. The estimated tap weights are shown in Figure 13.18. After an initial transient the Kalman filter “locks on” to the true weights and tracks them closely. The Kalman filter gains are shown in Figure 13.19. They appear to attain a periodic steady-state, although this behavior is different than the usual steady-state discussed previously since $\mathbf{v}[n]$ varies with time and so true steady-state is never attained. Also, at times the gain is zero, as for example in $[\mathbf{K}]_1 = K_1[n]$ for $0 \leq n \leq 4$. This is because at these times $\mathbf{v}[n]$ is zero due to the zero input and thus the observations contain only noise. The Kalman filter ignores these data samples by forcing the gain to be zero. Finally, the minimum MSEs are shown in Figure 13.20 and are seen to decrease monotonically, although this generally will not be the case for a Kalman filter. $\diamond$





Figure 13.16 Realization of TDL coefficients

Example 13.4 - Vehicle Tracking

In this example we use an extended Kalman filter to track the position and velocity of a vehicle moving in a nominal given direction and at a nominal speed. The measurements are noisy versions of the range and bearing. Such a track is shown in Figure 13.21. In arriving at a model for the dynamics of the vehicle we assume a constant velocity, perturbed only by wind gusts, slight speed corrections, etc., as might occur in an aircraft. We model these perturbations as noise inputs, so that the velocity components in the x and y directions at time n are

$$\begin{aligned} v_x[n] &= v_x[n-1] + u_x[n] \\ v_y[n] &= v_y[n-1] + u_y[n]. \end{aligned} \tag{13.74}$$

**Figure 13.17** Input-output waveforms of channel




Figure 13.18 Kalman filter estimate

Without the noise perturbations $u_x[n], u_y[n]$ the velocities would be constant, and hence the vehicle would be modeled as traveling in a straight line as indicated by the dashed line in Figure 13.21. From the equations of motion the position at time n is

$$\begin{aligned} r_x[n] &= r_x[n-1] + v_x[n-1]\Delta \\ r_y[n] &= r_y[n-1] + v_y[n-1]\Delta \end{aligned} \quad (13.75)$$

where Δ is the time interval between samples. In this discretized model of the equations of motion the vehicle is modeled as moving at the velocity of the previous time instant and then changing abruptly at the next time instant, an approximation to the true continuous behavior. Now, we choose the signal vector as consisting of the position and




Figure 13.19 Kalman filter gains

velocity components or

$$s[n] = \begin{bmatrix} r_x[n] \\ r_y[n] \\ v_x[n] \\ v_y[n] \end{bmatrix}$$

and from (13.74) and (13.75) it is seen to satisfy

$$\underbrace{\begin{bmatrix} r_x[n] \\ r_y[n] \\ v_x[n] \\ v_y[n] \end{bmatrix}}_{s[n]} = \underbrace{\begin{bmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} r_x[n-1] \\ r_y[n-1] \\ v_x[n-1] \\ v_y[n-1] \end{bmatrix}}_{s[n-1]} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ u_x[n] \\ u_y[n] \end{bmatrix}}_{u[n]}. \quad (13.76)$$




Figure 13.20 Kalman filter minimum MSE




Figure 13.21 Typical track of vehicle moving in given direction at constant speed

The measurements are noisy observations of the range and bearing

$$\begin{aligned}R[n] &= \sqrt{r_x^2[n] + r_y^2[n]} \\ \beta[n] &= \arctan \frac{r_y[n]}{r_x[n]}\end{aligned}$$

or

$$\begin{aligned}\hat{R}[n] &= R[n] + w_R[n] \\ \hat{\beta}[n] &= \beta[n] + w_\beta[n].\end{aligned}\tag{13.77}$$

In general terms the observation equation of (13.77) is

$$\mathbf{x}[n] = \mathbf{h}(s[n]) + \mathbf{w}[n]$$

where $\mathbf{h}$ is the function

$$\mathbf{h}(s[n]) = \begin{bmatrix} \sqrt{r_x^2[n] + r_y^2[n]} \\ \arctan \frac{r_y[n]}{r_x[n]} \end{bmatrix}.$$

Unfortunately, the measurement vector is nonlinear in the signal parameters. To estimate the signal vector we will need to apply an extended Kalman filter (see (13.67)–(13.71)). Since the state equation of (13.76) is linear, we need only determine

$$\mathbf{H}[n] = \left. \frac{\partial \mathbf{h}}{\partial s[n]} \right|_{s[n] = \hat{s}[n|n-1]}$$

because $\mathbf{A}[n]$ is just $\mathbf{A}$ as given in (13.76). Differentiating the observation equation, we have the Jacobian

$$\frac{\partial \mathbf{h}}{\partial s[n]} = \begin{bmatrix} \frac{r_x[n]}{R[n]} & \frac{r_y[n]}{R[n]} & 0 & 0 \\ -\frac{r_y[n]}{R^2[n]} & \frac{r_x[n]}{R^2[n]} & 0 & 0 \end{bmatrix}.$$

Finally, we need to specify the covariances of the driving noise and observation noise. If we assume that the wind gusts, speed corrections, etc., are just as likely to occur in any direction and with the same magnitude, then it seems reasonable to assign the same variances to $u_x[n]$ and $u_y[n]$ and to assume that they are independent. Call the common variance σ_u^2 . Then, we have

$$\mathbf{Q} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_u^2 & 0 \\ 0 & 0 & 0 & \sigma_u^2 \end{bmatrix}.$$

The exact value to use for σ_u^2 should depend on the possible change in the velocity component from sample to sample since $u_x[n] = v_x[n] - v_x[n-1]$. This is just the

acceleration times Δ and should be derivable from the physics of the vehicle. In specifying the variances of the measurement noise we note that the measurement error can be thought of as the estimation error of $\hat{R}[n]$ and $\hat{\beta}[n]$ as seen from (13.77). We usually assume the estimation errors $w_R[n]$, $w_\beta[n]$ to be zero mean. Then, the variance of $w_R[n]$, for example, is $E(w_R^2[n]) = E[(\hat{R}[n] - R[n])^2]$. This variance is sometimes derivable but in most instances is not. One possibility is to assume that $E(w_R^2[n])$ does not depend on the PDF of $R[n]$, so that $E(w_R^2[n]) = E[(\hat{R}[n] - R[n])^2 | R[n]]$. Equivalently, we could regard $R[n]$ as a *deterministic* parameter so that the variance of $w_R[n]$ is just the classical estimator variance. As such, if $\hat{R}[n]$ were the MLE, then assuming long data records and/or high SNRs, we could assume that the variance attains the CRLB. Using this approach, we could then make use of the CRLB for range and bearing such as was derived in Examples 3.13 and 3.15 to set the variances. For simplicity we usually assume the estimation errors to be independent and the variances to be time invariant (although this is not always valid). Hence, we have

$$\mathbf{C}[n] = \mathbf{C} = \begin{bmatrix} \sigma_R^2 & 0 \\ 0 & \sigma_\beta^2 \end{bmatrix}.$$

In summary, the extended Kalman filter equations for this problem are, from (13.67)-(13.71),

$$\begin{aligned} \hat{\mathbf{s}}[n|n-1] &= \mathbf{A}\hat{\mathbf{s}}[n-1|n-1] \\ \mathbf{M}[n|n-1] &= \mathbf{A}\mathbf{M}[n-1|n-1]\mathbf{A}^T + \mathbf{Q} \\ \mathbf{K}[n] &= \mathbf{M}[n|n-1]\mathbf{H}^T[n] (\mathbf{C} + \mathbf{H}[n]\mathbf{M}[n|n-1]\mathbf{H}^T[n])^{-1} \\ \hat{\mathbf{s}}[n|n] &= \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{h}(\hat{\mathbf{s}}[n|n-1])) \\ \mathbf{M}[n|n] &= (\mathbf{I} - \mathbf{K}[n]\mathbf{H}[n])\mathbf{M}[n|n-1] \end{aligned}$$

where

$$\begin{aligned} \mathbf{A} &= \begin{bmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ \mathbf{Q} &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_u^2 & 0 \\ 0 & 0 & 0 & \sigma_u^2 \end{bmatrix} \\ \mathbf{x}[n] &= \begin{bmatrix} \hat{R}[n] \\ \hat{\beta}[n] \end{bmatrix} \\ \mathbf{h}(\mathbf{s}[n]) &= \begin{bmatrix} \sqrt{r_x^2[n] + r_y^2[n]} \\ \arctan \frac{r_y[n]}{r_x[n]} \end{bmatrix} \end{aligned}$$




Figure 13.22 Realization of vehicle track

$$\mathbf{H}[n] = \begin{bmatrix} \frac{r_x[n]}{\sqrt{r_x^2[n] + r_y^2[n]}} & \frac{r_y[n]}{\sqrt{r_x^2[n] + r_y^2[n]}} & 0 & 0 \\ \frac{-r_y[n]}{r_x^2[n] + r_y^2[n]} & \frac{r_x[n]}{r_x^2[n] + r_y^2[n]} & 0 & 0 \end{bmatrix} \bigg|_{\mathbf{s}[n] = \hat{\mathbf{s}}[n|n-1]}$$

$$\mathbf{C} = \begin{bmatrix} \sigma_R^2 & 0 \\ 0 & \sigma_\beta^2 \end{bmatrix}$$

and the initial conditions are $\hat{\mathbf{s}}[-1|-1] = \boldsymbol{\mu}_s$, $\mathbf{M}[-1|-1] = \mathbf{C}_s$. As an example, consider the ideal straight line trajectory shown in Figure 13.22 as a dashed line. The coordinates are given by

$$\begin{aligned} r_x[n] &= 10 - 0.2n \\ r_y[n] &= -5 + 0.2n \end{aligned}$$

for $n = 0, 1, \dots, 100$, where we have assumed $\Delta = 1$ for convenience. From (13.75) this trajectory assumes $v_x = -0.2$, $v_y = 0.2$. To accommodate a more realistic vehicle track we introduce driving or plant noise, so that the vehicle state is described by (13.76) with a driving noise variance of $\sigma_u^2 = 0.0001$. With an initial state of

$$\mathbf{s}[-1] = \begin{bmatrix} 10 \\ -5 \\ -0.2 \\ 0.2 \end{bmatrix}, \quad (13.78)$$

which is identical to that of the initial state of the straight line trajectory, a realization of the vehicle position $[r_x[n] \ r_y[n]]^T$ is shown in Figure 13.22 as the solid curve. The state equation of (13.76) has been used to generate the realization. Note that as time

increases, the true trajectory gradually deviates from the straight line. It can be shown that the variances of $v_x[n]$ and $v_y[n]$ will eventually increase to infinity (see Problem 13.22), causing $r_x[n]$ and $r_y[n]$ to quickly become unbounded. Thus, this modeling is valid for only a portion of the trajectory. The true range and bearing are shown in Figure 13.23. We assume the measurement noise variances to be $\sigma_R^2 = 0.1$ and $\sigma_\beta^2 = 0.01$, where β is measured in radians. In Figure 13.24 we compare the true trajectory with the noise corrupted one as obtained from

$$\begin{aligned}\hat{r}_x[n] &= \hat{R}[n] \cos \hat{\beta}[n] \\ \hat{r}_y[n] &= \hat{R}[n] \sin \hat{\beta}[n].\end{aligned}$$

To employ an extended Kalman filter we must specify an initial state estimate. In practice, it is unlikely that we will have knowledge of the position and speed. Thus, for the sake of illustration we choose an initial state that is quite far from the true one or

$$\hat{\mathbf{s}}[-1|-1] = \begin{bmatrix} 5 \\ 5 \\ 0 \\ 0 \end{bmatrix}$$

and so as not to “bias” the extended Kalman filter we assume a large initial MSE or $\mathbf{M}[-1|-1] = 100\mathbf{I}$. The results of an extended Kalman filter are shown in Figure 13.25 as the solid curve. Initially, because of the poor state estimate, the error is large. This is also reflected in the MSE curves in Figure 13.26 (actually these are only estimates based on our linearization). However, after about 20 samples the extended Kalman filter attains the track. Also, it is interesting to note that the minimum MSE does not monotonically decrease as it did in the previous example. On the contrary, it increases for part of the time. This is explained by contrasting the Kalman filter with the sequential LMMSE estimator in Chapter 12. For the latter we estimated the same parameter as we received more and more data. Consequently, the minimum MSE decreased or at worst remained constant. Here, however, each time we receive a new data sample we are estimating a new parameter. The increased uncertainty of the new parameter due to the influence of the driving noise input may be large enough to offset the knowledge gained by observing a new data sample, causing the minimum MSE to increase (see Problem 13.23). As a final remark, in this simulation the extended Kalman filter appeared to be quite tolerant of linearization errors due to a poor initial state estimate. In general, however, we cannot expect to be so fortunate. $\diamond$





Figure 13.23 Range and bearing of true vehicle track




Figure 13.24 True and observed vehicle tracks




Figure 13.25 True and extended Kalman filter estimate





Figure 13.26 "Minimum" MSEs for $r_x[n]$ and $r_y[n]$

References

- Anderson, B.D.O., J.B. Moore, *Optimal Filtering*, Prentice-Hall, Englewood Cliffs, N.J., 1979.
 Chen, C.T., *Introduction to Linear System Theory*, Holt, Rinehart, and Winston, New York, 1970.
 Gelb, A., *Applied Optimal Estimation*, M.I.T. Press, Cambridge, Mass., 1974.
 Jazwinski, A.H., *Stochastic Processes and Filtering Theory*, Academic Press, New York, 1970.
 Kennedy, R.S., *Fading Dispersive Communication Channels*, J. Wiley, New York, 1969.
 Mendel, J.M., *Lessons in Digital Estimation Theory*, Prentice-Hall, Englewood Cliffs, N.J., 1987.
 Van Trees, H.L., *Detection, Estimation, and Modulation Theory III*, J. Wiley, New York, 1971.

Problems

- 13.1 A random process is Gaussian if for arbitrary samples $\{s[n_1], s[n_2], \dots, s[n_k]\}$ and for any k the random vector $\mathbf{s} = [s[n_1] s[n_2] \dots s[n_k]]^T$ is distributed according to a multivariate Gaussian PDF. If $s[n]$ is given by (13.2), prove that it is a Gaussian random process.
- 13.2 Consider a scalar Gauss-Markov process. Show that if $\mu_s = 0$ and $\sigma_s^2 = \sigma_u^2/(1 - a^2)$, then the process will be WSS for $n \geq 0$ and explain why this is so.
- 13.3 Plot the mean, variance, and steady-state covariance of a scalar Gauss-Markov process if $a = 0.98$, $\sigma_u^2 = 0.1$, $\mu_s = 5$, and $\sigma_s^2 = 1$. What is the PSD of the steady-state process?
- 13.4 For a scalar Gauss-Markov process derive a covariance propagation equation, i.e., a formula relating $c_s[m, n]$ to $c_s[n, n]$ for $m \geq n$. To do so first show that

$$c_s[m, n] = a^{m-n} \text{var}(s[n])$$

for $m \geq n$.

- 13.5 Verify (13.17) and (13.18) for the covariance matrix propagation of a vector Gauss-Markov process.
- 13.6 Show that the mean of a vector Gauss-Markov process in general will grow with n if any eigenvalue of $\mathbf{A}$ is greater than 1 in magnitude. What happens to the steady-state mean if all the eigenvalues are less than 1 in magnitude? To simplify matters assume that $\mathbf{A}$ is symmetric, so that it can be written as $\mathbf{A} = \sum_{i=1}^p \lambda_i \mathbf{v}_i \mathbf{v}_i^T$, where $\mathbf{v}_i$ is the i th eigenvector of $\mathbf{A}$ and λ_i is the corresponding real eigenvalue.
- 13.7 Show that

$$\mathbf{A}^n \mathbf{C}_s \mathbf{A}^{n^T} \rightarrow 0$$

as $n \rightarrow \infty$ if all the eigenvalues of $\mathbf{A}$ are less than 1 in magnitude. As in Problem 13.6, assume that $\mathbf{A}$ is symmetric. Hint: Examine $\mathbf{e}_i^T \mathbf{A}^n \mathbf{C}_s \mathbf{A}^{n^T} \mathbf{e}_j = [\mathbf{A}^n \mathbf{C}_s \mathbf{A}^{n^T}]_{ij}$, where $\mathbf{e}_i$ is a vector of all zeros except for a one in the i th element.

13.8 Consider the recursive difference equation (for $q < p$)

$$r[n] = - \sum_{k=1}^p a[k]r[n-k] + u[n] + \sum_{k=1}^q b[k]u[n-k] \quad n \geq 0$$

where $u[n]$ is WGN with variance σ_u^2 . (In steady-state this would be an ARMA process.) Let the state vector be defined as

$$\mathbf{s}[n-1] = \begin{bmatrix} r[n-p] \\ r[n-p+1] \\ \vdots \\ r[n-1] \end{bmatrix}$$

where $\mathbf{s}[-1] \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$ and is independent of $u[n]$. Also, define the vector driving noise sequence as

$$\mathbf{u}[n] = \begin{bmatrix} u[n-q] \\ u[n-q+1] \\ \vdots \\ u[n] \end{bmatrix}.$$

Rewrite the process as in (13.11). Explain why this is not a vector Gauss-Markov model by examining the assumptions on $\mathbf{u}[n]$.

13.9 For Problem 13.8 show that the process may alternatively be expressed by the difference equations

$$\begin{aligned} s[n] &= - \sum_{k=1}^p a[k]s[n-k] + u[n] \\ r[n] &= s[n] + \sum_{k=1}^q b[k]s[n-k] \end{aligned}$$

for $n \geq 0$. Assume that $\mathbf{s}[-1] = [s[-p] \ s[-p+1] \ \dots \ s[-1]]^T \sim \mathcal{N}(\boldsymbol{\mu}_s, \mathbf{C}_s)$ and that we observe $x[n] = r[n] + w[n]$, where $w[n]$ is WGN with variance σ_w^2 , and $\mathbf{s}[-1]$, $u[n]$, and $w[n]$ are independent. Show how to set up the vector state-scalar observation Kalman filter.

13.10 Assume that we observe $x[n] = A + w[n]$ for $n = 0, 1, \dots$, where A is the realization of a random variable with PDF $\mathcal{N}(0, \sigma_A^2)$ and $w[n]$ is WGN with variance σ^2 . Using the scalar state-scalar observation Kalman filter find a sequential estimator of A based on $\{x[0], x[1], \dots, x[n]\}$ or $\hat{A}[n]$. Solve explicitly for $\hat{A}[n]$, the Kalman gain, and the minimum MSE.

13.11 In this problem we implement a scalar state-scalar observation Kalman filter (see (13.38)–(13.42)). A computer solution is advised. If $a = 0.9$, $\sigma_u^2 = 1$, $\mu_s = 0$, $\sigma_s^2 = 1$, find the Kalman gain and minimum MSE if

- a. $\sigma_n^2 = (0.9)^n$
- b. $\sigma_n^2 = 1$
- c. $\sigma_n^2 = (1.1)^n$.

Explain your results. Using a Monte Carlo computer simulation generate a realization of the signal and noise and apply your Kalman filter to estimation of the signal for all three cases. Plot the signal as well as the Kalman filter estimate.

- 13.12** For the scalar state–scalar observation Kalman filter assume that $\sigma_n^2 = 0$ for all n so that we observe $s[n]$ directly. Find the innovation sequence. Is it white?
- 13.13** In this problem we show that the same set of equations result for the scalar state–scalar observation Kalman filter even if $E(s[-1]) \neq 0$. To do so let $s'[n] = s[n] - E(s[n])$ and $x'[n] = x[n] - E(x[n])$, so that equations (13.38)–(13.42) apply for $s'[n]$. Now determine the corresponding equations for $s[n]$. Recall that the MMSE estimator of $\theta + c$ for c a constant is $\hat{\theta} + c$, where $\hat{\theta}$ is the MMSE of θ .
- 13.14** Prove that for the scalar state–scalar observation Kalman filter

$$M[n|n-1] > M[n-1|n-1]$$

for large enough n or in steady-state. Why is this reasonable?

- 13.15** Prove that the optimal l -step predictor for a scalar Gauss–Markov process $s[n]$ is

$$\hat{s}[n+l|n] = a^l \hat{s}[n|n]$$

where $\hat{s}[n|n]$ and $\hat{s}[n+l|n]$ are based on $\{x[0], x[1], \dots, x[n]\}$.

- 13.16** Find the transfer function of the steady-state, scalar state–scalar observation Kalman filter if $a = 0.8$, $\sigma_u^2 = 1$, and $\sigma^2 = 1$. Give its time domain form as a recursive difference equation.
- 13.17** For the scalar state–scalar observation Kalman filter let $a = 0.9$, $\sigma_u^2 = 1$, $\sigma^2 = 1$ and find the steady-state gain and minimum MSE by running the filter until convergence, i.e., compute equations (13.39), (13.40), and (13.42). Compare your results to those given in Section 13.5.
- 13.18** Assume we observe the data

$$x[k] = Ar^k + w[k]$$

for $k = 0, 1, \dots, n$, where A is the realization of a random variable with PDF $\mathcal{N}(\mu_A, \sigma_A^2)$, $0 < r < 1$, and the $w[k]$'s are samples of WGN with variance σ^2 . Furthermore, assume that A is independent of the $w[k]$'s. Find the sequential MMSE estimator of A based on $\{x[0], x[1], \dots, x[n]\}$.

- 13.19** Consider the vector state–vector observation Kalman filter for which $\mathbf{H}[n]$ is assumed to be invertible. If a particular observation is noiseless, so that $\mathbf{C}[n] = 0$, find $\hat{s}[n|n]$ and explain your results. What happens if $\mathbf{C}[n] \rightarrow \infty$?

- 13.20** Prove that the optimal l -step predictor for a vector Gauss-Markov process $\mathbf{s}[n]$ is

$$\hat{\mathbf{s}}[n+l|n] = \mathbf{A}^l \hat{\mathbf{s}}[n|n]$$

where $\hat{\mathbf{s}}[n|n]$ and $\hat{\mathbf{s}}[n+l|n]$ are based on $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n]\}$. Use the vector state–vector observation Kalman filter.

- 13.21** In this problem we set up an extended Kalman filter for the frequency tracking application. Specifically, we wish to track the frequency of a sinusoid in noise. The frequency is assumed to follow the model

$$f_0[n] = a f_0[n-1] + u[n] \quad n \geq 0$$

where $u[n]$ is WGN with variance σ_u^2 and $f_0[-1]$ is distributed according to $\mathcal{N}(\mu_{f_0}, \sigma_{f_0}^2)$ and is independent of $u[n]$. The observed data are

$$x[n] = \cos(2\pi f_0[n]) + w[n] \quad n \geq 0$$

where $w[n]$ is WGN with variance σ^2 and is independent of $u[n]$ and $f_0[-1]$. Write down the extended Kalman filter equations for this problem.

- 13.22** For the vehicle position model in Example 13.4 we had

$$v_x[n] = v_x[n-1] + u_x[n]$$

where $u_x[n]$ is WGN with variance σ_u^2 . Find the variance of $v_x[n]$ to show that it increases with n . What might be a more suitable model for $v_x[n]$? (See also [Anderson and Moore 1979] for a further discussion of the modeling issue.)

- 13.23** For the scalar state–scalar observation Kalman filter find an expression relating $M[n|n]$ to $M[n-1|n-1]$. Now, let $a = 0.9$, $\sigma_u^2 = 1$, and $\sigma_x^2 = n+1$. If $M[-1|-1] = 1$, determine $M[n|n]$ for $n \geq 0$. Explain your results.

Appendix 13A

Vector Kalman Filter Derivation

We derive the vector state–vector observation Kalman filter with the vector state–scalar observation being a special case. Theorem 13.2 contains the modeling assumptions. In our derivation we will assume $\mu_s = 0$, so that all random variables are zero mean. With this assumption the vector space viewpoint is applicable. For $\mu_s \neq 0$ it can be shown that the same equations result. The reason for this is a straightforward generalization of the comments made at the end of Appendix 12A (see also Problem 13.13).

The properties of MMSE estimators that we will use are

1. The MMSE estimator of θ based on two uncorrelated data samples $\mathbf{x}_1$ and $\mathbf{x}_2$, assuming jointly Gaussian statistics (see Section 11.4), is

$$\begin{aligned}\hat{\theta} &= E(\theta|\mathbf{x}_1, \mathbf{x}_2) \\ &= E(\theta|\mathbf{x}_1) + E(\theta|\mathbf{x}_2)\end{aligned}$$

if $\theta, \mathbf{x}_1, \mathbf{x}_2$ are zero mean.

2. The MMSE estimator is linear in that if $\theta = \mathbf{A}_1\theta_1 + \mathbf{A}_2\theta_2$, then

$$\begin{aligned}\hat{\theta} &= E(\theta|\mathbf{x}) \\ &= E(\mathbf{A}_1\theta_1 + \mathbf{A}_2\theta_2|\mathbf{x}) \\ &= \mathbf{A}_1E(\theta_1|\mathbf{x}) + \mathbf{A}_2E(\theta_2|\mathbf{x}) \\ &= \mathbf{A}_1\hat{\theta}_1 + \mathbf{A}_2\hat{\theta}_2.\end{aligned}$$

The derivation follows that for the scalar state–scalar observation Kalman filter with obvious adjustments for vector quantities. We assume $\mu_s = 0$ so that all random vectors are zero mean. The MMSE estimator of $s[n]$ based on $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n]\}$ is the mean of the posterior PDF

$$\hat{s}[n|n] = E(s[n]|\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n]). \quad (13A.1)$$

But $s[n], \mathbf{x}[0], \dots, \mathbf{x}[n]$ are jointly Gaussian since from (13.13) $s[n]$ depends linearly on $\{s[-1], \mathbf{u}[0], \dots, \mathbf{u}[n]\}$ and $\mathbf{x}[n]$ depends linearly on $s[n], \mathbf{w}[n]$. Hence, we have a linear

dependence on the set of random vectors $S = \{s[-1], \mathbf{u}[0], \dots, \mathbf{u}[n], \mathbf{w}[0], \dots, \mathbf{w}[n]\}$, where each random vector is independent of the others. As a result, the vectors in S are jointly Gaussian and any linear transformation also produces jointly Gaussian random vectors. Now, from (10.24) with zero means we have

$$\hat{s}[n|n] = \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{x} \quad (13A.2)$$

where $\theta = s[n]$ and $\mathbf{x} = [\mathbf{x}^T[0] \mathbf{x}^T[1] \dots \mathbf{x}^T[n]]^T$ and is seen to be linear in $\mathbf{x}$. To determine a recursive in time algorithm we appeal to the vector space approach. Let

$$\mathbf{X}[n] = [\mathbf{x}^T[0] \mathbf{x}^T[1] \dots \mathbf{x}^T[n]]$$

and let $\tilde{\mathbf{x}}[n] = \mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1]$ be the innovation or the error incurred by linearly estimating $\mathbf{x}[n]$ based on $\mathbf{X}[n-1]$. Now we can replace (13A.1) by

$$\hat{s}[n|n] = E(s[n]|\mathbf{X}[n-1], \tilde{\mathbf{x}}[n])$$

since $\mathbf{x}[n]$ is recoverable from $\mathbf{X}[n-1]$ and $\tilde{\mathbf{x}}[n]$ (see comments in Section 13.4). Because $\mathbf{X}[n-1]$ and $\tilde{\mathbf{x}}[n]$ are uncorrelated, we have from property 1

$$\begin{aligned} \hat{s}[n|n] &= E(s[n]|\mathbf{X}[n-1]) + E(s[n]|\tilde{\mathbf{x}}[n]) \\ &= \hat{s}[n|n-1] + E(s[n]|\tilde{\mathbf{x}}[n]). \end{aligned}$$

The predicted sample is found from (13.12) as

$$\begin{aligned} \hat{s}[n|n-1] &= E(\mathbf{A}s[n-1] + \mathbf{B}\mathbf{u}[n]|\mathbf{X}[n-1]) \\ &= \mathbf{A}\hat{s}[n-1|n-1] + \mathbf{B}E(\mathbf{u}[n]|\mathbf{X}[n-1]) \\ &= \mathbf{A}\hat{s}[n-1|n-1] \end{aligned}$$

since the second term is zero. To see why, note that

$$E(\mathbf{u}[n]|\mathbf{X}[n-1]) = E(\mathbf{u}[n]) = \mathbf{0}$$

since from (13.13) $\mathbf{X}[n-1]$ depends on $\{s[-1], \mathbf{u}[0], \dots, \mathbf{u}[n-1], \mathbf{w}[0], \mathbf{w}[1], \dots, \mathbf{w}[n-1]\}$ which are all independent of $\mathbf{u}[n]$. We now have

$$\hat{s}[n|n] = \hat{s}[n|n-1] + E(s[n]|\tilde{\mathbf{x}}[n]) \quad (13A.3)$$

where

$$\hat{s}[n|n-1] = \mathbf{A}\hat{s}[n-1|n-1]. \quad (13A.4)$$

To determine $E(s[n]|\tilde{\mathbf{x}}[n])$ we recall that it is the MMSE estimator of $s[n]$ based on $\tilde{\mathbf{x}}[n]$. Since $s[n]$ and $\tilde{\mathbf{x}}[n]$ are jointly Gaussian, we have from (10.24)

$$E(s[n]|\tilde{\mathbf{x}}[n]) = \mathbf{C}_{s\tilde{x}} \mathbf{C}_{\tilde{x}\tilde{x}}^{-1} \tilde{\mathbf{x}}[n].$$

Letting $\mathbf{K}[n] = \mathbf{C}_{s\tilde{x}} \mathbf{C}_{\tilde{x}\tilde{x}}^{-1}$ be the gain matrix, this is rewritten as

$$E(s[n]|\tilde{\mathbf{x}}[n]) = \mathbf{K}[n](\mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1]).$$

But

$$\mathbf{x}[n] = \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n]$$

so that from property 2

$$\begin{aligned}\hat{\mathbf{x}}[n|n-1] &= \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1] + \hat{\mathbf{w}}[n|n-1] \\ &= \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]\end{aligned}\quad (13A.5)$$

since $\hat{\mathbf{w}}[n|n-1] = \mathbf{0}$ due to the independence of $\mathbf{w}[n]$ with $\{\mathbf{x}[0], \mathbf{x}[1], \dots, \mathbf{x}[n-1]\}$. Thus,

$$E(\mathbf{s}[n]|\hat{\mathbf{x}}[n]) = \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])$$

and from (13A.3) we now have

$$\hat{\mathbf{s}}[n|n] = \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]) \quad (13A.6)$$

where

$$\hat{\mathbf{s}}[n|n-1] = \mathbf{A}\hat{\mathbf{s}}[n-1|n-1].$$

The gain matrix $\mathbf{K}[n]$ is

$$\begin{aligned}\mathbf{K}[n] &= \mathbf{C}_{s\hat{x}}\mathbf{C}_{\hat{x}\hat{x}}^{-1} \\ &= E(\mathbf{s}[n]\hat{\mathbf{x}}^T[n])E(\hat{\mathbf{x}}[n]\hat{\mathbf{x}}^T[n])^{-1}.\end{aligned}$$

To evaluate it we will need the results

1.

$$\begin{aligned}E[\mathbf{s}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] &= \\ E[(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1])(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] &= 0.\end{aligned}$$

2.

$$E[\mathbf{w}[n](\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1])^T] = \mathbf{0}.$$

The first result follows from the innovation $\tilde{\mathbf{x}}[n] = \mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1] = \mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]$ being uncorrelated with the past data and hence with $\hat{\mathbf{s}}[n|n-1]$. The second result is due to the assumption that $\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1]$ is a linear combination of $\{\mathbf{s}[-1], \mathbf{u}[0], \dots, \mathbf{u}[n], \mathbf{w}[0], \dots, \mathbf{w}[n-1]\}$, which is independent of $\mathbf{w}[n]$. We first examine $\mathbf{C}_{s\hat{x}}$ using results 1 and 2 and also (13A.5).

$$\begin{aligned}\mathbf{C}_{s\hat{x}} &= E[\mathbf{s}[n](\mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1])^T] \\ &= E[\mathbf{s}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] \\ &= E[(\mathbf{s}[n] - \hat{\mathbf{s}}[n|n-1])(\mathbf{H}[n]\mathbf{s}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1] + \mathbf{w}[n])^T] \\ &= \mathbf{M}[n|n-1]\mathbf{H}^T[n]\end{aligned}$$

and using result 2 and (13A.5) again

$$\begin{aligned}
 \mathbf{C}_{\hat{x}\hat{x}} &= E [(\mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1])(\mathbf{x}[n] - \hat{\mathbf{x}}[n|n-1])^T] \\
 &= E [(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] \\
 &= E [(\mathbf{H}[n]\mathbf{s}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1] + \mathbf{w}[n]) \\
 &\quad \cdot (\mathbf{H}[n]\mathbf{s}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1] + \mathbf{w}[n])^T] \\
 &= \mathbf{H}[n]\mathbf{M}[n|n-1]\mathbf{H}^T[n] + \mathbf{C}[n].
 \end{aligned}$$

Thus, the Kalman gain matrix is

$$\mathbf{K}[n] = \mathbf{M}[n|n-1]\mathbf{H}^T[n](\mathbf{C}[n] + \mathbf{H}[n]\mathbf{M}[n|n-1]\mathbf{H}^T[n])^{-1}.$$

To evaluate this we require $\mathbf{M}[n|n-1]$, so that by using (13A.4)

$$\begin{aligned}
 \mathbf{M}[n|n-1] &= E [(s[n] - \hat{s}[n|n-1])(s[n] - \hat{s}[n|n-1])^T] \\
 &= E [(\mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] - \mathbf{A}\hat{\mathbf{s}}[n-1|n-1]) \\
 &\quad \cdot (\mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] - \mathbf{A}\hat{\mathbf{s}}[n-1|n-1])^T] \\
 &= E [(\mathbf{A}(s[n-1] - \hat{s}[n-1|n-1]) + \mathbf{B}\mathbf{u}[n]) \\
 &\quad \cdot (\mathbf{A}(s[n-1] - \hat{s}[n-1|n-1]) + \mathbf{B}\mathbf{u}[n])^T]
 \end{aligned}$$

and using the result

$$E [(s[n-1] - \hat{s}[n-1|n-1])\mathbf{u}^T[n]] = 0 \quad (13A.7)$$

we have

$$\mathbf{M}[n|n-1] = \mathbf{A}\mathbf{M}[n-1|n-1]\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T.$$

The equation (13A.7) is true because $s[n-1] - \hat{s}[n-1|n-1]$ depends on $\{s[-1], \mathbf{u}[0], \dots, \mathbf{u}[n-1], \mathbf{w}[0], \dots, \mathbf{w}[n-1]\}$, which are independent of $\mathbf{u}[n]$. Finally, to determine the recursion for $\mathbf{M}[n|n]$ we have from (13A.6)

$$\begin{aligned}
 \mathbf{M}[n|n] &= E [(s[n] - \hat{s}[n|n])(s[n] - \hat{s}[n|n])^T] \\
 &= E [(s[n] - \hat{s}[n|n-1] - \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])) \\
 &\quad \cdot (s[n] - \hat{s}[n|n-1] - \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]))^T] \\
 &= E [(s[n] - \hat{s}[n|n-1])(s[n] - \hat{s}[n|n-1])^T] \\
 &\quad - E [(s[n] - \hat{s}[n|n-1])(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] \mathbf{K}^T[n] \\
 &\quad - \mathbf{K}[n] E [(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])(s[n] - \hat{s}[n|n-1])^T] \\
 &\quad + \mathbf{K}[n] E [(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])(\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1])^T] \mathbf{K}^T[n].
 \end{aligned}$$

The first term is $\mathbf{M}[n|n-1]$, the second expectation is $\mathbf{C}_{s\hat{x}}$ using result 1, the third expectation is $\mathbf{C}_{\hat{s}\hat{x}}^T$, and the last expectation is $\mathbf{C}_{\hat{x}\hat{x}}$. From the derivation for $\mathbf{K}[n]$ we have that

$$\mathbf{K}[n] = \mathbf{C}_{s\hat{x}}\mathbf{C}_{\hat{x}\hat{x}}^{-1}$$

where

$$\mathbf{C}_{s\bar{x}} = \mathbf{M}[n|n-1]\mathbf{H}^T[n].$$

Hence,

$$\begin{aligned}\mathbf{M}[n|n] &= \mathbf{M}[n|n-1] - \mathbf{C}_{s\bar{x}}\mathbf{K}^T[n] - \mathbf{K}[n]\mathbf{C}_{s\bar{x}}^T + \mathbf{K}[n]\mathbf{C}_{\bar{x}\bar{x}}\mathbf{K}^T[n] \\ &= \mathbf{M}[n|n-1] - \mathbf{C}_{s\bar{x}}\mathbf{K}^T[n] - \mathbf{K}[n]\mathbf{C}_{s\bar{x}}^T + \mathbf{C}_{s\bar{x}}\mathbf{C}_{\bar{x}\bar{x}}^{-1}\mathbf{C}_{s\bar{x}}^T \\ &= \mathbf{M}[n|n-1] - \mathbf{K}[n]\mathbf{C}_{s\bar{x}}^T \\ &= (\mathbf{I} - \mathbf{K}[n]\mathbf{H}[n])\mathbf{M}[n|n-1]\end{aligned}$$

which completes the derivation.

Appendix 13B

Extended Kalman Filter Derivation

To derive the extended Kalman filter equations we need to first determine the equations for a modified state model that has a known deterministic input or

$$\mathbf{s}[n] = \mathbf{A}\mathbf{s}[n-1] + \mathbf{B}\mathbf{u}[n] + \mathbf{v}[n] \quad n \geq 0 \quad (13B.1)$$

where $\mathbf{v}[n]$ is known. The presence of $\mathbf{v}[n]$ at the input will produce a deterministic component in the output, so that the mean of $\mathbf{s}[n]$ will no longer be zero. We assume that $E(\mathbf{s}[-1]) = \mathbf{0}$, so that if $\mathbf{v}[n] = \mathbf{0}$, then $E(\mathbf{s}[n]) = \mathbf{0}$. Hence, the effect of the deterministic input is to produce a nonzero mean signal vector

$$\mathbf{s}[n] = \mathbf{s}'[n] + E(\mathbf{s}[n]) \quad (13B.2)$$

where $\mathbf{s}'[n]$ is the value of $\mathbf{s}[n]$ when $\mathbf{v}[n] = \mathbf{0}$. Then, we can write a state equation for the zero mean signal vector $\mathbf{s}'[n] = \mathbf{s}[n] - E(\mathbf{s}[n])$ as

$$\mathbf{s}'[n] = \mathbf{A}\mathbf{s}'[n-1] + \mathbf{B}\mathbf{u}[n].$$

Note that the mean satisfies

$$E(\mathbf{s}[n]) = \mathbf{A}E(\mathbf{s}[n-1]) + \mathbf{v}[n] \quad (13B.3)$$

which follows from (13B.1). Likewise, the observation equation is

$$\begin{aligned} \mathbf{x}[n] &= \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n] \\ &= \mathbf{H}[n]\mathbf{s}'[n] + \mathbf{w}[n] + \mathbf{H}[n]E(\mathbf{s}[n]) \end{aligned}$$

or letting

$$\mathbf{x}'[n] = \mathbf{x}[n] - \mathbf{H}[n]E(\mathbf{s}[n]),$$

we have the usual observation equation

$$\mathbf{x}'[n] = \mathbf{H}[n]\mathbf{s}'[n] + \mathbf{w}[n].$$

The Kalman filter for $\mathbf{s}'[n]$ can be found using (13.58)–(13.62) with $\mathbf{x}[n]$ replaced by $\mathbf{x}'[n]$. Then, the MMSE estimator of $\mathbf{s}[n]$ can easily be found by the relations

$$\begin{aligned}\hat{\mathbf{s}}'[n|n-1] &= \hat{\mathbf{s}}[n|n-1] - E(\mathbf{s}[n]) \\ \hat{\mathbf{s}}'[n-1|n-1] &= \hat{\mathbf{s}}[n-1|n-1] - E(\mathbf{s}[n-1]).\end{aligned}$$

Using these in (13.58), we have for the prediction equation

$$\hat{\mathbf{s}}'[n|n-1] = \mathbf{A}\hat{\mathbf{s}}'[n-1|n-1]$$

or

$$\hat{\mathbf{s}}[n|n-1] - E(\mathbf{s}[n]) = \mathbf{A}(\hat{\mathbf{s}}[n-1|n-1] - E(\mathbf{s}[n-1]))$$

and from (13B.3) this reduces to

$$\hat{\mathbf{s}}[n|n-1] = \mathbf{A}\hat{\mathbf{s}}[n-1|n-1] + \mathbf{v}[n].$$

For the correction we have from (13.61)

$$\hat{\mathbf{s}}'[n|n] = \hat{\mathbf{s}}'[n|n-1] + \mathbf{K}[n](\mathbf{x}'[n] - \mathbf{H}[n]\hat{\mathbf{s}}'[n|n-1])$$

or

$$\begin{aligned}\hat{\mathbf{s}}[n|n] - E(\mathbf{s}[n]) &= \hat{\mathbf{s}}[n|n-1] - E(\mathbf{s}[n]) + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]E(\mathbf{s}[n]) - \mathbf{H}[n](\hat{\mathbf{s}}[n|n-1] - E(\mathbf{s}[n])))\end{aligned}$$

This reduces to the usual equation

$$\hat{\mathbf{s}}[n|n] = \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]).$$

The Kalman gain as well as the MSE matrices remain the same since the MMSE estimator will not incur any additional error due to known constants. Hence, the only revision is in the prediction equation.

Returning to the extended Kalman filter, we also have the modified observation equation

$$\mathbf{x}[n] = \mathbf{H}[n]\mathbf{s}[n] + \mathbf{w}[n] + \mathbf{z}[n]$$

where $\mathbf{z}[n]$ is known. With $\mathbf{x}'[n] = \mathbf{x}[n] - \mathbf{z}[n]$ we have the usual observation equation. Hence, our two equations become, upon replacing $\mathbf{A}$ by $\mathbf{A}[n]$,

$$\begin{aligned}\hat{\mathbf{s}}[n|n-1] &= \mathbf{A}[n-1]\hat{\mathbf{s}}[n-1|n-1] + \mathbf{v}[n] \\ \hat{\mathbf{s}}[n|n] &= \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{z}[n] - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]).\end{aligned}$$

But

$$\begin{aligned}\mathbf{v}[n] &= \mathbf{a}(\hat{\mathbf{s}}[n-1|n-1]) - \mathbf{A}[n-1]\hat{\mathbf{s}}[n-1|n-1] \\ \mathbf{z}[n] &= \mathbf{h}(\hat{\mathbf{s}}[n|n-1]) - \mathbf{H}[n]\hat{\mathbf{s}}[n|n-1]\end{aligned}$$

so that finally

$$\begin{aligned}\hat{\mathbf{s}}[n|n-1] &= \mathbf{a}(\hat{\mathbf{s}}[n-1|n-1]) \\ \hat{\mathbf{s}}[n|n] &= \hat{\mathbf{s}}[n|n-1] + \mathbf{K}[n](\mathbf{x}[n] - \mathbf{h}(\hat{\mathbf{s}}[n|n-1])).\end{aligned}$$

Chapter 14

Summary of Estimators

14.1 Introduction

The choice of an estimator that will perform well for a particular application depends upon many considerations. Of primary concern is the selection of a good data model. It should be complex enough to describe the principal features of the data, but at the same time simple enough to allow an estimator that is optimal and easily implemented. We have seen that at times we were unable to determine the existence of an optimal estimator, an example being the search for the MVU estimator in classical estimation. In other instances, even though the optimal estimator could easily be found, it could not be implemented, an example being the MMSE estimator in Bayesian estimation. For a particular problem we are neither assured of finding an optimal estimator or, even if we are fortunate enough to do so, of being able to implement it. Therefore, it becomes critical to have at one's disposal a knowledge of the estimators that are optimal and easily implemented, and furthermore, to understand under what conditions we may justify their use. To this end we now summarize the approaches, assumptions, and for the linear data model, the explicit estimators obtained. Then, we will illustrate the decision making process that one must go through in order to choose a good estimator. Also, in our discussions we will highlight some relationships between the various estimators.

14.2 Estimation Approaches

We first summarize the classical approaches to estimation in which the unknown $p \times 1$ parameter vector θ is assumed to be a deterministic constant, followed by the Bayesian approaches in which θ is assumed to be the realization of a random vector. In the classical approach the data information is summarized by the probability density function (PDF) $p(\mathbf{x}; \theta)$, where the PDF is functionally dependent on θ . In contrast to this modeling, the Bayesian approach augments the data information with a prior PDF $p(\theta)$ which describes our knowledge about θ (before any data are observed). This is

summarized by the joint PDF $p(\mathbf{x}, \boldsymbol{\theta})$ or, equivalently, by the conditional PDF $p(\mathbf{x}|\boldsymbol{\theta})$ (data information) and the prior PDF $p(\boldsymbol{\theta})$ (prior information).

Classical Estimation Approaches

1. Cramer-Rao Lower Bound (CRLB)

a. *Data Model/Assumptions*

PDF $p(\mathbf{x}; \boldsymbol{\theta})$ is known.

b. *Estimator*

If the equality condition for the CRLB

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{I}(\boldsymbol{\theta})(\mathbf{g}(\mathbf{x}) - \boldsymbol{\theta})$$

is satisfied, then the estimator is

$$\hat{\boldsymbol{\theta}} = \mathbf{g}(\mathbf{x})$$

where $\mathbf{I}(\boldsymbol{\theta})$ is a $p \times p$ matrix dependent only on $\boldsymbol{\theta}$ and $\mathbf{g}(\mathbf{x})$ is a p -dimensional function of the data $\mathbf{x}$.

c. *Optimality/Error Criterion*

$\hat{\boldsymbol{\theta}}$ achieves the CRLB, the lower bound on the variance for any unbiased estimator (and hence is said to be efficient), and is therefore the minimum variance unbiased (MVU) estimator. The MVU estimator is the one whose variance for each component is minimum among all unbiased estimators.

d. *Performance*

It is unbiased or

$$E(\hat{\theta}_i) = \theta_i \quad i = 1, 2, \dots, p$$

and has minimum variance

$$\text{var}(\hat{\theta}_i) = [\mathbf{I}^{-1}(\boldsymbol{\theta})]_{ii} \quad i = 1, 2, \dots, p$$

where

$$[\mathbf{I}(\boldsymbol{\theta})]_{ij} = E \left[\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \theta_j} \right].$$

e. *Comments*

An efficient estimator may not exist, and hence this approach may fail.

f. *Reference*

Chapter 3

2. Rao-Blackwell-Lehmann-Scheffe

a. *Data Model/Assumptions*

PDF $p(\mathbf{x}; \boldsymbol{\theta})$ is known.

b. Estimator

- i. Find a sufficient statistic $\mathbf{T}(\mathbf{x})$ by factoring PDF as

$$p(\mathbf{x}; \boldsymbol{\theta}) = g(\mathbf{T}(\mathbf{x}), \boldsymbol{\theta})h(\mathbf{x})$$

where $\mathbf{T}(\mathbf{x})$ is a p -dimensional function of $\mathbf{x}$, g is a function depending only on $\mathbf{T}$ and $\boldsymbol{\theta}$, and h depends only on $\mathbf{x}$.

- ii. If $E[\mathbf{T}(\mathbf{x})] = \boldsymbol{\theta}$, then $\hat{\boldsymbol{\theta}} = \mathbf{T}(\mathbf{x})$. If not, we must find a p -dimensional function $\mathbf{g}$ so that $E[\mathbf{g}(\mathbf{T})] = \boldsymbol{\theta}$, and then $\hat{\boldsymbol{\theta}} = \mathbf{g}(\mathbf{T})$.

c. Optimality/Error Criterion

$\hat{\boldsymbol{\theta}}$ is the MVU estimator.

d. Performance

$\hat{\theta}_i$ for $i = 1, 2, \dots, p$ is unbiased. The variance depends on the PDF—no general formula is available.

e. Comments

Also, “completeness” of sufficient statistic must be checked. A p -dimensional sufficient statistic may not exist, so that this method may fail.

f. Reference

Chapter 5

3. Best Linear Unbiased Estimator (BLUE)**a. Data Model/Assumptions**

$$\mathbf{E}(\mathbf{x}) = \mathbf{H}\boldsymbol{\theta}$$

where $\mathbf{H}$ is an $N \times p$ ($N > p$) known matrix and $\mathbf{C}$, the covariance matrix of $\mathbf{x}$, is known. Equivalently, we have

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where $E(\mathbf{w}) = \mathbf{0}$ and $\mathbf{C}_w = \mathbf{C}$.

b. Estimator

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}.$$

c. Optimality/Error Criterion

$\hat{\theta}_i$ for $i = 1, 2, \dots, p$ has the minimum variance of all unbiased estimators that are linear in $\mathbf{x}$.

d. Performance

$\hat{\theta}_i$ for $i = 1, 2, \dots, p$ is unbiased. The variance is

$$\text{var}(\hat{\theta}_i) = [(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}]_{ii} \quad i = 1, 2, \dots, p.$$

e. *Comments*

If $\mathbf{w}$ is a Gaussian random vector so that $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, then $\hat{\boldsymbol{\theta}}$ is also the MVU estimator (for all linear or nonlinear functions of $\mathbf{x}$).

f. *Reference*

Chapter 6

4. Maximum Likelihood Estimator (MLE)a. *Data Model/Assumptions*

PDF $p(\mathbf{x}; \boldsymbol{\theta})$ is known.

b. *Estimator*

$\hat{\boldsymbol{\theta}}$ is the value of $\boldsymbol{\theta}$ maximizing $p(\mathbf{x}; \boldsymbol{\theta})$, where $\mathbf{x}$ is replaced by the observed data samples.

c. *Optimality/Error Criterion*

Not optimal in general. Under certain conditions on the PDF, however, the MLE is efficient for large data records or as $N \rightarrow \infty$ (asymptotically). Hence, asymptotically it is the MVU estimator.

d. *Performance*

For finite N depends on PDF—no general formula is available. Asymptotically, under certain conditions

$$\hat{\boldsymbol{\theta}} \stackrel{a}{\sim} \mathcal{N}(\boldsymbol{\theta}, I^{-1}(\boldsymbol{\theta})).$$

e. *Comments*

If an MVU estimator exists, the maximum likelihood procedure will produce it.

f. *Reference*

Chapter 7

5. Least Squares Estimator (LSE)a. *Data Model/Assumptions*

$$x[n] = s[n; \boldsymbol{\theta}] + w[n] \quad n = 0, 1, \dots, N-1$$

where the signal $s[n; \boldsymbol{\theta}]$ depends explicitly on the unknown parameters. Equivalently, the model is

$$\mathbf{x} = \mathbf{s}(\boldsymbol{\theta}) + \mathbf{w}$$

where $\mathbf{s}$ is a known N -dimensional function of $\boldsymbol{\theta}$ and the noise or perturbation $\mathbf{w}$ has zero mean.

b. *Estimator*

$\hat{\boldsymbol{\theta}}$ is the value of $\boldsymbol{\theta}$ that minimizes

$$\begin{aligned} J(\boldsymbol{\theta}) &= (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta}))^T (\mathbf{x} - \mathbf{s}(\boldsymbol{\theta})) \\ &= \sum_{n=0}^{N-1} (x[n] - s[n; \boldsymbol{\theta}])^2. \end{aligned}$$

c. *Optimality/Error Criterion*

None in general.

d. *Performance*Depends on PDF of $\mathbf{w}$ —no general formula is available.e. *Comments*

The fact that we are minimizing a LS error criterion does not in general translate into minimizing the estimation error. Also, if $\mathbf{w}$ is a Gaussian random vector with $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, then the LSE is equivalent to the MLE.

f. *Reference*

Chapter 8

6. Method of Momentsa. *Data Model/Assumptions*

There are p moments $\mu_i = E(x^i[n])$ for $i = 1, 2, \dots, p$, which depend on θ in a known way. The entire PDF need not be known.

b. *Estimator*

If $\boldsymbol{\mu} = \mathbf{h}(\boldsymbol{\theta})$, where $\mathbf{h}$ is an invertible p -dimensional function of $\boldsymbol{\theta}$ and $\boldsymbol{\mu} = [\mu_1 \mu_2 \dots \mu_p]^T$, then

$$\hat{\boldsymbol{\theta}} = \mathbf{h}^{-1}(\hat{\boldsymbol{\mu}})$$

where

$$\hat{\boldsymbol{\mu}} = \begin{bmatrix} \frac{1}{N} \sum_{n=0}^{N-1} x[n] \\ \frac{1}{N} \sum_{n=0}^{N-1} x^2[n] \\ \vdots \\ \frac{1}{N} \sum_{n=0}^{N-1} x^p[n] \end{bmatrix}.$$

c. *Optimality/Error Criterion*

None in general.

d. *Performance*

For finite N it depends on the PDF of $\mathbf{x}$. However, for large data records (asymptotically), if $\hat{\theta}_i = g_i(\hat{\boldsymbol{\mu}})$, then

$$\begin{aligned} E(\hat{\theta}_i) &= g_i(\boldsymbol{\mu}) \\ \text{var}(\hat{\theta}_i) &= \left. \frac{\partial g_i}{\partial \hat{\boldsymbol{\mu}}} \right|_{\hat{\boldsymbol{\mu}}=\boldsymbol{\mu}}^T \mathbf{C}_{\hat{\boldsymbol{\mu}}} \left. \frac{\partial g_i}{\partial \hat{\boldsymbol{\mu}}} \right|_{\hat{\boldsymbol{\mu}}=\boldsymbol{\mu}} \end{aligned}$$

for $i = 1, 2, \dots, p$.e. *Comments*

Usually very easy to implement.

f. *Reference*

Chapter 9

Bayesian Estimation Approaches

7. Minimum Mean Square Error (MMSE) Estimator

a. *Data Model/Assumptions*

The joint PDF of $\mathbf{x}, \boldsymbol{\theta}$ or $p(\mathbf{x}, \boldsymbol{\theta})$ is known, where $\boldsymbol{\theta}$ is now considered to be a random vector. Usually, $p(\mathbf{x}|\boldsymbol{\theta})$ is specified as the data model and $p(\boldsymbol{\theta})$ as the prior PDF for $\boldsymbol{\theta}$, so that $p(\mathbf{x}, \boldsymbol{\theta}) = p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta})$.

b. *Estimator*

$$\hat{\boldsymbol{\theta}} = E(\boldsymbol{\theta}|\mathbf{x})$$

where the expectation is with respect to the posterior PDF

$$p(\boldsymbol{\theta}|\mathbf{x}) = \frac{p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}) d\boldsymbol{\theta}}$$

If $\mathbf{x}, \boldsymbol{\theta}$ are jointly Gaussian,

$$\hat{\boldsymbol{\theta}} = E(\boldsymbol{\theta}) + \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} (\mathbf{x} - E(\mathbf{x})). \quad (14.1)$$

c. *Optimality/Error Criterion*

$\hat{\boldsymbol{\theta}}_i$ minimizes the Bayesian MSE

$$\text{Bmse}(\hat{\boldsymbol{\theta}}_i) = E[(\theta_i - \hat{\theta}_i)^2] \quad i = 1, 2, \dots, p \quad (14.2)$$

where the expectation is with respect to $p(\mathbf{x}, \boldsymbol{\theta}_i)$.

d. *Performance*

The error $\epsilon_i = \theta_i - \hat{\theta}_i$ has zero mean and variance

$$\text{var}(\epsilon_i) = \text{Bmse}(\hat{\theta}_i) = \int [\mathbf{C}_{\theta|x}]_{ii} p(\mathbf{x}) d\mathbf{x} \quad (14.3)$$

where $\mathbf{C}_{\theta|x}$ is the covariance matrix of $\boldsymbol{\theta}$ conditioned on $\mathbf{x}$, or of the posterior PDF $p(\boldsymbol{\theta}|\mathbf{x})$. If $\mathbf{x}, \boldsymbol{\theta}$ are jointly Gaussian, then the error is Gaussian with zero mean and variance

$$\text{var}(\epsilon_i) = \text{Bmse}(\hat{\theta}_i) = [\mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta x} \mathbf{C}_{xx}^{-1} \mathbf{C}_{x\theta}]_{ii}.$$

e. *Comments*

In the non-Gaussian case, this will be difficult to implement.

f. *Reference*

Chapters 10 and 11

8. Maximum A Posteriori (MAP) Estimator**a. Data Model/Assumptions**

Same as for the MMSE estimator.

b. Estimator

$\hat{\theta}$ is the value of θ that maximizes $p(\theta|\mathbf{x})$ or, equivalently, the value that maximizes $p(\mathbf{x}|\theta)p(\theta)$. If $\mathbf{x}, \theta$ are jointly Gaussian, then $\hat{\theta}$ is given by (14.1).

c. Optimality/Error Criterion

Minimizes the “hit-or-miss” cost function.

d. Performance

Depends on PDF—no general formula is available. If $\mathbf{x}, \theta$ are jointly Gaussian, then the performance is identical to that of the MMSE estimator.

e. Comments

For PDFs whose mean and mode (the location of the maximum) are the same, the MMSE and MAP estimators will be identical, i.e., the Gaussian PDF, for example.

f. Reference: Chapter 11**9. Linear Minimum Mean Square Error (LMMSE) Estimator****a. Data Model/Assumptions**

The first two moments of the joint PDF $p(\mathbf{x}, \theta)$ are known or the mean and covariance

$$\begin{bmatrix} E(\theta) \\ E(\mathbf{x}) \end{bmatrix} \quad \begin{bmatrix} \mathbf{C}_{\theta\theta} & \mathbf{C}_{\theta\mathbf{x}} \\ \mathbf{C}_{\mathbf{x}\theta} & \mathbf{C}_{\mathbf{x}\mathbf{x}} \end{bmatrix}.$$

b. Estimator

$$\hat{\theta} = E(\theta) + \mathbf{C}_{\theta\mathbf{x}} \mathbf{C}_{\mathbf{x}\mathbf{x}}^{-1} (\mathbf{x} - E(\mathbf{x})).$$

c. Optimality/Error Criterion

$\hat{\theta}_i$ has the minimum Bayesian MSE (see (14.2)) of all estimators that are linear functions of $\mathbf{x}$.

d. Performance

The error $\epsilon_i = \theta_i - \hat{\theta}_i$ has zero mean and variance

$$\text{var}(\epsilon_i) = \text{Bmse}(\hat{\theta}_i) = [\mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta\mathbf{x}} \mathbf{C}_{\mathbf{x}\mathbf{x}}^{-1} \mathbf{C}_{\mathbf{x}\theta}]_{ii}.$$

e. Comments

If $\mathbf{x}, \theta$ are jointly Gaussian, this is identical to the MMSE and MAP estimators.

f. Reference: Chapter 12

The reader may observe that we have omitted a summary of the Kalman filter. This is because it is a particular implementation of the MMSE estimator and so is contained within that discussion.

14.3 Linear Model

When the linear model can be used to describe the data, the various estimation approaches yield closed form estimators. In fact, by assuming the linear model we are able to determine the optimal estimator as well as its performance for both the classical and Bayesian approaches. We first consider the classical approach. The classical general linear model assumes the data to be described by

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \mathbf{w}$$

where $\mathbf{x}$ is an $N \times 1$ vector of observations, $\mathbf{H}$ is a known $N \times p$ observation matrix ($N > p$) of rank p , $\boldsymbol{\theta}$ is a $p \times 1$ vector of parameters to be estimated, and $\mathbf{w}$ is an $N \times 1$ noise vector with PDF $\mathcal{N}(\mathbf{0}, \mathbf{C})$. The PDF of $\mathbf{x}$ is

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\boldsymbol{\theta}) \right]. \quad (14.4)$$

1. Cramer-Rao Lower Bound

$$\frac{\partial \ln p(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$$

where

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x} \quad (14.5)$$

so that $\hat{\boldsymbol{\theta}}$ is the MVU estimator (and also efficient) and has minimum variance given by the diagonal elements of the covariance matrix

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \mathbf{I}^{-1}(\boldsymbol{\theta}) = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1}.$$

2. Rao-Blackwell-Lehmann-Scheffe

A factorization of the PDF will yield

$$p(\mathbf{x}; \boldsymbol{\theta}) = \underbrace{\frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C})} \exp \left\{ -\frac{1}{2} [(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T \mathbf{H}^T \mathbf{C}^{-1} \mathbf{H} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})] \right\}}_{g(\mathbf{T}(\mathbf{x}), \boldsymbol{\theta})} \cdot \underbrace{\exp \left\{ -\frac{1}{2} [(\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\hat{\boldsymbol{\theta}})] \right\}}_{h(\mathbf{x})}$$

where

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}.$$

The sufficient statistic is $\mathbf{T}(\mathbf{x}) = \hat{\boldsymbol{\theta}}$, which can be shown to be unbiased as well as complete. Thus, it is the MVU estimator.

3. Best Linear Unbiased Estimator

Now we will have the identical estimator as in the previous two cases since $\hat{\theta}$ is already linear in $\mathbf{x}$. However, if $\mathbf{w}$ were not Gaussian, $\hat{\theta}$ would still be the BLUE but not the MVU estimator. Note that the data modeling assumption for the BLUE is satisfied by the general linear model.

4. Maximum Likelihood Estimator

To find the MLE we maximize $p(\mathbf{x}; \theta)$ as given in (14.4) or, equivalently, we minimize

$$(\mathbf{x} - \mathbf{H}\theta)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{H}\theta).$$

This leads to

$$\hat{\theta} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$$

which we know to be the MVU estimator. Thus, as expected, since an efficient estimator exists (satisfies the CRLB), then the maximum likelihood procedure produces it.

5. Least Squares Estimator

Viewing $\mathbf{H}\theta$ as the signal vector $\mathbf{s}(\theta)$, we must minimize

$$\begin{aligned} J(\theta) &= (\mathbf{x} - \mathbf{s}(\theta))^T (\mathbf{x} - \mathbf{s}(\theta)) \\ &= (\mathbf{x} - \mathbf{H}\theta)^T (\mathbf{x} - \mathbf{H}\theta) \end{aligned}$$

which is identical to the maximum likelihood procedure if $\mathbf{C} = \sigma^2 \mathbf{I}$. The LSE is $\hat{\theta} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{x}$, which is also the MVU estimator if $\mathbf{C} = \sigma^2 \mathbf{I}$. If $\mathbf{C} \neq \sigma^2 \mathbf{I}$, then $\hat{\theta}$ will not be the MVU estimator. However, if we minimize the *weighted* LS criterion

$$J'(\theta) = (\mathbf{x} - \mathbf{H}\theta)^T \mathbf{W} (\mathbf{x} - \mathbf{H}\theta)$$

where the weighting matrix is $\mathbf{C}^{-1}$, then the resultant estimator will be the MVU estimator of (14.5). Finally, if $\mathbf{w}$ is not Gaussian, the weighted LSE would still be $\hat{\theta}$ as given by (14.5), but it would only be the BLUE.

6. Method of Moments

This approach is omitted since the optimal estimator is known.

The results are summarized in Table 14.1.

Proceeding to the Bayesian linear model, we assume that

$$\mathbf{x} = \mathbf{H}\theta + \mathbf{w}$$

where $\mathbf{x}$ is an $N \times 1$ vector of observations, $\mathbf{H}$ is a known $N \times p$ observation matrix (with $N \leq p$ possibly), θ is a $p \times 1$ random vector with PDF $\mathcal{N}(\mu_\theta, \mathbf{C}_\theta)$, and $\mathbf{w}$ is an $N \times 1$ noise vector independent of θ and having PDF $\mathcal{N}(0, \mathbf{C}_w)$. The conditional PDF of $\mathbf{x}$ is

$$p(\mathbf{x}|\theta) = \frac{1}{(2\pi)^{\frac{N}{2}} \det^{\frac{1}{2}}(\mathbf{C}_w)} \exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{H}\theta)^T \mathbf{C}_w^{-1} (\mathbf{x} - \mathbf{H}\theta) \right]$$

TABLE 14.1 Properties of $\hat{\theta}$ for Classical General Linear Model

Model: $\mathbf{x} = \mathbf{H}\theta + \mathbf{w}$		
Assumptions: $E(\mathbf{w}) = 0$		
$\mathbf{C}_w = \mathbf{C}$		
Estimator: $\hat{\theta} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{x}$		
Properties	$\mathbf{w} \sim \text{Gaussian}$ (linear model)	$\mathbf{w} \sim \text{Non-Gaussian}$
Efficient	*	
Sufficient Statistic	*	
MVU	*	
BLUE	*	*
MLE	*	
WLS ($\mathbf{W} = \mathbf{C}^{-1}$)	*	*

* Property holds.

and the prior PDF of θ is

$$p(\theta) = \frac{1}{(2\pi)^{\frac{p}{2}} \det^{\frac{1}{2}}(\mathbf{C}_\theta)} \exp \left[-\frac{1}{2} (\theta - \mu_\theta)^T \mathbf{C}_\theta^{-1} (\theta - \mu_\theta) \right].$$

The posterior PDF $p(\theta|\mathbf{x})$ is again Gaussian with mean and covariance

$$E(\theta|\mathbf{x}) = \mu_\theta + \mathbf{C}_\theta \mathbf{H}^T (\mathbf{H} \mathbf{C}_\theta \mathbf{H}^T + \mathbf{C}_w)^{-1} (\mathbf{x} - \mathbf{H} \mu_\theta) \quad (14.6)$$

$$= \mu_\theta + (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} (\mathbf{x} - \mathbf{H} \mu_\theta) \quad (14.7)$$

$$\mathbf{C}_{\theta|\mathbf{x}} = \mathbf{C}_\theta - \mathbf{C}_\theta \mathbf{H}^T (\mathbf{H} \mathbf{C}_\theta \mathbf{H}^T + \mathbf{C}_w)^{-1} \mathbf{H} \mathbf{C}_\theta \quad (14.8)$$

$$= (\mathbf{C}_\theta^{-1} + \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1}. \quad (14.9)$$

7. Minimum Mean Square Error Estimator

The MMSE estimator is just the mean of the posterior PDF given by (14.6) or (14.7). The minimum Bayesian MSE or $E((\theta_i - \hat{\theta}_i)^2)$, is from (14.3),

$$\text{Bmse}(\hat{\theta}_i) = [\mathbf{C}_{\theta|\mathbf{x}}]_{ii}$$

since $\mathbf{C}_{\theta|\mathbf{x}}$ does not depend on $\mathbf{x}$ (see (14.8)).

8. Maximum A Posteriori Estimator

Because the location of the peak (or mode) of the Gaussian PDF is equal to the mean, the MAP estimator is identical to the MMSE estimator.

9. Linear Minimum Mean Square Error Estimator

Since the MMSE estimator is linear in $\mathbf{x}$, the LMMSE estimator is just given by (14.6) or (14.7).

Hence, for the Bayesian linear model the MMSE estimator, MAP estimator, and the LMMSE estimator are identical. A last comment concerns the form of the estimator when there is no prior information. This may be modeled by letting $\mathbf{C}_\theta^{-1} \rightarrow \mathbf{0}$. Then, from (14.7) we have

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}_w^{-1} \mathbf{x}$$


which is recognized as having the identical form as the MVU estimator for the classical general linear model. Of course, the estimators cannot really be compared since they have been derived under different data modeling assumptions (see Problem 11.7). However, this apparent equivalence has often been identified by asserting that the Bayesian approach with no prior information is equivalent to the classical approach. *When viewed in its proper statistical context, this assertion is incorrect.*

14.4 Choosing an Estimator

We now illustrate the decision making process involved in choosing an estimator. In doing so our goal is always to find the optimal estimator for a given data model. If this is not possible, we consider suboptimal estimation approaches. We will consider our old friend the data set

$$\mathbf{x}[n] = A[n] + w[n] \quad n = 0, 1, \dots, N-1$$

where the unknown parameters are $\{A[0], A[1], \dots, A[N-1]\}$. We have allowed the parameter A to change with time, as most parameters normally change to some extent in real world problems. Depending on our assumptions on $A[n]$ and $w[n]$, the data may have the form of the classical or Bayesian linear model. If this is the case, the optimal estimator is easily found as explained previously. However, even if the estimator is optimal for the assumed data model, its performance may not be adequate. Thus, the data model may need to be modified, as we now discuss. We will refer to the flowchart in Figure 14.1 as we describe the considerations in the selection of an estimator. Since we are attempting to estimate as many parameters as data points, we can expect poor estimation performance due to a lack of averaging (see Figure 14.1a). With prior knowledge such as the PDF of $\boldsymbol{\theta} = [A[0], A[1], \dots, A[N-1]]^T$, we could use a Bayesian approach as detailed in Figure 14.1b. Based on the PDF $p(\mathbf{x}, \boldsymbol{\theta})$ we can in theory find the MMSE estimator. This will involve a multidimensional integration and may not in practice be possible. Failing to determine the MMSE estimator we could attempt to maximize the posterior PDF to produce the MAP estimator, either analytically or at least numerically. As a last resort, if the first two joint moments of $\mathbf{x}$ and $\boldsymbol{\theta}$ are available, we could determine the LMMSE estimator in explicit form. Even if dimensionality is not a problem, as for example if $A[n] = A$, the use of prior knowledge as embodied by the prior PDF will improve the estimation accuracy in the Bayesian sense. That is to say, the Bayesian MSE will be reduced. If no prior knowledge is available, we will be forced to reevaluate our data model or else obtain more data. For example, we might suppose that $A[n] = A$ or even $A[n] = A + Bn$, reducing the dimensionality of the problem. This may result in bias errors due to modeling inaccuracies, but at least




(a) Classical versus Bayesian


Figure 14.1 Decision-making process in estimator selection

the variability of any resultant estimator would be reduced. Then, we could resort to a classical approach (see Figure 14.1c). If the PDF is known, we first compute the equality condition for the CRLB, and if satisfied, an efficient and hence MVU estimator will be found. If not, we could attempt to find a sufficient statistic, make it unbiased, and if complete, this would produce the MVU estimator. If these approaches fail, a maximum likelihood approach could be tried if the likelihood function (PDF with $\mathbf{x}$ replaced by the observed data values) can be maximized analytically or at least numerically. Finally, the moments could be found and a method of moments estimator tried. Note that the entire PDF need not be known for a method of moments estimator. If the PDF is unknown but the problem is one of a signal in noise, then either a BLUE or LS approach could be tried. If the signal is linear in θ and the first two moments of the noise are known, a BLUE can be found. Otherwise, a LS and a possibly nonlinear LS estimator must be employed.

In general, the choice of an appropriate estimator for a signal processing problem should begin with the search for an optimal estimator that is computationally feasible. If the search proves to be futile, then suboptimal estimators should be investigated.



(b) Bayesian approach



(c) Classical approach

Chapter 15

Extensions for Complex Data and Parameters

15.1 Introduction

For many signal processing applications the data samples are more conveniently modeled as being complex—typically the concatenation of two time series of real data into a single time series of complex data. Also, the same technique can be useful in representing a real parameter vector of dimension 2×1 by a single complex parameter. In doing so it is found that the representation is considerably more intuitive, is analytically more tractable, and lends itself to easier manipulation by a digital computer. An analogous situation arises in the more convenient use of a complex Fourier series composed of complex exponentials for a real data signal as opposed to a real Fourier series of sines and cosines. Furthermore, once the complex Fourier series representation has been accepted, its extension to complex data is trivial, requiring only the Hermitian symmetry property of the Fourier coefficients to be relaxed. In this chapter we reformulate much of our previous theory to accommodate complex data and complex parameters. In doing so we do not present any new theory but only an algebra for manipulating complex data and parameters.

15.2 Summary

The need for complex data models with complex parameters is discussed in Section 15.3. A particularly important model is the complex envelope representation of a real bandpass signal as given by (15.2). Next, the tedious process of minimizing a function with respect to the real and imaginary parts of a complex parameter is illustrated in Example 15.2. It is shown how the introduction of a complex derivative greatly simplifies the algebra. Complex random variables are defined and their properties described in Section 15.4. The important complex Gaussian PDF for a complex random variable is given by (15.16), while for a complex random vector the corresponding PDF is given by (15.22). These PDFs assume that the real covariance matrix has the special form




Figure 15.1 Definition of complex envelope

of (15.19) or, equivalently, the covariances satisfy (15.20) for the case of two random variables. Properties of complex Gaussian random variables are summarized in that section as well. If a complex WSS random process has an autocorrelation and a cross-correlation that satisfy (15.33), then it is said to be a complex Gaussian WSS random process. An example is the complex envelope of a real bandpass Gaussian random process, as described in Example 15.5. The complex derivative of a real function with respect to a complex variable is formally defined in (15.40). The associated complex gradient can be used to minimize Hermitian functions by employing (15.44)–(15.46). If the Hermitian function to be minimized has a linear constraint on its parameters, then the solution is given by (15.51). Classical estimation based on complex Gaussian data with real parameters employs the CRLB of (15.52) and the equality condition of (15.53). When the Fisher information matrix has a special form, however, the equality condition of (15.54) can be used. This involves a complex Fisher information matrix. One important example is the complex linear model in Example 15.9, in which (15.58) is the efficient estimator and (15.59) is its corresponding covariance. Bayesian estimation for the complex Bayesian linear model is discussed in Section 15.8. The MMSE estimator is given by (15.64) or (15.65), while the minimum Bayesian MSE is given by (15.66) or (15.67). For large data records an approximate PDF for a complex Gaussian WSS random process is (15.68) and an approximate CRLB can be computed based on (15.69). These forms are frequently easier to use in deriving estimators.

15.3 Complex Data and Parameters

The most common example of the use of complex data in signal processing occurs in radar/sonar in which a bandpass signal is of interest. As shown in Figure 15.1a, if the Fourier transform $S(F)$ of a continuous-time real signal $s(t)$ is nonzero over a band $(F_0 - B/2, F_0 + B/2)$, then the essential information is contained in the *complex envelope* $\tilde{s}(t)$ whose Fourier transform $\tilde{S}(F)$ is shown in Figure 15.1b. (We will generally use a tilde to denote a complex quantity when necessary to avoid confusion. Otherwise, whether the quantity is real or complex will be evident from the context of the discussion.) The




Figure 15.2 Extraction of discrete-time complex envelope from bandpass signal

relationship between the Fourier transforms of the real bandpass signal and its complex envelope is

$$S(F) = \tilde{S}(F - F_0) + \tilde{S}^*(-(F + F_0)). \quad (15.1)$$

To obtain $S(F)$ we shift the complex envelope spectrum up to $F = F_0$ and also, after first “flipping it around” in F and conjugating it, down to $F = -F_0$. Taking the inverse Fourier transform of (15.1) produces

$$s(t) = \tilde{s}(t) \exp(j2\pi F_0 t) + [\tilde{s}(t) \exp(j2\pi F_0 t)]^*$$

or

$$s(t) = 2\text{Re} [\tilde{s}(t) \exp(j2\pi F_0 t)]. \quad (15.2)$$

Alternatively, if we let $\tilde{s}(t) = s_R(t) + js_I(t)$, where R and I refer to the real and imaginary parts of the complex envelope, we have

$$s(t) = 2s_R(t) \cos 2\pi F_0 t - 2s_I(t) \sin 2\pi F_0 t. \quad (15.3)$$

This is also referred to as the “narrowband” representation, although it is actually valid if $F_0 > B/2$. Note that $s_R(t), s_I(t)$ are bandlimited to $B/2$ Hz, as is evident from Figure 15.1b. Hence, the bandpass signal is completely described by its complex envelope. In processing $s(t)$ or a noise corrupted version we need only obtain the complex envelope. From a practical viewpoint the use of the complex envelope, which is a low-pass signal, allows us to sample at the lower rate of B complex samples/sec as opposed to the higher rate of $2(F_0 + B/2)$ real samples/sec for the real bandpass signal. This results in the widely adopted means of processing a real bandpass signal shown in Figure 15.2. The aim is to produce the samples of the complex envelope in a digital computer. Note from (15.3) that the output of the low-pass filter for the

“in-phase” channel is

$$\begin{aligned} & [2s_R(t) \cos^2 2\pi F_0 t - 2s_I(t) \sin 2\pi F_0 t \cos 2\pi F_0 t]_{\text{LPF}} \\ &= [s_R(t) + s_R(t) \cos 4\pi F_0 t - s_I(t) \sin 4\pi F_0 t]_{\text{LPF}} \\ &= s_R(t) \end{aligned}$$

since the other signals have spectra centered about $\pm 2F_0$. The “LPF” designation means the output of the low-pass filter in Figure 15.2. Similarly, the output of the low-pass filter in the “quadrature” channel is $s_I(t)$. Hence, the discrete signal at the output of the processor in Figure 15.2 is $s_R(n\Delta) + js_I(n\Delta) = \tilde{s}(n\Delta)$. It is seen that complex data naturally arise in bandpass systems. An important example follows.

Example 15.1 - Complex Envelope for Sinusoids

It frequently occurs that the bandpass signal of interest is sinusoidal, as for example,

$$s(t) = \sum_{i=1}^p A_i \cos(2\pi F_i t + \phi_i)$$

where it is known that $F_0 - B/2 \leq F_i \leq F_0 + B/2$ for all i . The complex envelope for this signal is easily found by noting that it may be written as

$$\begin{aligned} s(t) &= \text{Re} \left[\sum_{i=1}^p A_i \exp[j(2\pi F_i t + \phi_i)] \right] \\ &= 2 \text{Re} \left[\sum_{i=1}^p \frac{A_i}{2} \exp[j(2\pi(F_i - F_0)t + \phi_i)] \exp[j2\pi F_0 t] \right] \end{aligned}$$

so that from (15.2)

$$\tilde{s}(t) = \sum_{i=1}^p \frac{A_i}{2} \exp(j\phi_i) \exp[j2\pi(F_i - F_0)t].$$

Note that the complex envelope is composed of *complex* sinusoids with frequencies $F_i - F_0$, which may be positive or negative, and *complex* amplitudes $\frac{A_i}{2} \exp(j\phi_i)$. If we sample at the Nyquist rate of $F_s = 1/\Delta = B$, we have as our signal data

$$\tilde{s}(n\Delta) = \sum_{i=1}^p \frac{A_i}{2} \exp(j\phi_i) \exp[j2\pi(F_i - F_0)n\Delta].$$

Furthermore, we can let $\tilde{s}[n] = \tilde{s}(n\Delta)$, so that

$$\tilde{s}[n] = \sum_{i=1}^p \tilde{A}_i \exp(j2\pi f_i n)$$

where $\bar{A}_i = \frac{A_i}{2} \exp(j\phi_i)$ is the complex amplitude and $f_i = (F_i - F_0)\Delta$ is the frequency of the i th discrete sinusoid. For the purposes of designing a signal processor we may assume the data model

$$\tilde{x}[n] = \sum_{i=1}^p \bar{A}_i \exp(j2\pi f_i n) + \tilde{w}[n]$$

where $\tilde{w}[n]$ is a complex noise sequence. This is a commonly used model. $\diamond$

A complex signal can also arise when the analytic signal $\tilde{s}_a(t)$ is used to represent a real low-pass signal. It is formed as

$$\tilde{s}_a(t) = s(t) + j\mathcal{H}[s(t)]$$

where $\mathcal{H}$ denotes the Hilbert transform [Papoulis 1965]. The effect of forming this complex signal is to remove the redundant negative frequency components of the Fourier transform. Hence, if $S(F)$ is bandlimited to B Hz, then $S_a(F)$ will have a Fourier transform that is zero for $F < 0$ and so can be sampled at B complex samples/sec. As an example, if

$$s(t) = \sum_{i=1}^p A_i \cos(2\pi F_i t + \phi_i)$$

then it can be shown that

$$\tilde{s}_a(t) = \sum_{i=1}^p A_i \exp[j(2\pi F_i t + \phi_i)].$$

Now that the use of complex data has been shown to be a natural outgrowth of bandpass signal processing, how do complex parameters come about? Example 15.1 illustrates this possibility. Suppose we wanted to estimate the amplitudes and phases of the p sinusoids. Then, a possible parameter set would be $\{A_1, \phi_1, A_2, \phi_2, \dots, A_p, \phi_p\}$, which consists of $2p$ real parameters. But, equivalently, we could estimate the complex parameters $\{A_1 \exp(j\phi_1), A_2 \exp(j\phi_2), \dots, A_p \exp(j\phi_p)\}$, which is only a p -dimensional but complex parameter set. The equivalence of the two parameter sets is evident from the transformation

$$\bar{A} = A \exp(j\phi)$$

and inverse transformation

$$\begin{aligned} A &= \sqrt{A_R^2 + A_I^2} \\ \phi &= \arctan \frac{A_I}{A_R}. \end{aligned}$$

Another common example where complex parameters occur is in spectral modeling of a nonsymmetric PSD. For example, an AR model of the PSD of a complex process $\tilde{x}[n]$ would assume the form [Kay 1988]

$$P_{\tilde{x}\tilde{x}}(f) = \frac{\sigma_u^2}{|1 + a[1] \exp(-j2\pi f) + \dots + a[p] \exp(-j2\pi fp)|^2}.$$

If the PSD is nonsymmetric about $f = 0$, corresponding to a complex ACF and thus a complex random process [Papoulis 1965], then the AR filter parameters $\{a[1], a[2], \dots, a[p]\}$ will be complex. Otherwise, for real $a[k]$'s we would have $P_{\tilde{x}\tilde{x}}(-f) = P_{\tilde{x}\tilde{x}}(f)$. It seems, therefore, natural to let the filter parameters be complex to allow for all possible spectra.

In dealing with complex data and/or complex parameters we can always just decompose them into their real and imaginary parts and proceed as we normally would for real data vectors and/or real parameter vectors. That there is a distinct *disadvantage* in doing so is illustrated by the following example.

Example 15.2 - Least Squares Estimation of Amplitude

Suppose we wish to minimize the LS error

$$J(\tilde{A}) = \sum_{n=0}^{N-1} |\tilde{x}[n] - \tilde{A}\tilde{s}[n]|^2$$

over $\tilde{A}$, where $\tilde{x}[n], \tilde{A}, \tilde{s}[n]$ are all complex. A straightforward approach would decompose all complex quantities into their real and imaginary parts to yield

$$\begin{aligned} J'(A_R, A_I) &= \sum_{n=0}^{N-1} |x_R[n] + jx_I[n] - (A_R + jA_I)(s_R[n] + js_I[n])|^2 \\ &= \sum_{n=0}^{N-1} (x_R[n] - A_R s_R[n] + A_I s_I[n])^2 + (x_I[n] - A_R s_I[n] - A_I s_R[n])^2. \end{aligned}$$

This is a standard quadratic form in the real variables A_R and A_I . Hence, we can let $\mathbf{x}_R = [x_R[0] \ x_R[1] \ \dots \ x_R[N-1]]^T$, $\mathbf{x}_I = [x_I[0] \ x_I[1] \ \dots \ x_I[N-1]]^T$, $\mathbf{s}_R = [s_R[0] \ s_R[1] \ \dots \ s_R[N-1]]^T$, $\mathbf{s}_I = [s_I[0] \ s_I[1] \ \dots \ s_I[N-1]]^T$, so that

$$\begin{aligned} J'(A_R, A_I) &= (\mathbf{x}_R - A_R \mathbf{s}_R + A_I \mathbf{s}_I)^T (\mathbf{x}_R - A_R \mathbf{s}_R + A_I \mathbf{s}_I) \\ &\quad + (\mathbf{x}_I - A_R \mathbf{s}_I - A_I \mathbf{s}_R)^T (\mathbf{x}_I - A_R \mathbf{s}_I - A_I \mathbf{s}_R) \end{aligned}$$

or letting $\mathbf{s}_1 = [\mathbf{s}_R \ -\mathbf{s}_I]$, $\mathbf{s}_2 = [\mathbf{s}_I \ \mathbf{s}_R]$, and $\mathbf{A} = [A_R \ A_I]^T$

$$\begin{aligned} J'(\mathbf{A}) &= (\mathbf{x}_R - \mathbf{s}_1 \mathbf{A})^T (\mathbf{x}_R - \mathbf{s}_1 \mathbf{A}) + (\mathbf{x}_I - \mathbf{s}_2 \mathbf{A})^T (\mathbf{x}_I - \mathbf{s}_2 \mathbf{A}) \\ &= \mathbf{x}_R^T \mathbf{x}_R - \mathbf{x}_R^T \mathbf{s}_1 \mathbf{A} - \mathbf{A}^T \mathbf{s}_1^T \mathbf{x}_R + \mathbf{A}^T \mathbf{s}_1^T \mathbf{s}_1 \mathbf{A} \\ &\quad + \mathbf{x}_I^T \mathbf{x}_I - \mathbf{x}_I^T \mathbf{s}_2 \mathbf{A} - \mathbf{A}^T \mathbf{s}_2^T \mathbf{x}_I + \mathbf{A}^T \mathbf{s}_2^T \mathbf{s}_2 \mathbf{A}. \end{aligned}$$

Taking the gradients and using (4.3) yields

$$\frac{\partial J'}{\partial \mathbf{A}} = -2\mathbf{s}_1^T \mathbf{x}_R + 2\mathbf{s}_1^T \mathbf{s}_1 \mathbf{A} - 2\mathbf{s}_2^T \mathbf{x}_I + 2\mathbf{s}_2^T \mathbf{s}_2 \mathbf{A}.$$

Setting this equal to zero and solving produces

$$\begin{aligned}
 \hat{\mathbf{A}} &= (\mathbf{s}_1^T \mathbf{s}_1 + \mathbf{s}_2^T \mathbf{s}_2)^{-1} (\mathbf{s}_1^T \mathbf{x}_R + \mathbf{s}_2^T \mathbf{x}_I) \\
 &= \left(\begin{bmatrix} \mathbf{s}_R^T \mathbf{s}_R & -\mathbf{s}_R^T \mathbf{s}_I \\ -\mathbf{s}_I^T \mathbf{s}_R & \mathbf{s}_I^T \mathbf{s}_I \end{bmatrix} + \begin{bmatrix} \mathbf{s}_I^T \mathbf{s}_I & \mathbf{s}_I^T \mathbf{s}_R \\ \mathbf{s}_R^T \mathbf{s}_I & \mathbf{s}_R^T \mathbf{s}_R \end{bmatrix} \right)^{-1} \\
 &\quad \cdot \left(\begin{bmatrix} \mathbf{s}_R^T \mathbf{x}_R \\ -\mathbf{s}_I^T \mathbf{x}_R \end{bmatrix} + \begin{bmatrix} \mathbf{s}_I^T \mathbf{x}_I \\ \mathbf{s}_R^T \mathbf{x}_I \end{bmatrix} \right) \\
 &= \begin{bmatrix} \mathbf{s}_R^T \mathbf{s}_R + \mathbf{s}_I^T \mathbf{s}_I & 0 \\ 0 & \mathbf{s}_I^T \mathbf{s}_I + \mathbf{s}_R^T \mathbf{s}_R \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{s}_R^T \mathbf{x}_R + \mathbf{s}_I^T \mathbf{x}_I \\ \mathbf{s}_R^T \mathbf{x}_I - \mathbf{s}_I^T \mathbf{x}_R \end{bmatrix} \\
 &= \begin{bmatrix} \frac{\mathbf{s}_R^T \mathbf{x}_R + \mathbf{s}_I^T \mathbf{x}_I}{\mathbf{s}_R^T \mathbf{s}_R + \mathbf{s}_I^T \mathbf{s}_I} \\ \frac{\mathbf{s}_R^T \mathbf{x}_I - \mathbf{s}_I^T \mathbf{x}_R}{\mathbf{s}_R^T \mathbf{s}_R + \mathbf{s}_I^T \mathbf{s}_I} \end{bmatrix}
 \end{aligned}$$

which is the minimizing solution. However, if we rewrite $\hat{\mathbf{A}}$ in complex form as $\hat{A}_R + j\hat{A}_I = \hat{\tilde{A}}$, we have

$$\begin{aligned}
 \hat{\tilde{A}} &= \frac{\mathbf{s}_R^T \mathbf{x}_R + \mathbf{s}_I^T \mathbf{x}_I + j\mathbf{s}_R^T \mathbf{x}_I - j\mathbf{s}_I^T \mathbf{x}_R}{\mathbf{s}_R^T \mathbf{s}_R + \mathbf{s}_I^T \mathbf{s}_I} \\
 &= \frac{(\mathbf{x}_R + j\mathbf{x}_I)^T (\mathbf{s}_R - j\mathbf{s}_I)}{\mathbf{s}_R^T \mathbf{s}_R + \mathbf{s}_I^T \mathbf{s}_I} \\
 &= \frac{\sum_{n=0}^{N-1} \tilde{x}[n] \tilde{s}^*[n]}{\sum_{n=0}^{N-1} |\tilde{s}[n]|^2}
 \end{aligned}$$

a result analogous to the real case. We can simplify the minimization process using complex variables as follows. First, define a complex derivative as (see Section 15.6 for further discussions) [Brandwood 1983]

$$\frac{\partial J}{\partial \tilde{A}} = \frac{1}{2} \left(\frac{\partial J}{\partial A_R} - j \frac{\partial J}{\partial A_I} \right)$$

and note that

$$\frac{\partial J}{\partial \tilde{A}} = 0 \quad \text{if and only if} \quad \frac{\partial J}{\partial \mathbf{A}} = \mathbf{0}.$$

Then, we must establish that (see Section 15.6 for the definition of a derivative of a complex function)

$$\frac{\partial \tilde{A}}{\partial \tilde{A}} = 1 \tag{15.4}$$

$$\frac{\partial \tilde{A}^*}{\partial \tilde{A}} = 0 \quad (15.5)$$

$$\frac{\partial \tilde{A} \tilde{A}^*}{\partial \tilde{A}} = \tilde{A} \frac{\partial \tilde{A}^*}{\partial \tilde{A}} + \frac{\partial \tilde{A}}{\partial \tilde{A}} \tilde{A}^* = \tilde{A}^* \quad (15.6)$$

as we shall do in Section 15.6. Now, we can minimize J using the complex derivative

$$\begin{aligned} \frac{\partial J}{\partial \tilde{A}} &= \frac{\partial}{\partial \tilde{A}} \sum_{n=0}^{N-1} |\tilde{x}[n] - \tilde{A} \tilde{s}[n]|^2 \\ &= \sum_{n=0}^{N-1} \frac{\partial}{\partial \tilde{A}} \left(|\tilde{x}[n]|^2 - \tilde{x}[n] \tilde{A}^* \tilde{s}^*[n] - \tilde{A} \tilde{s}[n] \tilde{x}^*[n] + \tilde{A} \tilde{A}^* |\tilde{s}[n]|^2 \right) \\ &= \sum_{n=0}^{N-1} \left(0 - 0 - \tilde{s}[n] \tilde{x}^*[n] + \tilde{A}^* |\tilde{s}[n]|^2 \right). \end{aligned}$$

Setting this equal to zero and solving produces the same results. Thus, we see that with the use of some easily established identities the algebra required for minimization becomes quite simple. $\diamond$

15.4 Complex Random Variables and PDFs

As we have seen, it is quite natural to encounter complex signal models. To formulate complex noise models as well we now extend many of our standard definitions and results for real random variables to complex ones. A complex random variable $\tilde{x}$ is defined to be $\tilde{x} = u + jv$, where u, v , the real and imaginary parts of $\tilde{x}$, are real random variables. The tilde denotes a complex random variable since we will have need to distinguish it from a real random variable. We will assume that the real random vector $[u \ v]^T$ possesses a joint PDF, allowing us to define moments of the complex random variable. The mean of $\tilde{x}$ is defined as

$$E(\tilde{x}) = E(u) + jE(v) \quad (15.7)$$

where the expectation is with respect to the marginal PDFs or $p(u)$ and $p(v)$, and the second moment is defined as

$$E(|\tilde{x}|^2) = E(u^2) + E(v^2). \quad (15.8)$$

The variance is defined as

$$\text{var}(\tilde{x}) = E(|\tilde{x} - E(\tilde{x})|^2) \quad (15.9)$$

which can easily be shown to reduce to

$$\text{var}(\tilde{x}) = E(|\tilde{x}|^2) - |E(\tilde{x})|^2. \quad (15.10)$$

If we now have two complex random variables $\tilde{x}_1$ and $\tilde{x}_2$ and a joint PDF for the real random vector $[u_1 u_2 v_1 v_2]^T$, then we can define a cross-moment as

$$\begin{aligned} E(\tilde{x}_1^* \tilde{x}_2) &= E[(u_1 - jv_1)(u_2 + jv_2)] \\ &= [E(u_1 u_2) + E(v_1 v_2)] + j[E(u_1 v_2) - E(u_2 v_1)] \end{aligned}$$

which is seen to involve all the possible real cross-moments. The covariance between $\tilde{x}_1$ and $\tilde{x}_2$ is defined as

$$\text{cov}(\tilde{x}_1, \tilde{x}_2) = E[(\tilde{x}_1 - E(\tilde{x}_1))^*(\tilde{x}_2 - E(\tilde{x}_2))] \quad (15.11)$$

and can be shown to reduce to

$$\text{cov}(\tilde{x}_1, \tilde{x}_2) = E(\tilde{x}_1^* \tilde{x}_2) - E^*(\tilde{x}_1)E(\tilde{x}_2). \quad (15.12)$$

Just as in the real case, if $\tilde{x}_1$ is independent of $\tilde{x}_2$, which is to say $[u_1 v_1]^T$ is independent of $[u_2 v_2]^T$, then $\text{cov}(\tilde{x}_1, \tilde{x}_2) = 0$ (see Problem 15.1). Note that the covariance reduces to the variance if $\tilde{x}_1 = \tilde{x}_2$. We can easily extend these definitions to complex random vectors. For example, if $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \dots \tilde{x}_n]^T$, then the mean of $\tilde{\mathbf{x}}$ is defined as

$$E(\tilde{\mathbf{x}}) = \begin{bmatrix} E(\tilde{x}_1) \\ E(\tilde{x}_2) \\ \vdots \\ E(\tilde{x}_n) \end{bmatrix}$$

and the covariance of $\tilde{\mathbf{x}}$ is defined as

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{x}}} &= E[(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H] \\ &= E \left\{ \begin{bmatrix} \tilde{x}_1 - E(\tilde{x}_1) \\ \tilde{x}_2 - E(\tilde{x}_2) \\ \vdots \\ \tilde{x}_n - E(\tilde{x}_n) \end{bmatrix} \begin{bmatrix} \tilde{x}_1^* - E^*(\tilde{x}_1) & \tilde{x}_2^* - E^*(\tilde{x}_2) & \dots & \tilde{x}_n^* - E^*(\tilde{x}_n) \end{bmatrix} \right\} \\ &= \begin{bmatrix} \text{var}(\tilde{x}_1) & \text{cov}(\tilde{x}_1, \tilde{x}_2) & \dots & \text{cov}(\tilde{x}_1, \tilde{x}_n) \\ \text{cov}(\tilde{x}_2, \tilde{x}_1) & \text{var}(\tilde{x}_2) & \dots & \text{cov}(\tilde{x}_2, \tilde{x}_n) \\ \vdots & \vdots & \ddots & \vdots \\ \text{cov}(\tilde{x}_n, \tilde{x}_1) & \text{cov}(\tilde{x}_n, \tilde{x}_2) & \dots & \text{var}(\tilde{x}_n) \end{bmatrix}^* \end{aligned} \quad (15.13)$$

where H denotes the conjugate transpose of a matrix. Note that the covariance matrix is Hermitian or the diagonal elements are real and the off-diagonal elements are complex conjugates of each other so that $\mathbf{C}_{\tilde{\mathbf{x}}}^H = \mathbf{C}_{\tilde{\mathbf{x}}}$. Also, $\mathbf{C}_{\tilde{\mathbf{x}}}$ can be shown to be positive semidefinite (see Problem 15.2).

We frequently need to compute the moments of linear transformations of random vectors. For example, if $\tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}} + \mathbf{b}$, where $\tilde{\mathbf{x}}$ is a complex $n \times 1$ random vector, $\mathbf{A}$ is

a complex $m \times n$ matrix, $\mathbf{b}$ is a complex $m \times 1$ vector, so that $\tilde{\mathbf{y}}$ is a complex $m \times 1$ random vector, then

$$\begin{aligned} E(\tilde{\mathbf{y}}) &= \mathbf{A}E(\tilde{\mathbf{x}}) + \mathbf{b} \\ \mathbf{C}_{\tilde{\mathbf{y}}} &= \mathbf{A}\mathbf{C}_{\tilde{\mathbf{x}}}\mathbf{A}^H. \end{aligned}$$

The first result follows from considering

$$\tilde{y}_i = \sum_{j=1}^n [\mathbf{A}]_{ij} \tilde{x}_j + b_i$$

so that

$$E(\tilde{y}_i) = \sum_{j=1}^n [\mathbf{A}]_{ij} E(\tilde{x}_j) + b_i$$

and expressing the results in matrix form. The second result uses the first to yield

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{y}}} &= E[(\tilde{\mathbf{y}} - E(\tilde{\mathbf{y}}))(\tilde{\mathbf{y}} - E(\tilde{\mathbf{y}}))^H] \\ &= E[\mathbf{A}(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H \mathbf{A}^H]. \end{aligned}$$

But

$$\begin{aligned} [\mathbf{C}_{\tilde{\mathbf{y}}}]_{ij} &= E\left\{[\mathbf{A}(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H \mathbf{A}^H]_{ij}\right\} \\ &= E\left\{\sum_{k=1}^n \sum_{l=1}^n [\mathbf{A}]_{ik} [(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H]_{kl} [\mathbf{A}^H]_{lj}\right\} \\ &= \sum_{k=1}^n \sum_{l=1}^n [\mathbf{A}]_{ik} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{kl} [\mathbf{A}^H]_{lj} \end{aligned}$$

and expressing this in matrix form yields the desired result. It is also of interest to determine the first two moments of a positive definite Hermitian form or of

$$Q = \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}$$

where $\tilde{\mathbf{x}}$ is a complex $n \times 1$ random vector and $\mathbf{A}$ is an $n \times n$ positive definite hermitian ($\mathbf{A}^H = \mathbf{A}$) matrix. Note that Q is real since

$$\begin{aligned} Q^* = Q^H &= (\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}})^H \\ &= \tilde{\mathbf{x}}^H \mathbf{A}^H \tilde{\mathbf{x}} \\ &= \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}} \\ &= Q. \end{aligned}$$

Also, since $\mathbf{A}$ is assumed to be positive definite, we have $Q > 0$ for all $\tilde{\mathbf{x}} \neq \mathbf{0}$. To find the moments of Q we assume $E(\tilde{\mathbf{x}}) = \mathbf{0}$. (If this is not the case, we can easily modify the results by replacing $\tilde{\mathbf{x}}$ by $\tilde{\mathbf{y}} = \tilde{\mathbf{x}} - E(\tilde{\mathbf{x}})$ and then evaluating the expressions.) Thus,

$$\begin{aligned} E(Q) &= E(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}) \\ &= E(\text{tr}(\mathbf{A} \tilde{\mathbf{x}} \tilde{\mathbf{x}}^H)) \end{aligned}$$

since $\tilde{\mathbf{x}}^H \tilde{\mathbf{y}} = \text{tr}(\tilde{\mathbf{y}} \tilde{\mathbf{x}}^H)$. But then

$$\begin{aligned} E(Q) &= \text{tr}[E(\mathbf{A} \tilde{\mathbf{x}} \tilde{\mathbf{x}}^H)] \\ &= \text{tr}[\mathbf{A} E(\tilde{\mathbf{x}} \tilde{\mathbf{x}}^H)] \\ &= \text{tr}(\mathbf{A} \mathbf{C}_{\tilde{\mathbf{x}}}). \end{aligned} \quad (15.14)$$

To find the second moment we need to evaluate

$$E(Q^2) = E(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}} \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}) \quad (15.15)$$

which will require fourth-order moments of $\tilde{\mathbf{x}}$ (see Example 15.4). Recall that for a real Gaussian PDF the fourth-order moments were functions of second-order moments, considerably simplifying the evaluation. We will show next that we can define a *complex* Gaussian PDF for complex random variables and that it will have many of the same properties as the real one. Then, we will be able to evaluate complicated expressions such as (15.15). Furthermore, the complex Gaussian PDF arises naturally from a consideration of the distribution of the complex envelope of a bandpass process.

We begin our discussion by defining the complex Gaussian PDF of a scalar complex random variable $\tilde{x}$. Since $\tilde{x} = u + jv$, any complete statistical description will involve the joint PDF of u and v . The complex Gaussian PDF assumes u to be *independent* of v . Furthermore, it assumes that the real and imaginary parts are distributed as $\mathcal{N}(\mu_u, \sigma^2/2)$ and $\mathcal{N}(\mu_v, \sigma^2/2)$, respectively. Hence, the joint PDF of the real random variables becomes

$$\begin{aligned} p(u, v) &= \frac{1}{\sqrt{2\pi\frac{\sigma^2}{2}}} \exp\left[-\frac{1}{2(\frac{\sigma^2}{2})}(u - \mu_u)^2\right] \\ &\quad \cdot \frac{1}{\sqrt{2\pi\frac{\sigma^2}{2}}} \exp\left[-\frac{1}{2(\frac{\sigma^2}{2})}(v - \mu_v)^2\right] \\ &= \frac{1}{\pi\sigma^2} \exp\left[-\frac{1}{\sigma^2}((u - \mu_u)^2 + (v - \mu_v)^2)\right]. \end{aligned}$$

But letting $\tilde{\mu} = E(\tilde{x}) = \mu_u + j\mu_v$, we have in more succinct form

$$p(\tilde{x}) = \frac{1}{\pi\sigma^2} \exp\left[-\frac{1}{\sigma^2}|\tilde{x} - \tilde{\mu}|^2\right]. \quad (15.16)$$

Since the joint PDF depends on u and v only through $\tilde{x}$, we can view the PDF to be that of the scalar random variable $\tilde{x}$, as our notation suggests. This is called the *complex Gaussian PDF* for a scalar complex random variable and is denoted by $\mathcal{CN}(\tilde{\mu}, \sigma^2)$. Note the similarity to the usual real Gaussian PDF. We would expect $p(\tilde{x})$ to have many of the same algebraic properties as the real Gaussian PDF, and indeed it does. To extend these results we next consider a complex random vector $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \dots \tilde{x}_n]^T$. Assume the components of $\tilde{\mathbf{x}}$ are each distributed as $\mathcal{CN}(\tilde{\mu}_i, \sigma_i^2)$ and are also independent. By independence we mean that the real random vectors $[u_1 \ v_1]^T, [u_2 \ v_2]^T, \dots, [u_n \ v_n]^T$ are

independent. Then, the multivariate complex Gaussian PDF is just the product of the marginal PDFs or

$$p(\tilde{\mathbf{x}}) = \prod_{i=1}^n p(\tilde{x}_i)$$

which follows from the usual property of PDFs for real independent random variables. From (15.16) this can be written as

$$p(\tilde{\mathbf{x}}) = \frac{1}{\pi^n \prod_{i=1}^n \sigma_i^2} \exp \left[- \sum_{i=1}^n \frac{1}{\sigma_i^2} |\tilde{x}_i - \tilde{\mu}_i|^2 \right].$$

But noting that for independent complex random variables the covariances are zero, we have the covariance matrix for $\tilde{\mathbf{x}}$

$$\mathbf{C}_{\tilde{\mathbf{x}}} = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_n^2)$$

so that the PDF becomes

$$p(\tilde{\mathbf{x}}) = \frac{1}{\pi^n \det(\mathbf{C}_{\tilde{\mathbf{x}}})} \exp [- (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})]. \quad (15.17)$$

This is the *multivariate complex Gaussian PDF*, and it is denoted by $\mathcal{CN}(\tilde{\boldsymbol{\mu}}, \mathbf{C}_{\tilde{\mathbf{x}}})$. Again note the similarity to the usual real multivariate Gaussian PDF. But (15.17) is more general than we have assumed in our derivation. It is actually valid for complex covariance matrices other than just diagonal ones. To define the general complex Gaussian PDF we need to restrict the form of the underlying *real* covariance matrix, as we now show. Recall that for a scalar complex Gaussian random variable $\tilde{x} = u + jv$ the *real* covariance matrix of $[u \ v]^T$ is

$$\begin{bmatrix} \frac{\sigma_1^2}{2} & 0 \\ 0 & \frac{\sigma_1^2}{2} \end{bmatrix}.$$

For a 2×1 complex random vector $\tilde{\mathbf{x}} = [\tilde{x}_1 \ \tilde{x}_2]^T = [u_1 + jv_1 \ u_2 + jv_2]^T$ with $\tilde{x}_1$ independent of $\tilde{x}_2$, as we have assumed, the real covariance matrix of $[u_1 \ v_1 \ u_2 \ v_2]^T$ is

$$\frac{1}{2} \begin{bmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_1^2 & 0 & 0 \\ 0 & 0 & \sigma_2^2 & 0 \\ 0 & 0 & 0 & \sigma_2^2 \end{bmatrix}.$$

If we rearrange the real random vector as $\mathbf{x} = [u_1 \ u_2 \ v_1 \ v_2]^T$, then the real covariance matrix becomes

$$\mathbf{C}_x = \frac{1}{2} \left[\begin{array}{cc|cc} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_1^2 & 0 & 0 \\ \hline 0 & 0 & \sigma_2^2 & 0 \\ 0 & 0 & 0 & \sigma_2^2 \end{array} \right]. \quad (15.18)$$

This is a special case of the more general form of the real 4×4 covariance matrix

$$\mathbf{C}_x = \frac{1}{2} \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \quad (15.19)$$

where $\mathbf{A}$ and $\mathbf{B}$ are each 2×2 matrices. This form of $\mathbf{C}_x$ allows us to define a complex Gaussian PDF for complex Gaussian random variables that are *correlated*. By letting $\mathbf{u} = [u_1 \ u_2]^T$ and $\mathbf{v} = [v_1 \ v_2]^T$ we see that the submatrices of $\mathbf{C}_x$ are

$$\begin{aligned} \frac{1}{2}\mathbf{A} &= E[(\mathbf{u} - E(\mathbf{u}))(\mathbf{u} - E(\mathbf{u}))^T] = E[(\mathbf{v} - E(\mathbf{v}))(\mathbf{v} - E(\mathbf{v}))^T] \\ \frac{1}{2}\mathbf{B} &= E[(\mathbf{v} - E(\mathbf{v}))(\mathbf{u} - E(\mathbf{u}))^T] = -E[(\mathbf{u} - E(\mathbf{u}))(\mathbf{v} - E(\mathbf{v}))^T]. \end{aligned}$$

Thus $\mathbf{A}$ is symmetric and $\mathbf{B}$ is skew-symmetric or $\mathbf{B}^T = -\mathbf{B}$. Explicitly, the real covariance matrix is

$$\mathbf{C}_x = \begin{bmatrix} \text{var}(u_1) & \text{cov}(u_1, u_2) & \text{cov}(u_1, v_1) & \text{cov}(u_1, v_2) \\ \text{cov}(u_2, u_1) & \text{var}(u_2) & \text{cov}(u_2, v_1) & \text{cov}(u_2, v_2) \\ \text{cov}(v_1, u_1) & \text{cov}(v_1, u_2) & \text{var}(v_1) & \text{cov}(v_1, v_2) \\ \text{cov}(v_2, u_1) & \text{cov}(v_2, u_2) & \text{cov}(v_2, v_1) & \text{var}(v_2) \end{bmatrix}.$$

For $\mathbf{C}_x$ to have the form of (15.19) we require (in addition to variances being equal or $\text{var}(u_i) = \text{var}(v_i)$ for $i = 1, 2$, and the covariance between the real and imaginary parts being zero or $\text{cov}(u_i, v_i) = 0$ for $i = 1, 2$, as we have assumed for the scalar complex Gaussian random variable) the covariances to satisfy

$$\begin{aligned} \text{cov}(u_1, u_2) &= \text{cov}(v_1, v_2) \\ \text{cov}(u_1, v_2) &= -\text{cov}(u_2, v_1). \end{aligned} \quad (15.20)$$

The covariance between the real parts is identical to that between the imaginary parts, and the covariance between the real part of $\tilde{x}_1$ and the imaginary part of $\tilde{x}_2$ is the negative of the covariance between the real part of $\tilde{x}_2$ and the imaginary part of $\tilde{x}_1$. With this form of the real covariance matrix, as in (15.19), we can prove the following (see Appendix 15A) for $\tilde{\mathbf{x}} = [\tilde{x}_1 \ \tilde{x}_2 \ \dots \ \tilde{x}_n] = [u_1 + jv_1 \ u_2 + jv_2 \ \dots \ u_n + jv_n]^T$ and for $\mathbf{A}$ and $\mathbf{B}$ both having dimensions $n \times n$.

1. The complex covariance matrix of $\tilde{\mathbf{x}}$ is $\mathbf{C}_{\tilde{\mathbf{x}}} = \mathbf{A} + j\mathbf{B}$. As an example, for $n = 2$ and independent $\tilde{x}_1, \tilde{x}_2$, $\mathbf{C}_x$ is given by (15.18) and thus

$$\mathbf{C}_{\tilde{\mathbf{x}}} = \mathbf{A} + j\mathbf{B} = \mathbf{A} = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}.$$

2. The quadratic form in the exponent of the real multivariate Gaussian PDF may be expressed as

$$(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}_x^{-1} (\mathbf{x} - \boldsymbol{\mu}) = 2(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})$$

where $\tilde{\boldsymbol{\mu}} = E(\tilde{\mathbf{x}}) = \boldsymbol{\mu}_u + j\boldsymbol{\mu}_v$. As an example, for $n = 2$ and independent $\tilde{x}_1, \tilde{x}_2$, we have from (15.18) and (15.19)

$$\mathbf{C}_x = \begin{bmatrix} \frac{\mathbf{A}}{2} & \mathbf{0} \\ \mathbf{0} & \frac{\mathbf{A}}{2} \end{bmatrix}$$

so that

$$\begin{aligned} & (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}_x^{-1} (\mathbf{x} - \boldsymbol{\mu}) \\ &= [(\mathbf{u} - \boldsymbol{\mu}_u)^T \quad (\mathbf{v} - \boldsymbol{\mu}_v)^T] 2 \begin{bmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{u} - \boldsymbol{\mu}_u \\ \mathbf{v} - \boldsymbol{\mu}_v \end{bmatrix} \\ &= 2[(\mathbf{u} - \boldsymbol{\mu}_u)^T \mathbf{A}^{-1} (\mathbf{u} - \boldsymbol{\mu}_u)] + [(\mathbf{v} - \boldsymbol{\mu}_v)^T \mathbf{A}^{-1} (\mathbf{v} - \boldsymbol{\mu}_v)] \\ &= 2[(\mathbf{u} - \boldsymbol{\mu}_u) - j(\mathbf{v} - \boldsymbol{\mu}_v)]^T \mathbf{C}_{\tilde{x}}^{-1} [(\mathbf{u} - \boldsymbol{\mu}_u) + j(\mathbf{v} - \boldsymbol{\mu}_v)] \\ &= 2(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\tilde{x}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}) \end{aligned}$$

since $\mathbf{C}_{\tilde{x}} = \mathbf{A}$, where $\mathbf{A}$ is real.

3. The determinant of $\mathbf{C}_x$ is

$$\det(\mathbf{C}_x) = \frac{1}{2^{2n}} \det^2(\mathbf{C}_{\tilde{x}}).$$

Note that $\det(\mathbf{C}_{\tilde{x}})$ is real and positive since $\mathbf{C}_{\tilde{x}}$ is Hermitian and positive definite. As an example, for $n = 2$ and independent $\tilde{x}_1, \tilde{x}_2$ we have $\mathbf{C}_{\tilde{x}} = \mathbf{A}$, and from (15.18) and (15.19)

$$\mathbf{C}_x = \begin{bmatrix} \frac{\mathbf{A}}{2} & \mathbf{0} \\ \mathbf{0} & \frac{\mathbf{A}}{2} \end{bmatrix}$$

so that

$$\begin{aligned} \det(\mathbf{C}_x) &= \det^2 \left(\frac{\mathbf{A}}{2} \right) \\ &= \frac{1}{2^4} \det^2(\mathbf{A}) \\ &= \frac{1}{2^4} \det^2(\mathbf{C}_{\tilde{x}}). \end{aligned}$$

With these results we can rewrite the real multivariate Gaussian PDF of $\mathbf{x} = [u_1 \dots u_n v_1 \dots v_n]^T$

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{2n}{2}} \det^{\frac{1}{2}}(\mathbf{C}_x)} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}_x^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]$$

as

$$\begin{aligned} p(\mathbf{x}) &= p(\tilde{\mathbf{x}}) \\ &= \frac{1}{\pi^n \det(\mathbf{C}_{\tilde{\mathbf{x}}})} \exp [-(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})] \end{aligned}$$

where $\tilde{\boldsymbol{\mu}}$ is the mean and $\mathbf{C}_{\tilde{\mathbf{x}}}$ is the covariance matrix of $\tilde{\mathbf{x}}$. Whereas the real Gaussian PDF involves $2n \times 1$ real vectors and $2n \times 2n$ real matrices, the complex Gaussian PDF involves $n \times 1$ complex vectors and $n \times n$ complex matrices. We now summarize our results in a theorem.

Theorem 15.1 (Complex Multivariate Gaussian PDF) *If a real random vector $\mathbf{x}$ of dimension $2n \times 1$ can be partitioned as*

$$\mathbf{x} = \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} n \times 1 \\ n \times 1 \end{bmatrix}$$

where $\mathbf{u}, \mathbf{v}$ are real random vectors and $\mathbf{x}$ has the PDF

$$\mathbf{x} \sim \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu}_u \\ \boldsymbol{\mu}_v \end{bmatrix}, \begin{bmatrix} \mathbf{C}_{uu} & \mathbf{C}_{uv} \\ \mathbf{C}_{vu} & \mathbf{C}_{vv} \end{bmatrix} \right)$$

and $\mathbf{C}_{uu} = \mathbf{C}_{vv}$ and $\mathbf{C}_{uv} = -\mathbf{C}_{vu}$ or

$$\begin{aligned} \text{cov}(u_i, u_j) &= \text{cov}(v_i, v_j) \\ \text{cov}(u_i, v_j) &= -\text{cov}(v_i, u_j) \end{aligned} \quad (15.21)$$

then defining the $n \times 1$ complex random vector $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$, $\tilde{\mathbf{x}}$ has the complex multivariate Gaussian PDF

$$\tilde{\mathbf{x}} \sim \mathcal{CN}(\tilde{\boldsymbol{\mu}}, \mathbf{C}_{\tilde{\mathbf{x}}})$$

where

$$\begin{aligned} \tilde{\boldsymbol{\mu}} &= \boldsymbol{\mu}_u + j\boldsymbol{\mu}_v \\ \mathbf{C}_{\tilde{\mathbf{x}}} &= 2(\mathbf{C}_{uu} + j\mathbf{C}_{vu}) \end{aligned}$$

or more explicitly

$$p(\tilde{\mathbf{x}}) = \frac{1}{\pi^n \det(\mathbf{C}_{\tilde{\mathbf{x}}})} \exp [-(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})]. \quad (15.22)$$

It is important to realize that the complex Gaussian PDF is actually just a different algebraic form of the real Gaussian PDF. What we have essentially done is to replace 2×1 vectors by complex numbers or $[u \ v]^T \rightarrow \tilde{x} = u + jv$. This will simplify subsequent estimator calculations, but no new theory should be expected. The deeper reason that allowed us to perform this trick is the *isomorphism* that exists between the vector spaces M^2 (the space of all special 2×2 real matrices) and C^1 (the space of all complex

numbers). We can transform matrices from M^2 to complex numbers in C^1 , perform calculations in C^1 , and after we're done transform back to M^2 . As an example, if we wish to multiply the matrices

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} e & -f \\ f & e \end{bmatrix} = \begin{bmatrix} ae - bf & -af - be \\ af + be & ae - bf \end{bmatrix}$$

we can, equivalently, multiply the complex numbers

$$(a + jb)(e + jf) = (ae - bf) + j(af + be)$$

and then transform back to M^2 . We will explore this further in Problem 15.5.

There are many properties of the complex Gaussian PDF that mirror those of the real one. We now summarize the important properties with the proofs given in Appendix 15B.

1. Any subvector of a complex Gaussian random vector is also complex Gaussian. In particular, the marginals are complex Gaussian.
2. If $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \dots \tilde{x}_n]^T$ is complex Gaussian and $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n\}$ are uncorrelated, then they are also independent.
3. If $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n\}$ are independent, each one being complex Gaussian, then $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \dots \tilde{x}_n]^T$ is complex Gaussian.
4. Linear (actually affine) transformations of complex Gaussian random vectors are again complex Gaussian. More specifically, if $\tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}} + \mathbf{b}$, where $\mathbf{A}$ is a complex $m \times n$ matrix with $m \leq n$ and full rank (so that $\mathbf{C}_{\tilde{\mathbf{y}}}$ is invertible), $\mathbf{b}$ is a complex $m \times 1$ vector, and if $\tilde{\mathbf{x}} \sim \mathcal{CN}(\tilde{\boldsymbol{\mu}}, \mathbf{C}_{\tilde{\mathbf{x}}})$, then

$$\tilde{\mathbf{y}} \sim \mathcal{CN}(\mathbf{A}\tilde{\boldsymbol{\mu}} + \mathbf{b}, \mathbf{A}\mathbf{C}_{\tilde{\mathbf{x}}}\mathbf{A}^H). \quad (15.23)$$

5. The sum of independent complex Gaussian random variables is also complex Gaussian.
6. If $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \tilde{x}_3 \tilde{x}_4]^T \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{\mathbf{x}}})$, then

$$E(\tilde{x}_1^* \tilde{x}_2 \tilde{x}_3^* \tilde{x}_4) = E(\tilde{x}_1^* \tilde{x}_2)E(\tilde{x}_3^* \tilde{x}_4) + E(\tilde{x}_1^* \tilde{x}_4)E(\tilde{x}_2 \tilde{x}_3^*). \quad (15.24)$$

It is interesting to note that $E(\tilde{x}_1^* \tilde{x}_3^*) = E(\tilde{x}_2 \tilde{x}_4) = 0$ (as we will show shortly), so that this result is analogous to the real case, with the last product being equal to zero.

7. If $[\tilde{\mathbf{x}}^T \tilde{\mathbf{y}}^T]^T$ is a complex Gaussian random vector, where $\tilde{\mathbf{x}}$ is $k \times 1$ and $\tilde{\mathbf{y}}$ is $l \times 1$ and

$$\begin{bmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{y}} \end{bmatrix} \sim \mathcal{CN} \left(\begin{bmatrix} E(\tilde{\mathbf{x}}) \\ E(\tilde{\mathbf{y}}) \end{bmatrix}, \begin{bmatrix} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} & \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} \\ \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} & \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{y}}} \end{bmatrix} \right)$$

then the conditional PDF $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$ is also complex Gaussian with

$$E(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = E(\tilde{\mathbf{y}}) + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}})) \quad (15.25)$$

$$\mathbf{C}_{\tilde{\mathbf{y}}|\tilde{\mathbf{x}}} = \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{y}}} - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}}. \quad (15.26)$$

This is identical in form to the real case (see (10.24) and (10.25)).

These properties appear reasonable, and the reader will be inclined to accept them. The only one that seems strange is property 6, which seems to imply that for a zero mean complex Gaussian random vector we have $E(\tilde{x}_1 \tilde{x}_2) = 0$. To better appreciate this conjecture we determine this cross-moment as follows:

$$\begin{aligned} E(\tilde{x}_1 \tilde{x}_2) &= E[(u_1 + jv_1)(u_2 + jv_2)] \\ &= E(u_1 u_2) - E(v_1 v_2) + j[E(u_1 v_2) + E(v_1 u_2)] \\ &= \text{cov}(u_1, u_2) - \text{cov}(v_1, v_2) + j[\text{cov}(u_1, v_2) + \text{cov}(u_2, v_1)] \end{aligned}$$

which from (15.20) is zero. In fact, we have actually *defined* $E(\tilde{x}_1 \tilde{x}_2) = 0$ in arriving at the form of the covariance matrix in (15.19) (see Problem 15.9). It is this constraint that leads to the special form assumed for the real covariance matrix. We will now illustrate some of the foregoing properties by applying them in an important example.

Example 15.3 - PDF of Discrete Fourier Transform of WGN

Consider the real data $\{x[0], x[1], \dots, x[N-1]\}$ which are samples of a WGN process with zero mean and variance $\sigma^2/2$. We take the discrete Fourier transform (DFT) of the samples to produce

$$X(f_k) = \sum_{n=0}^{N-1} x[n] \exp(-j2\pi f_k n)$$

for $f_k = k/N$ and $k = 0, 1, \dots, N-1$. We inquire as to the PDF of the complex DFT outputs for $k = 1, 2, \dots, N/2 - 1$, neglecting those for $k = N/2 + 1, N/2 + 2, \dots, N-1$ since these are related to the chosen outputs by $X(f_k) = X^*(f_{N-k})$. We also neglect $X(f_0)$ and $X(f_{N/2})$ since these are purely real and are usually of little interest anyway, occurring at DC and at the Nyquist frequency. We will show that the PDF is

$$\mathbf{X}(f) = \begin{bmatrix} X(f_1) \\ \vdots \\ X(f_{\frac{N}{2}-1}) \end{bmatrix} \sim \mathcal{CN}\left(\mathbf{0}, \frac{N\sigma^2}{2} \mathbf{I}\right). \quad (15.27)$$

To verify this we first note that the real random vector

$$\left[\text{Re}(X(f_1)) \quad \dots \quad \text{Re}(X(f_{\frac{N}{2}-1})) \quad \text{Im}(X(f_1)) \quad \dots \quad \text{Im}(X(f_{\frac{N}{2}-1})) \right]^T$$

is a real Gaussian random vector since it is a linear transformation of $\mathbf{x}$. Thus, we need only show that the special form of the real covariance matrix of $\mathbf{X}$ is satisfied. First, we show that the real and imaginary parts of $X(f_k)$ are uncorrelated. Let $X(f_k) = U(f_k) + jV(f_k)$, where U and V are both real. Then

$$\begin{aligned}U(f_k) &= \sum_{n=0}^{N-1} x[n] \cos 2\pi f_k n \\V(f_k) &= -\sum_{n=0}^{N-1} x[n] \sin 2\pi f_k n\end{aligned}$$

and therefore

$$E(U(f_k)) = \sum_{n=0}^{N-1} E(x[n]) \cos 2\pi f_k n = 0$$

and similarly $E(V(f_k)) = 0$. Now

$$\begin{aligned}\text{cov}(U(f_k), U(f_l)) &= E(U(f_k)U(f_l)) \\&= E \left[\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m]x[n] \cos 2\pi f_k m \cos 2\pi f_l n \right] \\&= \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \frac{\sigma^2}{2} \delta[m-n] \cos 2\pi f_k m \cos 2\pi f_l n \\&= \frac{\sigma^2}{2} \sum_{n=0}^{N-1} \cos 2\pi f_k n \cos 2\pi f_l n \\&= \frac{\sigma^2}{2} \sum_{n=0}^{N-1} \frac{1}{2} [\cos 2\pi(f_k + f_l)n + \cos 2\pi(f_k - f_l)n] \\&= \frac{\sigma^2}{4} \sum_{n=0}^{N-1} \cos 2\pi(f_k - f_l)n\end{aligned}$$

due to the DFT relation in Appendix 1. Now if $f_k \neq f_l$, $\text{cov}(U(f_k), U(f_l)) = 0$ by the same DFT relation. If $f_k = f_l$, we have $\text{cov}(U(f_k), U(f_k)) = N\sigma^2/4$. A similar calculation shows that

$$\text{cov}(V(f_k), V(f_l)) = \begin{cases} 0 & k \neq l \\ \frac{N\sigma^2}{4} & k = l \end{cases}$$

and so

$$\mathbf{C}_{UU} = \mathbf{C}_{VV} = \frac{N\sigma^2}{4} \mathbf{I}$$

as required. To verify that $\mathbf{C}_{UV} = -\mathbf{C}_{VU}$ we show that these are both zero.

$$\begin{aligned}\text{cov}(U(f_k), V(f_l)) &= E(U(f_k)V(f_l)) \\ &= E\left(\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m]x[n] \cos 2\pi f_k m \sin 2\pi f_l n\right) \\ &= \frac{\sigma^2}{2} \sum_{n=0}^{N-1} \cos 2\pi f_k n \sin 2\pi f_l n \\ &= 0 \quad \text{for all } k, l\end{aligned}$$

by the DFT relation in Appendix 1. Similarly, $\text{cov}(V(f_k), U(f_l)) = 0$ for all k, l . Hence, $\mathbf{C}_{UV} = -\mathbf{C}_{VU} = \mathbf{0}$. From Theorem 15.1 the complex covariance matrix becomes

$$\begin{aligned}\mathbf{C}_X &= 2(\mathbf{C}_{UU} + j\mathbf{C}_{VU}) \\ &= \frac{N\sigma^2}{2} \mathbf{I}\end{aligned}$$

and the assertion is proved. Observe that the DFT coefficients are independent and identically distributed. As an application of this result we may wish to estimate σ^2 using

$$\hat{\sigma}^2 = \frac{1}{c} \sum_{k=1}^{\frac{N}{2}-1} |X(f_k)|^2.$$

This is frequently required for normalization purposes in determining a threshold for a detector [Knight, Pridham, and Kay 1981]. For the estimator to be unbiased we need to choose an appropriate value for c . But

$$E(\hat{\sigma}^2) = \frac{1}{c} \sum_{k=1}^{\frac{N}{2}-1} E(|X(f_k)|^2).$$

Since $X(f_k) \sim \mathcal{CN}(0, N\sigma^2/2)$, we have that

$$\text{var}(X(f_k)) = E(|X(f_k)|^2) = \frac{N\sigma^2}{2}$$

and

$$E(\hat{\sigma}^2) = \frac{1}{c} \sum_{k=1}^{\frac{N}{2}-1} \frac{N\sigma^2}{2} = \frac{N\sigma^2(\frac{N}{2}-1)}{c}.$$

Hence, c should be chosen to be $N(N/2 - 1)/2$. ◇

Example 15.4 - Variance of Hermitian Form

As a second application, we now derive the variance of the Hermitian form

$$Q = \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}$$

where $\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{\mathbf{x}}})$ and $\mathbf{A}$ is Hermitian. The mean has already been given in (15.14) as $\text{tr}(\mathbf{A}\mathbf{C}_{\tilde{\mathbf{x}}})$. We first determine the second moment

$$\begin{aligned} E(Q^2) &= E(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}} \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}) \\ &= E\left(\sum_{i=1}^n \sum_{j=1}^n \tilde{x}_i^* [\mathbf{A}]_{ij} \tilde{x}_j \sum_{k=1}^n \sum_{l=1}^n \tilde{x}_k^* [\mathbf{A}]_{kl} \tilde{x}_l\right) \\ &= \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n [\mathbf{A}]_{ij} [\mathbf{A}]_{kl} E(\tilde{x}_i^* \tilde{x}_j \tilde{x}_k^* \tilde{x}_l). \end{aligned}$$

Using the fourth-order moment property (15.24), we have

$$\begin{aligned} E(\tilde{x}_i^* \tilde{x}_j \tilde{x}_k^* \tilde{x}_l) &= E(\tilde{x}_i^* \tilde{x}_j) E(\tilde{x}_k^* \tilde{x}_l) + E(\tilde{x}_i^* \tilde{x}_l) E(\tilde{x}_k^* \tilde{x}_j) \\ &= [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{ij} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kl} + [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{il} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kj}. \end{aligned}$$

Thus,

$$\begin{aligned} E(Q^2) &= \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n [\mathbf{A}]_{ij} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{ij} [\mathbf{A}]_{kl} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kl} \\ &\quad + \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n [\mathbf{A}]_{ij} [\mathbf{A}]_{kl} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{il} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kj} \\ &= \sum_{i=1}^n \sum_{j=1}^n [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{ij} [\mathbf{A}]_{ij} \sum_{k=1}^n \sum_{l=1}^n [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kl} [\mathbf{A}]_{kl} \\ &\quad + \sum_{i=1}^n \sum_{k=1}^n \left(\sum_{j=1}^n [\mathbf{A}]_{ij} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{kj} \right) \left(\sum_{l=1}^n [\mathbf{A}]_{kl} [\mathbf{C}_{\tilde{\mathbf{x}}}^T]_{il} \right). \quad (15.28) \end{aligned}$$

But

$$\text{tr}(\mathbf{B}^T \mathbf{A}) = \text{tr} \left(\begin{bmatrix} \mathbf{b}_1^T \\ \vdots \\ \mathbf{b}_n^T \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_n \end{bmatrix} \right)$$

where $\mathbf{a}_i$, $\mathbf{b}_i$ are the i th columns of $\mathbf{A}$ and $\mathbf{B}$, respectively. More explicitly this becomes

$$\begin{aligned}
 \text{tr}(\mathbf{B}^T \mathbf{A}) &= \sum_{i=1}^n \mathbf{b}_i^T \mathbf{a}_i \\
 &= \sum_{i=1}^n \sum_{j=1}^n [\mathbf{B}]_{ji} [\mathbf{A}]_{ji} \\
 &= \sum_{i=1}^n \sum_{j=1}^n [\mathbf{B}]_{ij} [\mathbf{A}]_{ij}
 \end{aligned}$$

so that the first term in (15.28) is just $\text{tr}^2(\mathbf{C}_{\tilde{x}} \mathbf{A})$ or $\text{tr}^2(\mathbf{A} \mathbf{C}_{\tilde{x}})$, the squared mean. Let $\mathbf{D} = \mathbf{A} \mathbf{C}_{\tilde{x}}$. Then, the second term is

$$\sum_{i=1}^n \sum_{k=1}^n [\mathbf{D}]_{ik} [\mathbf{D}]_{ki} = \sum_{k=1}^n \sum_{i=1}^n [\mathbf{D}^T]_{ki} [\mathbf{D}]_{ki}$$

which is

$$\text{tr}(\mathbf{D} \mathbf{D}) = \text{tr}(\mathbf{A} \mathbf{C}_{\tilde{x}} \mathbf{A} \mathbf{C}_{\tilde{x}})$$

so that finally we have for $\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{x}})$ and $\mathbf{A}$ Hermitian

$$E(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}) = \text{tr}(\mathbf{A} \mathbf{C}_{\tilde{x}}) \quad (15.29)$$

$$\text{var}(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}) = \text{tr}(\mathbf{A} \mathbf{C}_{\tilde{x}} \mathbf{A} \mathbf{C}_{\tilde{x}}). \quad (15.30)$$

An application of this result is given in Problem 15.10. $\diamond$

15.5 Complex WSS Random Processes

We are now in a position to discuss the statistical modeling of the complex random process $\tilde{x}[n] = u[n] + jv[n]$. It will be assumed that $\tilde{x}[n]$ is WSS, which, as for real processes, implies that the mean does not depend on n and the covariance between $\tilde{x}[n]$ and $\tilde{x}[n+k]$ depends only on the lag k . The mean is defined as

$$E(\tilde{x}[n]) = E(u[n]) + jE(v[n])$$

and will be assumed to be zero for all n . Since $\tilde{x}[n]$ is WSS, we can define an ACF as

$$r_{\tilde{x}\tilde{x}}[k] = E(\tilde{x}^*[n] \tilde{x}[n+k]) \quad (15.31)$$

and as usual the PSD is the Fourier transform of $r_{\tilde{x}\tilde{x}}[k]$. If the random vector formed from any subset of samples, say $\{\tilde{x}[n_1], \tilde{x}[n_2], \dots, \tilde{x}[n_k]\}$, has a multivariate complex Gaussian PDF, then we say that $\tilde{x}[n]$ is a WSS complex Gaussian random process. Its

complete statistical description is then summarized by the ACF, which determines the covariance matrix, or, equivalently, by the PSD (since the mean is zero). Since the real covariance matrix is constrained to be of the form (15.19), this translates into a constraint on $r_{\tilde{x}\tilde{x}}[k]$. To interpret this constraint recall that we require

$$\begin{aligned} \mathbf{C}_{uu} &= \mathbf{C}_{vv} \\ \mathbf{C}_{uv} &= -\mathbf{C}_{vu} \end{aligned}$$

for all $\mathbf{u} = [u[n_1] \dots u[n_k]]^T$ and $\mathbf{v} = [v[n_1] \dots v[n_k]]^T$. This means that

$$\begin{aligned} E(u[n_i]u[n_j]) &= E(v[n_i]v[n_j]) \\ E(u[n_i]v[n_j]) &= -E(v[n_i]u[n_j]) \end{aligned}$$

for all i, j . By defining the ACF of $\tilde{x}[n]$ it is implicit that $u[n], v[n]$ are both WSS, as well as jointly WSS. If we define the cross-correlation function (CCF) as

$$r_{uv}[k] = E(u[n]v[n+k]) \quad (15.32)$$

then we must have

$$\begin{aligned} r_{uu}[k] &= r_{vv}[k] \\ r_{uv}[k] &= -r_{vu}[k] \end{aligned} \quad (15.33)$$

for all k , where we have let $k = n_j - n_i$. Alternatively, we require

$$\begin{aligned} P_{uu}(f) &= P_{vv}(f) \\ P_{uv}(f) &= -P_{vu}(f) \end{aligned} \quad (15.34)$$

where $P_{uu}(f), P_{vv}(f)$ are the auto-PSDs and $P_{uv}(f), P_{vu}(f)$ are the cross-PSDs. Since the cross-PSDs are complex conjugates of each other, the second condition is equivalent to requiring the cross-PSD to be purely imaginary. Then, if $u[n], v[n]$ are real jointly Gaussian random processes with zero mean and PSDs satisfying (15.34), $\tilde{x}[n]$ will be a complex Gaussian random process. Furthermore, the PSD of $\tilde{x}[n]$ can be related to those of $u[n]$ and $v[n]$ as follows. From (15.31) and (15.33)

$$\begin{aligned} r_{\tilde{x}\tilde{x}}[k] &= E[(u[n] - jv[n])(u[n+k] + jv[n+k])] \\ &= r_{uu}[k] + jr_{uv}[k] - jr_{vu}[k] + r_{vv}[k]. \end{aligned}$$

But from (15.33) this reduces to

$$r_{\tilde{x}\tilde{x}}[k] = 2r_{uu}[k] + 2jr_{uv}[k] \quad (15.35)$$

and thus

$$P_{\tilde{x}\tilde{x}}(f) = 2(P_{uu}(f) + jP_{uv}(f)). \quad (15.36)$$

The reader at this point may recognize that the same relationships (15.33) and (15.34) occur for a process and its Hilbert transform or the real and imaginary parts of a

bandpass noise process. Thus, the complex envelope or the analytic signal of a WSS Gaussian random process is a *complex WSS Gaussian random process* (see Problem 15.12). Since these results are standard fare in many textbooks [Papoulis 1965], we only briefly illustrate them. We consider a bandpass noise process using the complex envelope representation. What we will show is that by concatenating the in-phase and quadrature processes together into a single complex envelope process, a complex WSS Gaussian random process results if the bandpass process is a real WSS Gaussian process.

Example 15.5 - Bandpass Gaussian Noise Process

Consider a real continuous bandpass WSS Gaussian process $x(t)$ with zero mean. It can be represented by its complex envelope as (using an argument similar to that in Section 15.3 but replacing Fourier transforms by PSDs)

$$x(t) = 2\text{Re}[\tilde{x}(t) \exp(j2\pi F_0 t)] \quad (15.37)$$

where $\tilde{x}(t) = u(t) + jv(t)$ and

$$\begin{aligned} u(t) &= [x(t) \cos 2\pi F_0 t]_{\text{LPF}} \\ v(t) &= [x(t) \sin 2\pi F_0 t]_{\text{LPF}}. \end{aligned}$$

Then, as shown in Section 15.3, after sampling we have the discrete process

$$\tilde{x}[n] = u[n] + jv[n]$$

where $u[n] = u(n\Delta)$, $v[n] = v(n\Delta)$, and $\Delta = 1/B$. But $u(t), v(t)$ are jointly Gaussian WSS processes because they are obtained by linear transformations of $x(t)$. Hence, $u[n], v[n]$ are jointly Gaussian also. Furthermore, from (15.37)

$$x(t) = 2u(t) \cos 2\pi F_0 t - 2v(t) \sin 2\pi F_0 t$$

so that because $E(x(t)) = 0$ for all t , it follows that $E(u(t)) = E(v(t)) = 0$ or $E(u[n]) = E(v[n]) = 0$ for all n . To show that $u[n], v[n]$ are jointly WSS we determine the ACF and the CCF. The ACF of $x(t)$ is

$$\begin{aligned} r_{xx}(\tau) &= E[x(t)x(t+\tau)] \\ &= 4E[(u(t) \cos 2\pi F_0 t - v(t) \sin 2\pi F_0 t) \\ &\quad \cdot (u(t+\tau) \cos 2\pi F_0(t+\tau) - v(t+\tau) \sin 2\pi F_0(t+\tau))] \\ &= 4[(r_{uu}(\tau) \cos 2\pi F_0 t \cos 2\pi F_0(t+\tau) + r_{vv}(\tau) \sin 2\pi F_0 t \sin 2\pi F_0(t+\tau)) \\ &\quad - (r_{uv}(\tau) \cos 2\pi F_0 t \sin 2\pi F_0(t+\tau) + r_{vu}(\tau) \sin 2\pi F_0 t \cos 2\pi F_0(t+\tau))] \\ &= 2[(r_{uu}(\tau) + r_{vv}(\tau)) \cos 2\pi F_0 \tau + (r_{uu}(\tau) - r_{vv}(\tau)) \cos 2\pi F_0(2t + \tau) \\ &\quad - (r_{uv}(\tau) - r_{vu}(\tau)) \sin 2\pi F_0 \tau - (r_{uv}(\tau) + r_{vu}(\tau)) \sin 2\pi F_0(2t + \tau)]. \end{aligned}$$

Since this must not depend on t ($x(t)$ was assumed to be WSS), we must have

$$\begin{aligned} r_{uu}(\tau) &= r_{vv}(\tau) \\ r_{uv}(\tau) &= -r_{vu}(\tau). \end{aligned} \quad (15.38)$$




Figure 15.3 Bandpass “white” Gaussian noise

Then

$$r_{xx}(\tau) = 4(r_{uu}(\tau) \cos 2\pi F_0 \tau + r_{vu}(\tau) \sin 2\pi F_0 \tau) \quad (15.39)$$

and since

$$\begin{aligned} r_{x_1 x_2}[k] &= E(x_1[n]x_2[n+k]) \\ &= E(x_1(n\Delta)x_2(n\Delta + k\Delta)) \\ &= r_{x_1 x_2}(k\Delta) \end{aligned}$$

we have finally

$$\begin{aligned} r_{uu}[k] &= r_{vv}[k] \\ r_{uv}[k] &= -r_{vu}[k] \end{aligned}$$

as required. Thus, $\tilde{x}[n]$ is a complex Gaussian WSS random process.

As an example, if $x(t)$ is the result of filtering a continuous-time WGN process with a PSD of $N_0/2$ over a band of B Hz as shown in Figure 15.2, then

$$r_{xx}(\tau) = N_0 B \frac{\sin \pi B \tau}{\pi B \tau} \cos 2\pi F_0 \tau.$$

By comparison with (15.39)

$$\begin{aligned} r_{uu}(\tau) &= \frac{N_0 B \sin \pi B \tau}{4 \pi B \tau} \\ r_{vu}(\tau) &= 0. \end{aligned}$$

If we sample the complex envelope $\tilde{x}(t) = u(t) + jv(t)$ at the Nyquist rate of $F_s = 1/\Delta = B$, we have

$$\begin{aligned} r_{uu}(k\Delta) &= \frac{N_0 B \sin \pi k}{4 \pi k} \\ &= \frac{N_0 B}{4} \delta[k]. \end{aligned}$$

Hence,

$$\begin{aligned} r_{uu}[k] &= \frac{N_0 B}{4} \delta[k] \\ r_{vu}[k] &= 0 \end{aligned}$$

and finally, $\tilde{x}[n]$ is *complex white Gaussian noise* (CWGN) with (see (15.35) and (15.36))

$$\begin{aligned} r_{\tilde{x}\tilde{x}}[k] &= \frac{N_0 B}{2} \delta[k] \\ P_{\tilde{x}\tilde{x}}(f) &= \frac{N_0 B}{2}. \end{aligned}$$

Note that because the samples $\tilde{x}[n]$ are complex Gaussian and *uncorrelated*, they are also independent. In summary, for CWGN we have

$$\tilde{x}[n] \sim \mathcal{CN}(0, \sigma^2)$$

where all the samples of the process are independent. $\diamond$

Before concluding this section let us point out that a real Gaussian random process is not a special case of a complex Gaussian random process. As a counterexample, consider the second-order moment $E(\tilde{x}[m]\tilde{x}[n])$, which is known to be zero for all m and n . If $\tilde{x}[n]$ is a real Gaussian random process, then the second moment need not be identically zero.

15.6 Derivatives, Gradients, and Optimization

In order to determine estimators such as the LSE or the MLE we must optimize a function over a complex parameter space. We explored one such problem in Example 15.2. We now extend that approach so that a general method can be used for easily finding estimators. The complex derivative of a *real* scalar function J with respect to a complex parameter θ was defined to be

$$\frac{\partial J}{\partial \theta} = \frac{1}{2} \left(\frac{\partial J}{\partial \alpha} - j \frac{\partial J}{\partial \beta} \right) \quad (15.40)$$

where $\theta = \alpha + j\beta$ with α, β the real and imaginary parts of θ , respectively. In making this definition we note that by the properties of complex numbers

$$\frac{\partial J}{\partial \theta} = 0 \quad \text{if and only if} \quad \frac{\partial J}{\partial \alpha} = \frac{\partial J}{\partial \beta} = 0.$$

Hence, in optimizing J over all α and β we can, equivalently, find the θ that satisfies $\partial J / \partial \theta = 0$. A typical real scalar function J is the LS error criterion in Example 15.2. The reader should note that other definitions of the complex derivative are possible, as for example [Monzingo and Miller 1980]

$$\frac{\partial J}{\partial \theta} = \frac{\partial J}{\partial \alpha} + j \frac{\partial J}{\partial \beta}$$

and would accomplish the same goal. Our choice results in differentiation formulas that are analogous to some commonly encountered in real calculus but unfortunately others

that are somewhat strange, as we will see shortly (see (15.46)). The reason for having to define the complex derivative is that a real function of complex variables is not an analytic function. As an example, if $J = |\theta|^2$, then

$$J = \alpha^2 + \beta^2$$

and using standard real calculus techniques and our definition of the complex derivative

$$\begin{aligned}\frac{\partial J}{\partial \theta} &= \frac{1}{2} \left(\frac{\partial J}{\partial \alpha} - j \frac{\partial J}{\partial \beta} \right) \\ &= \frac{1}{2} (2\alpha - j2\beta) \\ &= \theta^*.\end{aligned}$$

But it is well known that the complex function (actually a real function) $|\theta|^2$ is not an analytic function and so cannot be differentiated. Yet by our definition a stationary point of the real function can be found by solving $\partial J / \partial \theta = 0$, producing $\theta = 0$ or $\alpha = \beta = 0$ as expected. Note that if we were to rewrite J as $\theta\theta^*$, then the same result would have been obtained by considering θ^* to be a constant in the partial differentiation. Specifically, consider J to be a function of the two independent complex variables θ and θ^* , so that we denote it by $J(\theta, \theta^*)$. It is easily shown for this example that J is an analytic function with respect to θ , holding θ^* constant, and also with respect to θ^* , holding θ constant. Hence, applying the chain rule, we have

$$\frac{\partial J(\theta, \theta^*)}{\partial \theta} = \frac{\partial J}{\partial \theta} \theta^* + \theta \frac{\partial \theta^*}{\partial \theta}.$$

To evaluate $\partial \theta / \partial \theta$ we use the same definition of $\partial / \partial \theta$ to yield

$$\begin{aligned}\frac{\partial \theta}{\partial \theta} &= \frac{1}{2} \left(\frac{\partial}{\partial \alpha} - j \frac{\partial}{\partial \beta} \right) (\alpha + j\beta) \\ &= \frac{1}{2} \left(\frac{\partial \alpha}{\partial \alpha} + j \frac{\partial \beta}{\partial \alpha} - j \frac{\partial \alpha}{\partial \beta} + \frac{\partial \beta}{\partial \beta} \right) \\ &= \frac{1}{2} (1 + j0 - j0 + 1)\end{aligned}$$

and finally

$$\frac{\partial \theta}{\partial \theta} = 1 \tag{15.41}$$

which is consistent with known differentiation results for analytic functions. Similarly, it is easily shown that (see Problem 15.13)

$$\frac{\partial \theta^*}{\partial \theta} = 0 \tag{15.42}$$

which again is also consistent, assuming that θ and θ^* are independent complex variables. Finally,

$$\frac{\partial J(\theta, \theta^*)}{\partial \theta} = 1 \cdot \theta^* + \theta \cdot 0 = \theta^*$$

consistent with our earlier results. In summary, with the definition of the complex derivative given by (15.40) stationary points of $J(\theta, \theta^*)$ can be found as follows. We consider $J(\theta, \theta^*)$ as a function of the two independent complex variables θ and θ^* . We then evaluate the partial derivative of J with respect to θ , where θ^* is a constant, by applying the usual differentiation rules for analytic functions. Setting $\partial J/\partial\theta = 0$ produces the stationary points. Equivalently, we could have considered $\partial J/\partial\theta^*$ in obtaining a stationary point. In this case θ is held constant, so that

$$\begin{aligned}\frac{\partial J}{\partial\theta^*} &= \frac{\partial J}{\partial(\alpha + j(-\beta))} \\ &= \frac{1}{2} \left(\frac{\partial J}{\partial\alpha} - j \frac{\partial J}{\partial(-\beta)} \right) \\ &= \frac{1}{2} \left(\frac{\partial J}{\partial\alpha} + j \frac{\partial J}{\partial\beta} \right) \\ &= \left(\frac{\partial J}{\partial\theta} \right)^*\end{aligned}\tag{15.43}$$

and setting $\partial J/\partial\theta^* = 0$ will produce the same solutions as $\partial J/\partial\theta = 0$.

The complex gradient of a real function J with respect to the complex vector parameter θ is defined as

$$\frac{\partial J}{\partial\theta} = \begin{bmatrix} \frac{\partial J}{\partial\theta_1} \\ \frac{\partial J}{\partial\theta_2} \\ \vdots \\ \frac{\partial J}{\partial\theta_p} \end{bmatrix}$$

where each element is defined by (15.40). Again note that the complex gradient is zero if and only if each element is zero or $\partial J/\partial\theta_i = 0$ for $i = 1, 2, \dots, p$, and hence if and only if $\partial J/\partial\alpha_i = \partial J/\partial\beta_i = 0$ for $i = 1, 2, \dots, p$. Stationary points of J may be found by setting the complex gradient equal to zero and solving.

For the most part we will be interested in differentiating linear and Hermitian forms such as $\theta^H \mathbf{b}$ and $\theta^H \mathbf{A} \theta$ (where $\mathbf{A}$ is Hermitian), respectively. We now derive the complex gradient of these functions. First, consider $l(\theta) = \mathbf{b}^H \theta$, where we note that l is complex. Then

$$l(\theta) = \sum_{i=1}^p b_i^* \theta_i$$

and using linearity as well as (15.41), we have

$$\frac{\partial l}{\partial\theta_k} = b_k^* \frac{\partial\theta_k}{\partial\theta_k} = b_k^*$$

so that

$$\frac{\partial \mathbf{b}^H \theta}{\partial\theta} = \mathbf{b}^*.\tag{15.44}$$

We let the reader show that (see Problem 15.14)

$$\frac{\partial \theta^H \mathbf{b}}{\partial \theta} = 0. \quad (15.45)$$

Next, consider the Hermitian form $J = \theta^H \mathbf{A} \theta$. Since $\mathbf{A}^H = \mathbf{A}$, J must be real, which follows from $J^H = \theta^H \mathbf{A}^H \theta = \theta^H \mathbf{A} \theta = J$. Now

$$J = \sum_{i=1}^p \sum_{j=1}^p \theta_i^* [\mathbf{A}]_{ij} \theta_j$$

and using (15.41) and (15.42)

$$\begin{aligned} \frac{\partial J}{\partial \theta_k} &= \sum_{i=1}^p \sum_{j=1}^p \left(\theta_i^* [\mathbf{A}]_{ij} \frac{\partial \theta_j}{\partial \theta_k} + \frac{\partial \theta_i^*}{\partial \theta_k} [\mathbf{A}]_{ij} \theta_j \right) \\ &= \sum_{i=1}^p \sum_{j=1}^p \theta_i^* [\mathbf{A}]_{ij} \delta_{jk} \\ &= \sum_{i=1}^p \theta_i^* [\mathbf{A}]_{ik} \\ &= \sum_{i=1}^p [\mathbf{A}^T]_{ki} \theta_i^* \end{aligned}$$

and therefore

$$\frac{\partial J}{\partial \theta} = \mathbf{A}^T \theta^* = (\mathbf{A} \theta)^*. \quad (15.46)$$

As alluded to earlier, this result appears somewhat strange since if θ were real, we would have had $\partial J / \partial \theta = 2 \mathbf{A} \theta$ (see (4.3)). The real case then is *not* a special case.

Because we will wish to differentiate the likelihood function corresponding to the complex Gaussian PDF to determine the CRLB as well as the MLE, the following formulas are quite useful. If the covariance matrix $\mathbf{C}_{\tilde{x}}$ depends upon a number of *real* parameters $\{\xi_1, \xi_2, \dots, \xi_p\}$, then denoting the covariance matrix as $\mathbf{C}_{\tilde{x}}(\xi)$, it can be shown that

$$\frac{\partial \ln \det(\mathbf{C}_{\tilde{x}}(\xi))}{\partial \xi_i} = \text{tr} \left(\mathbf{C}_{\tilde{x}}^{-1}(\xi) \frac{\partial \mathbf{C}_{\tilde{x}}(\xi)}{\partial \xi_i} \right) \quad (15.47)$$

$$\frac{\partial \tilde{x}^H \mathbf{C}_{\tilde{x}}^{-1}(\xi) \tilde{x}}{\partial \xi_i} = -\tilde{x}^H \mathbf{C}_{\tilde{x}}^{-1}(\xi) \frac{\partial \mathbf{C}_{\tilde{x}}(\xi)}{\partial \xi_i} \mathbf{C}_{\tilde{x}}^{-1}(\xi) \tilde{x}. \quad (15.48)$$

These are derived in Appendix 3C for a real covariance matrix. For a complex covariance matrix the derivations are easily extended. The definitions and formulas are summarized in Table 15.1. We now illustrate their use with an example.

TABLE 15.1 Useful Formulas

Definitions
θ : Complex scalar parameter ($\theta = \alpha + j\beta$) J : Real scalar function of θ $\frac{\partial J}{\partial \theta} = \frac{1}{2} \left(\frac{\partial J}{\partial \alpha} - j \frac{\partial J}{\partial \beta} \right)$ θ : Complex vector parameter $\frac{\partial J}{\partial \theta} = \begin{bmatrix} \frac{\partial J}{\partial \theta_1} \\ \vdots \\ \frac{\partial J}{\partial \theta_p} \end{bmatrix}$
Formulas
$\frac{\partial \theta}{\partial \theta} = 1 \quad \frac{\partial \theta^*}{\partial \theta} = 0$ $\frac{\partial \mathbf{b}^H \theta}{\partial \theta} = \mathbf{b}^* \quad \frac{\partial \theta^H \mathbf{b}}{\partial \theta} = 0$ $\frac{\partial \theta^H \mathbf{A} \theta}{\partial \theta} = (\mathbf{A} \theta)^*$, where $\mathbf{A}^H = \mathbf{A}$ $\frac{\partial \ln \det(\mathbf{C}_x(\xi))}{\partial \xi_i} = \text{tr} \left(\mathbf{C}_x^{-1}(\xi) \frac{\partial \mathbf{C}_x(\xi)}{\partial \xi_i} \right)$, ξ_i real $\frac{\partial \bar{\mathbf{x}}^H \mathbf{C}_x^{-1}(\xi) \bar{\mathbf{x}}}{\partial \xi_i} = -\bar{\mathbf{x}}^H \mathbf{C}_x^{-1}(\xi) \frac{\partial \mathbf{C}_x(\xi)}{\partial \xi_i} \mathbf{C}_x^{-1}(\xi) \bar{\mathbf{x}}$

Example 15.6 - Minimization of Hermitian Functions

Suppose we wish to minimize the LS error

$$J = (\tilde{\mathbf{x}} - \mathbf{H}\theta)^H \mathbf{C}^{-1} (\tilde{\mathbf{x}} - \mathbf{H}\theta)$$

where $\tilde{\mathbf{x}}$ is a complex $N \times 1$ vector, $\mathbf{H}$ is a complex $N \times p$ matrix with $N > p$ and full rank, $\mathbf{C}$ is a complex $N \times N$ covariance matrix, and θ is a complex $p \times 1$ parameter vector. We first note that J is a real function since $J^H = J$ (recall that $\mathbf{C}^H = \mathbf{C}$). To find the value of θ that minimizes J we expand the function and then make use of our formulas.

$$J = \tilde{\mathbf{x}}^H \mathbf{C}^{-1} \tilde{\mathbf{x}} - \tilde{\mathbf{x}}^H \mathbf{C}^{-1} \mathbf{H} \theta - \theta^H \mathbf{H}^H \mathbf{C}^{-1} \tilde{\mathbf{x}} + \theta^H \mathbf{H}^H \mathbf{C}^{-1} \mathbf{H} \theta$$

$$\begin{aligned} \frac{\partial J}{\partial \theta} &= \mathbf{0} - (\mathbf{H}^H \mathbf{C}^{-1} \tilde{\mathbf{x}})^* - \mathbf{0} + (\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H} \theta)^* \\ &= -[\mathbf{H}^H \mathbf{C}^{-1} (\tilde{\mathbf{x}} - \mathbf{H} \theta)]^* \end{aligned} \quad (15.49)$$

using (15.44), (15.45), and (15.46). We can find the minimizing solution by setting the complex gradient equal to zero, yielding the LS solution

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^H \mathbf{C}^{-1} \bar{\mathbf{x}}. \quad (15.50)$$

To confirm that $\hat{\boldsymbol{\theta}}$ produces the global minimum the reader should verify that

$$\begin{aligned} J &= (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^H \mathbf{H}^H \mathbf{C}^{-1} \mathbf{H} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) + (\bar{\mathbf{x}} - \mathbf{H} \hat{\boldsymbol{\theta}})^H \mathbf{C}^{-1} (\bar{\mathbf{x}} - \mathbf{H} \hat{\boldsymbol{\theta}}) \\ &\geq (\bar{\mathbf{x}} - \mathbf{H} \hat{\boldsymbol{\theta}})^H \mathbf{C}^{-1} (\bar{\mathbf{x}} - \mathbf{H} \hat{\boldsymbol{\theta}}) \end{aligned}$$

with equality if and only if $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$. As an example, using (15.50), the solution of the LS problem in Example 15.2 is easily found by letting $\boldsymbol{\theta} = \hat{\mathbf{A}}$, $\mathbf{H} = [\bar{s}[0] \bar{s}[1] \dots \bar{s}[N-1]]^T$, and $\mathbf{C} = \mathbf{I}$, so that from (15.50)

$$\begin{aligned} \hat{\mathbf{A}} &= (\mathbf{H}^H \mathbf{H})^{-1} \mathbf{H}^H \bar{\mathbf{x}} \\ &= \frac{\sum_{n=0}^{N-1} \bar{x}[n] \bar{s}^*[n]}{\sum_{n=0}^{N-1} |\bar{s}[n]|^2}. \end{aligned}$$

◇

A final useful result concerns the minimization of a Hermitian form subject to linear constraints. Consider the real function

$$g(\mathbf{a}) = \mathbf{a}^H \mathbf{W} \mathbf{a}$$

where $\mathbf{a}$ is a complex $n \times 1$ vector and $\mathbf{W}$ is a complex $n \times n$ positive definite (and Hermitian) matrix. To avoid the trivial solution $\mathbf{a} = \mathbf{0}$ we assume that $\mathbf{a}$ satisfies the linear constraint $\mathbf{B} \mathbf{a} = \mathbf{b}$, where $\mathbf{B}$ is a complex $r \times p$ matrix with $r < p$ and full rank and $\mathbf{b}$ is a complex $r \times 1$ vector. A solution to this constrained minimization problem requires the use of Lagrangian multipliers. To extend the use of real Lagrangian multipliers to the complex case we let $\mathbf{B} = \mathbf{B}_R + j\mathbf{B}_I$, $\mathbf{a} = \mathbf{a}_R + j\mathbf{a}_I$, and $\mathbf{b} = \mathbf{b}_R + j\mathbf{b}_I$. Then the constraint equation $\mathbf{B} \mathbf{a} = \mathbf{b}$ is equivalent to

$$(\mathbf{B}_R + j\mathbf{B}_I)(\mathbf{a}_R + j\mathbf{a}_I) = \mathbf{b}_R + j\mathbf{b}_I$$

or

$$\begin{aligned} \mathbf{B}_R \mathbf{a}_R - \mathbf{B}_I \mathbf{a}_I &= \mathbf{b}_R \\ \mathbf{B}_I \mathbf{a}_R + \mathbf{B}_R \mathbf{a}_I &= \mathbf{b}_I. \end{aligned}$$

Since there are two sets of r constraints, we can let the Lagrangian be

$$\begin{aligned} J(\mathbf{a}) &= \mathbf{a}^H \mathbf{W} \mathbf{a} + \boldsymbol{\lambda}_R^T (\mathbf{B}_R \mathbf{a}_R - \mathbf{B}_I \mathbf{a}_I - \mathbf{b}_R) \\ &\quad + \boldsymbol{\lambda}_I^T (\mathbf{B}_I \mathbf{a}_R + \mathbf{B}_R \mathbf{a}_I - \mathbf{b}_I) \\ &= \mathbf{a}^H \mathbf{W} \mathbf{a} + \boldsymbol{\lambda}_R^T \text{Re}(\mathbf{B} \mathbf{a} - \mathbf{b}) + \boldsymbol{\lambda}_I^T \text{Im}(\mathbf{B} \mathbf{a} - \mathbf{b}) \end{aligned}$$

where λ_R, λ_I are each real $r \times 1$ Lagrangian multiplier vectors. But letting $\lambda = \lambda_R + j\lambda_I$ be a complex $r \times 1$ Lagrangian multiplier vector, the constraint equation is simplified to yield

$$\begin{aligned} J(\mathbf{a}) &= \mathbf{a}^H \mathbf{W} \mathbf{a} + \operatorname{Re} [(\lambda_R - j\lambda_I)^T (\mathbf{B} \mathbf{a} - \mathbf{b})] \\ &= \mathbf{a}^H \mathbf{W} \mathbf{a} + \frac{1}{2} \lambda^H (\mathbf{B} \mathbf{a} - \mathbf{b}) + \frac{1}{2} \lambda^T (\mathbf{B}^* \mathbf{a}^* - \mathbf{b}^*). \end{aligned}$$

We can now perform the constrained minimization. Using (15.44), (15.45), and (15.46), we have

$$\frac{\partial J}{\partial \mathbf{a}} = (\mathbf{W} \mathbf{a})^* + \left(\frac{\mathbf{B}^H \lambda}{2} \right)^*$$

and setting it equal to zero produces

$$\mathbf{a}_{\text{opt}} = -\mathbf{W}^{-1} \mathbf{B}^H \frac{\lambda}{2}.$$

Imposing the constraint $\mathbf{B} \mathbf{a}_{\text{opt}} = \mathbf{b}$, we have

$$\mathbf{B} \mathbf{a}_{\text{opt}} = -\mathbf{B} \mathbf{W}^{-1} \mathbf{B}^H \frac{\lambda}{2}.$$

Since $\mathbf{B}$ is assumed to be full rank and $\mathbf{W}$ to be positive definite, $\mathbf{B} \mathbf{W}^{-1} \mathbf{B}^H$ is invertible and

$$\frac{\lambda}{2} = -(\mathbf{B} \mathbf{W}^{-1} \mathbf{B}^H)^{-1} \mathbf{b}$$

so that finally

$$\mathbf{a}_{\text{opt}} = \mathbf{W}^{-1} \mathbf{B}^H (\mathbf{B} \mathbf{W}^{-1} \mathbf{B}^H)^{-1} \mathbf{b}. \quad (15.51)$$

To show that this is indeed the global minimizing solution we need only verify the identity

$$\begin{aligned} \mathbf{a}^H \mathbf{W} \mathbf{a} &= (\mathbf{a} - \mathbf{a}_{\text{opt}})^H \mathbf{W} (\mathbf{a} - \mathbf{a}_{\text{opt}}) + \mathbf{a}_{\text{opt}}^H \mathbf{W} \mathbf{a}_{\text{opt}} \\ &\geq \mathbf{a}_{\text{opt}}^H \mathbf{W} \mathbf{a}_{\text{opt}}, \end{aligned}$$

which holds with equality if and only if $\mathbf{a} = \mathbf{a}_{\text{opt}}$. We illustrate this constrained minimization with an example.

Example 15.7 - BLUE of Mean of Complex Colored Random Process

Assume that we observe

$$\tilde{x}[n] = \tilde{A} + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A}$ is a complex parameter to be estimated and $\tilde{w}[n]$ is complex noise with zero mean and covariance matrix $\mathbf{C}$. The complex equivalent of the BLUE is

$$\hat{\tilde{A}} = \mathbf{a}^H \tilde{\mathbf{x}}$$

where $\hat{A}$ is an unbiased estimator and has minimum variance among all linear estimators. To be unbiased we require

$$E(\hat{A}) = \mathbf{a}^H E(\tilde{\mathbf{x}}) = \mathbf{a}^H \tilde{A} \mathbf{1} = \tilde{A}$$

or $\mathbf{a}^H \mathbf{1} = 1$. This constraint can also be written as $\mathbf{1}^T \mathbf{a} = 1$, which can be verified by replacing $\mathbf{a}$ by its real and imaginary parts. The variance of $\hat{A}$ is

$$\begin{aligned} \text{var}(\hat{A}) &= E \left[|\hat{A} - E(\hat{A})|^2 \right] \\ &= E \left[|\mathbf{a}^H \tilde{\mathbf{x}} - \mathbf{a}^H \tilde{A} \mathbf{1}|^2 \right] \\ &= E \left[\mathbf{a}^H (\tilde{\mathbf{x}} - \tilde{A} \mathbf{1}) (\tilde{\mathbf{x}} - \tilde{A} \mathbf{1})^H \mathbf{a} \right] \\ &= \mathbf{a}^H \mathbf{C} \mathbf{a}. \end{aligned}$$

We now wish to minimize the variance subject to the linear constraint on $\mathbf{a}$. Letting $\mathbf{B} = \mathbf{1}^T$, $\mathbf{b} = 1$, and $\mathbf{W} = \mathbf{C}$, we have from (15.51) that

$$\begin{aligned} \mathbf{a}_{\text{opt}} &= \mathbf{C}^{-1} \mathbf{1} (\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1})^{-1} \\ &= \frac{\mathbf{C}^{-1} \mathbf{1}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}. \end{aligned}$$

The BLUE of $\tilde{A}$ is then

$$\begin{aligned} \hat{A} &= \mathbf{a}_{\text{opt}}^H \tilde{\mathbf{x}} \\ &= \frac{\mathbf{1}^T \mathbf{C}^{-1} \tilde{\mathbf{x}}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}} \end{aligned}$$

which is identical in form to the real case (see Example 6.2). A subtle difference, however, is that $\hat{A}$ minimizes $\text{var}(\hat{A})$, or if $\hat{A} = \hat{A}_R + j\hat{A}_I$,

$$\begin{aligned} \text{var}(\hat{A}) &= E \left(|\hat{A} - E(\hat{A})|^2 \right) \\ &= E \left(|(\hat{A}_R - E(\hat{A}_R)) + j(\hat{A}_I - E(\hat{A}_I))|^2 \right) \\ &= \text{var}(\hat{A}_R) + \text{var}(\hat{A}_I) \end{aligned}$$

which is actually the sum of the variances for each component of the estimator. $\diamond$

15.7 Classical Estimation with Complex Data

We will restrict our discussion to complex data vectors that have the complex Gaussian PDF or

$$p(\tilde{\mathbf{x}}; \boldsymbol{\theta}) = \frac{1}{\pi^N \det(\mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\theta}))} \exp \left[-(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\theta}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\theta}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\theta})) \right].$$

The parameter vector θ is to be estimated based on the complex data $\tilde{\mathbf{x}} = [\tilde{x}[0] \tilde{x}[1] \dots \tilde{x}[N-1]]^T$ and may have components that are real and/or complex. Many of the results encountered for real data/real parameter estimation theory can be extended to the complex case. A full account, however, would require another book—hence we will mention the results we believe to be most useful. The reader will undoubtedly wish to extend other real estimators and is certainly encouraged to do so.

We first consider the MVU estimator of θ . Because θ may have complex and real components, the most general approach assumes that θ is a real vector. To avoid confusion with the case when θ is purely complex we will denote the vector of real parameters as ξ . For example, if we wish to estimate the complex amplitude $\tilde{A}$ and the frequency f_0 of a sinusoid, we let $\xi = [A_R A_I f_0]^T$. Then, the MVU estimator has its usual meaning—it is the unbiased estimator having the minimum variance or

$$\begin{aligned} E(\hat{\xi}_i) &= \xi_i & i &= 1, 2, \dots, p \\ \text{var}(\hat{\xi}_i) &\text{ is minimum} & i &= 1, 2, \dots, p. \end{aligned}$$

A good starting point is the CRLB, which for a complex Gaussian PDF results in a Fisher information matrix (see Appendix 15C)

$$\begin{aligned} [\mathbf{I}(\xi)]_{ij} &= \text{tr} \left[\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\xi) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\xi)}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\xi) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\xi)}{\partial \xi_j} \right] \\ &\quad + 2\text{Re} \left[\frac{\partial \tilde{\boldsymbol{\mu}}^H(\xi)}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\xi) \frac{\partial \tilde{\boldsymbol{\mu}}(\xi)}{\partial \xi_j} \right] \end{aligned} \quad (15.52)$$

for $i, j = 1, 2, \dots, p$. The derivatives are defined, as usual, as the matrix or vector of partial derivatives (see Section 3.9), where we differentiate each complex element $g = g_R + jg_I$ as

$$\frac{\partial g}{\partial \xi_i} = \frac{\partial g_R}{\partial \xi_i} + j \frac{\partial g_I}{\partial \xi_i}.$$

The equality condition for the CRLB to be attained is, from (3.25),

$$\frac{\partial \ln p(\tilde{\mathbf{x}}; \xi)}{\partial \xi} = \mathbf{I}(\xi)(g(\tilde{\mathbf{x}}) - \xi) \quad (15.53)$$

where $\hat{\xi} = g(\tilde{\mathbf{x}})$ is the efficient estimator of ξ . Note that $p(\tilde{\mathbf{x}}; \xi)$ is the same as $p(\tilde{\mathbf{x}}; \theta)$ since the former is the complex Gaussian PDF parameterized using real parameters. The example to follow illustrates the computation of the CRLB for the estimation of a real parameter based on complex data. This situation frequently occurs when the covariance matrix depends on an unknown real parameter.

Example 15.8 - Random Sinusoid in CWGN

Assume that

$$\tilde{x}[n] = \tilde{A} \exp(j2\pi f_0 n) + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{w}[n]$ is CWGN with variance σ^2 , $\tilde{A} \sim \mathcal{CN}(0, \sigma_A^2)$, and $\tilde{A}$ is independent of $\tilde{w}[n]$. In matrix form we have

$$\tilde{\mathbf{x}} = \tilde{\mathbf{A}}\mathbf{e} + \tilde{\mathbf{w}}$$

where $\mathbf{e} = [1 \exp(j2\pi f_0) \dots \exp(j2\pi f_0(N-1))]^T$. Since sums of independent complex Gaussian random variables are also complex Gaussian (see Section 15.4), we have

$$\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{\mathbf{x}}}).$$

To find the covariance matrix

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{x}}} &= E(\tilde{\mathbf{x}}\tilde{\mathbf{x}}^H) \\ &= E((\tilde{\mathbf{A}}\mathbf{e} + \tilde{\mathbf{w}})(\tilde{\mathbf{A}}\mathbf{e} + \tilde{\mathbf{w}})^H) \\ &= E(|\tilde{A}|^2)\mathbf{e}\mathbf{e}^H + \sigma^2\mathbf{I} \\ &= \sigma_A^2\mathbf{e}\mathbf{e}^H + \sigma^2\mathbf{I} \end{aligned}$$

since $\tilde{A}$ and $\tilde{\mathbf{w}}$ are independent. We now examine the problem of estimating σ_A^2 , the variance of the complex amplitude, assuming f_0 and σ^2 are known. The PDF is

$$p(\tilde{\mathbf{x}}; \sigma_A^2) = \frac{1}{\pi^N \det(\mathbf{C}_{\tilde{\mathbf{x}}}(\sigma_A^2))} \exp[-\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \tilde{\mathbf{x}}].$$

To see if the CRLB is attained, we compute

$$\frac{\partial \ln p(\tilde{\mathbf{x}}; \sigma_A^2)}{\partial \sigma_A^2} = -\frac{\partial \ln \det(\mathbf{C}_{\tilde{\mathbf{x}}}(\sigma_A^2))}{\partial \sigma_A^2} - \frac{\partial \tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \tilde{\mathbf{x}}}{\partial \sigma_A^2}.$$

Using formulas (15.47) and (15.48)

$$\begin{aligned} \frac{\partial \ln p(\tilde{\mathbf{x}}; \sigma_A^2)}{\partial \sigma_A^2} &= -\text{tr}\left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\sigma_A^2)}{\partial \sigma_A^2}\right) + \tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\sigma_A^2)}{\partial \sigma_A^2} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \tilde{\mathbf{x}} \\ &= -\text{tr}(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \mathbf{e}\mathbf{e}^H) + \tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \mathbf{e}\mathbf{e}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \tilde{\mathbf{x}} \\ &= -\mathbf{e}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \mathbf{e} + |\mathbf{e}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \tilde{\mathbf{x}}|^2. \end{aligned}$$

But

$$\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) = \frac{1}{\sigma^2} \mathbf{I} - \frac{1}{\sigma^4} \frac{\sigma_A^2 \mathbf{e}\mathbf{e}^H}{1 + \frac{N\sigma_A^2}{\sigma^2}}$$

so that

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\sigma_A^2) \mathbf{e} &= \frac{1}{\sigma^2} \mathbf{e} - \frac{N\sigma_A^2}{\sigma^4 + N\sigma_A^2\sigma^2} \mathbf{e} \\ &= \frac{\mathbf{e}}{N\sigma_A^2 + \sigma^2}. \end{aligned}$$

Hence

$$\begin{aligned}\frac{\partial \ln p(\tilde{\mathbf{x}}; \sigma_A^2)}{\partial \sigma_A^2} &= -\frac{N}{N\sigma_A^2 + \sigma^2} + \frac{|\tilde{\mathbf{x}}^H \mathbf{e}|^2}{(N\sigma_A^2 + \sigma^2)^2} \\ &= \left(\frac{N}{N\sigma_A^2 + \sigma^2} \right)^2 \left(\frac{|\tilde{\mathbf{x}}^H \mathbf{e}|^2}{N^2} - \frac{\sigma^2}{N} - \sigma_A^2 \right).\end{aligned}$$

The CRLB is satisfied with

$$\begin{aligned}\hat{\sigma}_A^2 &= \frac{|\tilde{\mathbf{x}}^H \mathbf{e}|^2}{N^2} - \frac{\sigma^2}{N} \\ &= \frac{1}{N^2} \left| \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_0 n) \right|^2 - \frac{\sigma^2}{N}\end{aligned}$$

which is an efficient estimator and has minimum variance

$$\text{var}(\hat{\sigma}_A^2) = \left(\sigma_A^2 + \frac{\sigma^2}{N} \right)^2.$$

Note that the variance does not decrease to zero with N , and hence $\hat{\sigma}_A^2$ is not consistent. The reader may wish to review the results of Problem 3.14 in explaining why this is so. $\diamond$

When the parameter to be estimated is complex, we can sometimes determine if an efficient estimator exists more easily than by expressing the parameter as a real vector parameter and using (15.53). We now examine this special case. First, consider the complex parameter $\boldsymbol{\theta}$, where $\boldsymbol{\theta} = \boldsymbol{\alpha} + j\boldsymbol{\beta}$. Then, concatenate the real and imaginary parts into the real parameter vector $\boldsymbol{\xi} = [\boldsymbol{\alpha}^T \boldsymbol{\beta}^T]^T$. If the Fisher information matrix for $\boldsymbol{\xi}$ takes the form

$$\mathbf{I}(\boldsymbol{\xi}) = 2 \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix}$$

then the equality condition of the CRLB may be simplified. From (15.53) we have

$$\begin{aligned}\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \boldsymbol{\xi}} &= \begin{bmatrix} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \boldsymbol{\alpha}} \\ \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \boldsymbol{\beta}} \end{bmatrix} \\ &= 2 \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha} \\ \hat{\boldsymbol{\beta}} - \boldsymbol{\beta} \end{bmatrix} \\ &= 2 \begin{bmatrix} \mathbf{E}(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) - \mathbf{F}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \\ \mathbf{F}(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) + \mathbf{E}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \end{bmatrix}\end{aligned}$$

and using the definition of the complex gradient

$$\begin{aligned}
 \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} &= \frac{1}{2} \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\alpha}} - j \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right] \\
 &= (\mathbf{E} - j\mathbf{F})(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) - j(\mathbf{E} - j\mathbf{F})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \\
 &= (\mathbf{E} - j\mathbf{F})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^* \\
 &= \mathbf{I}^*(\boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^*
 \end{aligned}$$

where we have defined the complex Fisher information matrix as $\mathbf{I}(\boldsymbol{\theta}) = \mathbf{E} + j\mathbf{F}$. Or finally, by using (15.43) this becomes

$$\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} = \mathbf{I}(\boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}). \quad (15.54)$$

The latter form is easier to evaluate, as is shown in the next example. To test whether the Fisher information matrix has the desired form we can partition the real Fisher information matrix as

$$\mathbf{I}(\boldsymbol{\xi}) = \begin{bmatrix} \mathbf{I}_{\alpha\alpha} & \mathbf{I}_{\alpha\beta} \\ \mathbf{I}_{\beta\alpha} & \mathbf{I}_{\beta\beta} \end{bmatrix}$$

and note that we must have

$$E \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*}^T \right] = \mathbf{0} \quad (15.55)$$

for then

$$\begin{aligned}
 \frac{1}{4} E \left[\left(\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\alpha}} + j \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\alpha}} + j \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right)^T \right] &= \mathbf{0} \\
 \frac{1}{4} (\mathbf{I}_{\alpha\alpha} - \mathbf{I}_{\beta\beta}) + \frac{1}{4} j (\mathbf{I}_{\alpha\beta} + \mathbf{I}_{\beta\alpha}) &= \mathbf{0}
 \end{aligned}$$

and thus

$$\begin{aligned}
 \mathbf{I}_{\alpha\alpha} &= \mathbf{I}_{\beta\beta} = 2\mathbf{E} \\
 \mathbf{I}_{\beta\alpha} &= -\mathbf{I}_{\alpha\beta} = 2\mathbf{F}.
 \end{aligned}$$

When (15.54) is true, $\mathbf{I}^{-1}(\boldsymbol{\theta})$ is the covariance matrix of the efficient estimator $\hat{\boldsymbol{\theta}}$. (The efficient estimator for $\boldsymbol{\theta}$ is defined as $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\alpha}} + j\hat{\boldsymbol{\beta}}$, where $[\hat{\boldsymbol{\alpha}}^T \hat{\boldsymbol{\beta}}^T]^T$ is the efficient estimator of $[\boldsymbol{\alpha}^T \boldsymbol{\beta}^T]^T$.) To verify this we have from (15.54)

$$\hat{\boldsymbol{\theta}} - \boldsymbol{\theta} = \mathbf{I}^{-1}(\boldsymbol{\theta}) \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*}$$

and thus

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \mathbf{I}^{-1}(\boldsymbol{\theta}) E \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*}^H \right] \mathbf{I}^{-1}(\boldsymbol{\theta})$$

and

$$\begin{aligned}
 E \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})^H}{\partial \boldsymbol{\theta}^*} \right] &= \frac{1}{4} E \left[\left(\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\alpha}} + j \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\alpha}} - j \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right)^T \right] \\
 &= \frac{1}{4} [\mathbf{I}_{\alpha\alpha} + \mathbf{I}_{\beta\beta} + j(\mathbf{I}_{\beta\alpha} - \mathbf{I}_{\alpha\beta})] \\
 &= \frac{1}{4} (4\mathbf{E} + j4\mathbf{F}) \\
 &= \mathbf{I}(\boldsymbol{\theta})
 \end{aligned}$$

and thus

$$\mathbf{C}_{\hat{\boldsymbol{\theta}}} = \mathbf{I}^{-1}(\boldsymbol{\theta}). \quad (15.56)$$

An important application of this result is the complex linear model.

Example 15.9 - Complex Classical Linear Model

Assume that we have the complex extension of the classical linear model or

$$\tilde{\mathbf{x}} = \mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}}$$

where $\mathbf{H}$ is a known complex $N \times p$ matrix with $N > p$ and full rank, $\boldsymbol{\theta}$ is a complex $p \times 1$ parameter vector to be estimated, and $\tilde{\mathbf{w}}$ is a complex $N \times 1$ random vector with PDF $\tilde{\mathbf{w}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C})$. Then, by the properties of the complex Gaussian PDF

$$\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{H}\boldsymbol{\theta}, \mathbf{C})$$

so that $\tilde{\boldsymbol{\mu}} = \mathbf{H}\boldsymbol{\theta}$ and $\mathbf{C}(\boldsymbol{\theta}) = \mathbf{C}$ (not dependent on $\boldsymbol{\theta}$). The PDF is

$$p(\tilde{\mathbf{x}}; \boldsymbol{\theta}) = \frac{1}{\pi^N \det(\mathbf{C})} \exp [-(\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\theta})^H \mathbf{C}^{-1} (\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\theta})].$$

To check the equality condition we have

$$\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} = - \frac{\partial (\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\theta})^H \mathbf{C}^{-1} (\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*}$$

and from (15.49) this becomes (recall (15.43))

$$\begin{aligned}
 \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}^*} &= \mathbf{H}^H \mathbf{C}^{-1} (\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\theta}) \\
 &= \mathbf{H}^H \mathbf{C}^{-1} \mathbf{H} [(\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^H \mathbf{C}^{-1} \tilde{\mathbf{x}} - \boldsymbol{\theta}].
 \end{aligned} \quad (15.57)$$

Thus, the equality condition is satisfied, and

$$\hat{\boldsymbol{\theta}} = (\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^H \mathbf{C}^{-1} \tilde{\mathbf{x}} \quad (15.58)$$

is an efficient estimator, and hence the MVU estimator of θ . Also, from (15.54), (15.56), and (15.57)

$$\mathbf{C}_\theta = (\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1} \quad (15.59)$$

is its covariance matrix. The covariance matrix of the real parameter vector ξ is

$$\begin{aligned} \mathbf{C}_\xi &= \mathbf{I}^{-1}(\xi) \\ &= \left(2 \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix} \right)^{-1} \\ &= \frac{1}{2} \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \end{aligned}$$

because of the special form of $\mathbf{I}(\xi)$ (see property 4 in Appendix 15A). But $(\mathbf{A} + j\mathbf{B})(\mathbf{E} + j\mathbf{F}) = \mathbf{I}$, so that

$$\begin{aligned} \mathbf{A} + j\mathbf{B} &= (\mathbf{E} + j\mathbf{F})^{-1} \\ &= \mathbf{I}^{-1}(\theta) \\ &= \mathbf{H}^H \mathbf{C}^{-1} \mathbf{H} \end{aligned}$$

and hence

$$\mathbf{C}_\xi = \frac{1}{2} \begin{bmatrix} \text{Re} [(\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1}] & -\text{Im} [(\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1}] \\ \text{Im} [(\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1}] & \text{Re} [(\mathbf{H}^H \mathbf{C}^{-1} \mathbf{H})^{-1}] \end{bmatrix}.$$

To check that the real Fisher information matrix has the special form we have from (15.55) and (15.57)

$$\begin{aligned} E \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \theta)}{\partial \theta^*} \frac{\partial \ln p(\tilde{\mathbf{x}}; \theta)}{\partial \theta^*}^T \right] &= \mathbf{H}^H \mathbf{C}^{-1} E [(\tilde{\mathbf{x}} - \mathbf{H}\theta)(\tilde{\mathbf{x}} - \mathbf{H}\theta)^T] \mathbf{C}^{-1T} \mathbf{H}^* \\ &= \mathbf{H}^H \mathbf{C}^{-1} E(\tilde{\mathbf{w}}\tilde{\mathbf{w}}^T) \mathbf{C}^{-1T} \mathbf{H}^*. \end{aligned}$$

But as shown in Section 15.4, $E(\tilde{\mathbf{w}}\tilde{\mathbf{w}}^T) = \mathbf{0}$ for $\tilde{\mathbf{w}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C})$. Hence, the condition is satisfied. Also, $\hat{\theta}$ as given by (15.58) is also the MLE, as can be verified by setting the gradient of the log-likelihood function (15.57) equal to zero. If $\mathbf{W} = \mathbf{C}^{-1}$ in the weighted LSE, then $\hat{\theta}$ is also the weighted LSE (see Example 15.6). If $\tilde{\mathbf{w}}$ is not a complex Gaussian random vector, then $\hat{\theta}$ can be shown to be the BLUE for this problem (see Problem 15.20). $\diamond$

When the CRLB is not satisfied, we can always resort to an MLE. In practice, the use of sufficient statistics for real or complex data appears to be of limited value. Hence, our approach in Chapter 5, although it should be tried, frequently does not produce a practical estimator. The MLE, on the other hand, being a “turn-the-crank” procedure, is usually of more value. The general MLE equations for real parameters of a complex Gaussian PDF can be expressed in closed form. For complex parameters there does not seem to be any general simplification of these equations. Letting

$$p(\tilde{\mathbf{x}}; \xi) = \frac{1}{\pi^N \det(\mathbf{C}_{\tilde{\mathbf{x}}}(\xi))} \exp \left[-(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\xi))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\xi) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\xi)) \right]$$

we differentiate using the formulas (15.47) and (15.48). As shown in Appendix 15C, this leads to

$$\begin{aligned} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \xi_i} &= -\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \\ &\quad + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \\ &\quad + 2 \text{Re} \left[(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} \right] \end{aligned} \quad (15.60)$$

which when set equal to zero can sometimes be used to find the MLE. An example follows.

Example 15.10 - Phase of a Complex Sinusoid in CWGN

Consider the data

$$\tilde{x}[n] = A \exp[j(2\pi f_0 n + \phi)] + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where the real amplitude A and frequency f_0 are known, the phase ϕ is to be estimated, and $\tilde{w}[n]$ is CWGN with variance σ^2 . Then, $\tilde{x}[n]$ is a complex Gaussian process (although not WSS due to the nonstationary mean), so that to find the MLE of ϕ we can use (15.60). But $\mathbf{C}_{\tilde{\mathbf{x}}} = \sigma^2 \mathbf{I}$ does not depend on ϕ , and

$$\tilde{\boldsymbol{\mu}}(\phi) = \begin{bmatrix} A \exp[j\phi] \\ A \exp[j(2\pi f_0 + \phi)] \\ \vdots \\ A \exp[j(2\pi f_0(N-1) + \phi)] \end{bmatrix}.$$

Differentiating, we have

$$\frac{\partial \tilde{\boldsymbol{\mu}}(\phi)}{\partial \phi} = \begin{bmatrix} jA \exp[j\phi] \\ jA \exp[j(2\pi f_0 + \phi)] \\ \vdots \\ jA \exp[j(2\pi f_0(N-1) + \phi)] \end{bmatrix}$$

and from (15.60)

$$\begin{aligned} \frac{\partial \ln p(\tilde{\mathbf{x}}; \phi)}{\partial \phi} &= \frac{2}{\sigma^2} \text{Re} \left[(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\phi))^H \frac{\partial \tilde{\boldsymbol{\mu}}(\phi)}{\partial \phi} \right] \\ &= \frac{2}{\sigma^2} \text{Re} \left[\sum_{n=0}^{N-1} (\tilde{x}^*[n] - A \exp[-j(2\pi f_0 n + \phi)]) jA \exp[j(2\pi f_0 n + \phi)] \right] \\ &= -\frac{2A}{\sigma^2} \text{Im} \left[\sum_{n=0}^{N-1} \tilde{x}^*[n] \exp[j(2\pi f_0 n + \phi)] - NA \right] \end{aligned}$$

$$= -\frac{2A}{\sigma^2} \operatorname{Im} \left[\exp(j\phi) \sum_{n=0}^{N-1} \tilde{x}^*[n] \exp(j2\pi f_0 n) - NA \right].$$

Setting this equal to zero, we have upon letting

$$X(f_0) = \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_0 n)$$

denote the Fourier transform of the data

$$\begin{aligned} \operatorname{Im}[\exp(j\phi)X^*(f_0)] &= 0 \\ \operatorname{Im}[(\cos\phi + j\sin\phi)(\operatorname{Re}(X(f_0)) - j\operatorname{Im}(X(f_0)))] &= 0 \end{aligned}$$

or

$$\sin\phi \operatorname{Re}(X(f_0)) = \cos\phi \operatorname{Im}(X(f_0)).$$

The MLE follows as

$$\hat{\phi} = \arctan \left[\frac{\operatorname{Im}(X(f_0))}{\operatorname{Re}(X(f_0))} \right]$$

which may be compared to the results in Example 7.6. $\diamond$

15.8 Bayesian Estimation

The usual Bayesian estimators, the MMSE and MAP estimators, extend naturally to the complex case. We will assume that $\boldsymbol{\theta}$ is a complex $p \times 1$ random vector. As we saw in Section 15.4, if $\tilde{\mathbf{x}}, \boldsymbol{\theta}$ are jointly complex Gaussian, then the posterior PDF $p(\boldsymbol{\theta}|\tilde{\mathbf{x}})$ will also be complex Gaussian. By “jointly complex Gaussian” we mean that if $\tilde{\mathbf{x}} = \mathbf{x}_R + j\mathbf{x}_I$ and $\tilde{\mathbf{y}} = \mathbf{y}_R + j\mathbf{y}_I$, then $[\mathbf{x}_R^T \mathbf{y}_R^T \mathbf{x}_I^T \mathbf{y}_I^T]^T$ has a real multivariate Gaussian PDF, and furthermore, if $\mathbf{u} = [\mathbf{x}_R^T \mathbf{y}_R^T]^T$ and $\mathbf{v} = [\mathbf{x}_I^T \mathbf{y}_I^T]^T$, then the real covariance matrix of $[\mathbf{u}^T \mathbf{v}^T]^T$ has the special form given in Theorem 15.1. The mean and covariance of the posterior PDF are (see (15.25) and (15.26))

$$E(\boldsymbol{\theta}|\tilde{\mathbf{x}}) = E(\boldsymbol{\theta}) + \mathbf{C}_{\theta\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}})) \quad (15.61)$$

$$\mathbf{C}_{\theta|\tilde{\mathbf{x}}} = \mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\theta} \quad (15.62)$$

where the covariance matrices are defined as

$$\begin{aligned} \begin{bmatrix} \mathbf{C}_{\theta\theta} & \mathbf{C}_{\theta\tilde{\mathbf{x}}} \\ \mathbf{C}_{\tilde{\mathbf{x}}\theta} & \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} \end{bmatrix} &= \begin{bmatrix} E[(\boldsymbol{\theta} - E(\boldsymbol{\theta}))(\boldsymbol{\theta} - E(\boldsymbol{\theta}))^H] & E[(\boldsymbol{\theta} - E(\boldsymbol{\theta}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H] \\ E[(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\boldsymbol{\theta} - E(\boldsymbol{\theta}))^H] & E[(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H] \end{bmatrix} \\ &= \begin{bmatrix} p \times p & p \times N \\ N \times p & N \times N \end{bmatrix}. \end{aligned}$$

Because of the form of the posterior PDF, the MAP and MMSE estimators are identical. It can further be shown that the MMSE estimator based on $\tilde{\mathbf{x}}$ is just $E(\theta_i|\tilde{\mathbf{x}})$ and that it minimizes the Bayesian MSE

$$\text{Bmse}(\hat{\theta}_i) = E \left[|\theta_i - \hat{\theta}_i|^2 \right]$$

where the expectation is with respect to $p(\tilde{\mathbf{x}}, \theta_i)$. Furthermore, if $\hat{\theta}_i = \hat{\alpha}_i + j\hat{\beta}_i$, then $\hat{\alpha}_i$ minimizes $E[(\alpha_i - \hat{\alpha}_i)^2]$ and $\hat{\beta}_i$ minimizes $E[(\beta_i - \hat{\beta}_i)^2]$ (see Problem 15.23).

An important example is the complex Bayesian linear model. It is defined as

$$\tilde{\mathbf{x}} = \mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}} \quad (15.63)$$

where $\mathbf{H}$ is a complex $N \times p$ matrix with $N \leq p$ possibly, $\boldsymbol{\theta}$ is a complex $p \times 1$ random vector with $\boldsymbol{\theta} \sim \mathcal{CN}(\boldsymbol{\mu}_\theta, \mathbf{C}_{\theta\theta})$, and $\tilde{\mathbf{w}}$ is a complex $N \times 1$ random vector with $\tilde{\mathbf{w}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\tilde{w}})$ and is independent of $\boldsymbol{\theta}$. The MMSE estimator of $\boldsymbol{\theta}$ is given by $E(\boldsymbol{\theta}|\tilde{\mathbf{x}})$, so that we need only find the first two moments of the PDF $p(\tilde{\mathbf{x}}, \boldsymbol{\theta})$, according to (15.61) and (15.62). The mean is

$$\begin{aligned} E(\tilde{\mathbf{x}}) &= \mathbf{H}E(\boldsymbol{\theta}) + E(\tilde{\mathbf{w}}) \\ &= \mathbf{H}\boldsymbol{\mu}_\theta \end{aligned}$$

and the covariances are

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} &= E[(\mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}} - \mathbf{H}\boldsymbol{\mu}_\theta)(\mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}} - \mathbf{H}\boldsymbol{\mu}_\theta)^H] \\ &= \mathbf{H}E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)^H]\mathbf{H}^H + E(\tilde{\mathbf{w}}\tilde{\mathbf{w}}^H) \\ &= \mathbf{H}\mathbf{C}_{\theta\theta}\mathbf{H}^H + \mathbf{C}_{\tilde{w}} \\ \mathbf{C}_{\theta\tilde{\mathbf{x}}} &= E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}} - \mathbf{H}\boldsymbol{\mu}_\theta)^H] \\ &= E[(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)(\mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \tilde{\mathbf{w}})^H] \\ &= \mathbf{C}_{\theta\theta}\mathbf{H}^H \end{aligned}$$

since $\boldsymbol{\theta}$ and $\tilde{\mathbf{w}}$ are independent and hence uncorrelated. Finally, we should verify that $\tilde{\mathbf{x}}, \boldsymbol{\theta}$ are jointly complex Gaussian. Because $\boldsymbol{\theta}, \tilde{\mathbf{w}}$ are independent, they are also jointly complex Gaussian. Due to the linear transformation

$$\begin{bmatrix} \tilde{\mathbf{x}} \\ \boldsymbol{\theta} \end{bmatrix} = \begin{bmatrix} \mathbf{H} & \mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta} \\ \tilde{\mathbf{w}} \end{bmatrix}$$

$\tilde{\mathbf{x}}, \boldsymbol{\theta}$ are therefore jointly complex Gaussian. Finally, we have as the MMSE estimator for the complex Bayesian linear model, from (15.61),

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\mu}_\theta + \mathbf{C}_{\theta\theta}\mathbf{H}^H(\mathbf{H}\mathbf{C}_{\theta\theta}\mathbf{H}^H + \mathbf{C}_{\tilde{w}})^{-1}(\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\mu}_\theta) \quad (15.64)$$

$$= \boldsymbol{\mu}_\theta + (\mathbf{C}_{\theta\theta}^{-1} + \mathbf{H}^H\mathbf{C}_{\tilde{w}}^{-1}\mathbf{H})^{-1}\mathbf{H}^H\mathbf{C}_{\tilde{w}}^{-1}(\tilde{\mathbf{x}} - \mathbf{H}\boldsymbol{\mu}_\theta) \quad (15.65)$$

and the minimum Bayesian MSE can be shown to be, from (15.62),

$$\text{Bmse}(\hat{\theta}_i) = \left[\mathbf{C}_{\theta\theta} - \mathbf{C}_{\theta\theta} \mathbf{H}^H (\mathbf{H} \mathbf{C}_{\theta\theta} \mathbf{H}^H + \mathbf{C}_{\tilde{w}})^{-1} \mathbf{H} \mathbf{C}_{\theta\theta} \right]_{ii} \quad (15.66)$$

$$= \left[(\mathbf{C}_{\theta\theta}^{-1} + \mathbf{H}^H \mathbf{C}_{\tilde{w}}^{-1} \mathbf{H})^{-1} \right]_{ii} \quad (15.67)$$

which is just the $[i, i]$ element of $\mathbf{C}_{\theta|\tilde{z}}$. The reader will note that the results are identical to those for the real case with the transposes replaced by conjugate transposes (see Section 11.6). An example follows.

Example 15.11 - Random Amplitude Signal in CWGN

Consider the data set

$$\tilde{x}[n] = \tilde{A}\tilde{s}[n] + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A} \sim \mathcal{CN}(0, \sigma_A^2)$, $\tilde{s}[n]$ is a known signal, and $\tilde{w}[n]$ is CWGN with variance σ^2 and is independent of $\tilde{A}$. This is a standard model for a slowly fluctuating point target [Van Trees 1971]. If $\tilde{s}[n] = \exp(j2\pi f_0 n)$, we have the case examined in Example 15.8 in which we estimated σ_A^2 . Now, however, suppose we wish to estimate the *realization* of $\tilde{A}$. Then, we have the complex Bayesian linear model with

$$\begin{aligned} \mathbf{H} &= \tilde{\mathbf{s}} = [\tilde{s}[0] \ \tilde{s}[1] \ \dots \ \tilde{s}[N-1]]^T \\ \boldsymbol{\theta} &= \tilde{A} \\ \boldsymbol{\mu}_\theta &= 0 \\ \mathbf{C}_{\theta\theta} &= \sigma_A^2 \\ \mathbf{C}_{\tilde{w}} &= \sigma^2 \mathbf{I}. \end{aligned}$$

Hence, from (15.65)

$$\begin{aligned} \hat{\tilde{A}} &= \left((\sigma_A^2)^{-1} + \tilde{\mathbf{s}}^H \frac{1}{\sigma^2} \mathbf{I} \tilde{\mathbf{s}} \right)^{-1} \tilde{\mathbf{s}}^H \frac{1}{\sigma^2} \mathbf{I} \tilde{\mathbf{x}} \\ &= \frac{\tilde{\mathbf{s}}^H \tilde{\mathbf{x}}}{\frac{\sigma_A^2}{\sigma_A^2} + \tilde{\mathbf{s}}^H \tilde{\mathbf{s}}}. \end{aligned}$$

The minimum Bayesian MSE is, from (15.67),

$$\text{Bmse}(\hat{\tilde{A}}) = \left((\sigma_A^2)^{-1} + \frac{\tilde{\mathbf{s}}^H \tilde{\mathbf{s}}}{\sigma^2} \right)^{-1}.$$

The results are analogous to the real case, and in fact, if $\tilde{\mathbf{s}} = \mathbf{1}$ (the signal is a DC level)

$$\begin{aligned} \hat{\tilde{A}} &= \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{\tilde{x}} \\ \text{Bmse}(\hat{\tilde{A}}) &= \frac{\frac{\sigma^2}{N} \sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \end{aligned}$$

where $\bar{x}$ is the sample mean of the $\tilde{x}[n]$'s. This is identical in form to the real case (see (10.31)). $\diamond$

15.9 Asymptotic Complex Gaussian PDF

Similar to the real case, if the data record is large and the data are a segment from a complex WSS Gaussian random process with zero mean, then the PDF simplifies. In essence the covariance matrix becomes an autocorrelation matrix which if large enough in dimension is expressible in terms of the PSD. As we now show, the log-PDF becomes approximately

$$\ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi}) = -N \ln \pi - N \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\ln P_{\tilde{x}\tilde{x}}(f) + \frac{I(f)}{P_{\tilde{x}\tilde{x}}(f)} \right] df \quad (15.68)$$

where $I(f)$ is the periodogram or

$$I(f) = \frac{1}{N} \left| \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f n) \right|^2.$$

This approximate form of the PDF then leads to the asymptotic CRLB, which is computed from the elements of the Fisher information matrix

$$[\mathbf{I}(\boldsymbol{\xi})]_{ij} = N \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{\partial \ln P_{\tilde{x}\tilde{x}}(f; \boldsymbol{\xi})}{\partial \xi_i} \frac{\partial \ln P_{\tilde{x}\tilde{x}}(f; \boldsymbol{\xi})}{\partial \xi_j} df \quad (15.69)$$

where the explicit dependence of the PDF on the unknown real parameters is shown. These expressions simplify many CRLB and MLE calculations. To derive (15.68) we could extend the derivation in Appendix 3D with obvious modifications for complex data. However, a simpler and more intuitive proof will be given, which depends on the properties of complex Gaussian PDFs. By doing so we will be able to illustrate some useful properties of Fourier transforms of complex Gaussian random processes. For a more rigorous proof and the conditions under which (15.68) applies, see [Dzhaparidze 1986]. In practice, the asymptotic results hold if the data record length N is much greater than the "correlation time" (see Appendix 3D for a further discussion).

Consider the complex data set $\{\tilde{x}[0], \tilde{x}[1], \dots, \tilde{x}[N-1]\}$, where $\tilde{x}[n]$ is a complex zero mean WSS Gaussian random process. The PSD is $P_{\tilde{x}\tilde{x}}(f)$ and need not be symmetric about $f = 0$. The discrete Fourier transform (DFT) of the data is

$$X(f_k) = \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_k n)$$

where $f_k = k/N$ for $k = 0, 1, \dots, N-1$. We will first find the PDF of the DFT coefficients $X(f_k)$ and then transform back to $\tilde{x}[n]$ to yield the PDF $p(\tilde{\mathbf{x}}; \boldsymbol{\xi})$. In matrix form the DFT can be expressed as

$$\mathbf{X} = \mathbf{E} \tilde{\mathbf{x}}$$

where $\mathbf{X} = [X(f_0) X(f_1) \dots X(f_{N-1})]^T$ and $[\mathbf{E}]_{kn} = \exp(-j2\pi f_k n)$. If we normalize $\mathbf{E}$ by letting $\mathbf{U} = \mathbf{E}/\sqrt{N}$, so that $\mathbf{X} = \sqrt{N}\mathbf{U}\tilde{\mathbf{x}}$, then we note that $\mathbf{U}$ is a unitary matrix or $\mathbf{U}^H = \mathbf{U}^{-1}$. Next, since $\mathbf{X}$ is the result of a linear transformation of $\tilde{\mathbf{x}}$, then $\mathbf{X}$ is also complex Gaussian and has zero mean and covariance matrix $N\mathbf{U}\mathbf{C}_{\tilde{\mathbf{x}}}\mathbf{U}^H$. As an example, if $\tilde{\mathbf{x}}$ is CWGN so that $\mathbf{C}_{\tilde{\mathbf{x}}} = \sigma^2\mathbf{I}$, then $\mathbf{X}$ will have the PDF

$$\mathbf{X} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_X)$$

where

$$\begin{aligned}\mathbf{C}_X &= N\mathbf{U}\sigma^2\mathbf{I}\mathbf{U}^H \\ &= N\sigma^2\mathbf{I}.\end{aligned}$$

The reader may wish to compare this to the results in Example 15.3 to determine the reason for the difference in the covariance matrices. In general, to find the covariance matrix of the DFT coefficients we consider the $[k, l]$ element of $\mathbf{C}_X$ or $E[X(f_k)X^*(f_l)]$. Because the DFT is periodic with period N , we can rewrite it as (for N even)

$$X(f_k) = \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \tilde{x}[n] \exp(-j2\pi f_k n).$$

Then,

$$\begin{aligned}E[X(f_k)X^*(f_l)] &= \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} E[\tilde{x}[m]\tilde{x}^*[n]] \exp[-j2\pi(f_k m - f_l n)] \\ &= \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} r_{\tilde{x}\tilde{x}}[m-n] \exp[-j2\pi(f_k m - f_l n)].\end{aligned}$$

Letting $i = m - n$, we have

$$E[X(f_k)X^*(f_l)] = \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} \left[\sum_{i=m-(\frac{N}{2}-1)}^{m+\frac{N}{2}} r_{\tilde{x}\tilde{x}}[i] \exp(-j2\pi f_l i) \right] \exp[-j2\pi(f_k - f_l)m].$$

But as $N \rightarrow \infty$, the term in brackets approaches the PSD, or for large N

$$E[X(f_k)X^*(f_l)] \approx \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} \exp[-j2\pi(f_k - f_l)m] P_{\tilde{x}\tilde{x}}(f_l)$$

and by the orthogonality of the complex exponentials we have

$$E[X(f_k)X^*(f_l)] \approx NP_{\tilde{x}\tilde{x}}(f_k)\delta_{kl}.$$

Thus we have asymptotically

$$E[X(f_k)X^*(f_l)] = \begin{cases} NP_{\bar{x}\bar{x}}(f_k) & k = l \\ 0 & \text{otherwise.} \end{cases} \quad (15.70)$$

The DFT coefficients are approximately uncorrelated, and since they are jointly complex Gaussian, they are also approximately independent. In summary, we have

$$X(f_k) \stackrel{a}{\sim} \mathcal{CN}(0, NP_{\bar{x}\bar{x}}(f_k)) \quad k = 0, 1, \dots, N-1 \quad (15.71)$$

where “ a ” denotes the asymptotic PDF and the $X(f_k)$ ’s are asymptotically independent. Recall that for CWGN these properties are exact, not requiring $N \rightarrow \infty$. This is due to the scaled identity matrix of $\bar{\mathbf{x}}$ as well as to the unitary transformation of the DFT. In general, however, $\mathbf{C}_X$ is only approximately diagonal due to the property that the columns of $\mathbf{U}$ or, equivalently, $\mathbf{E}$ are only approximately the eigenvectors of $\mathbf{C}_{\bar{x}}$, which is a Hermitian Toeplitz matrix [Grenander and Szego 1958]. In other words, $\mathbf{C}_X = \mathbf{U}(\mathbf{N}\mathbf{C}_{\bar{x}})\mathbf{U}^H$ is approximately an eigendecomposition of $\mathbf{N}\mathbf{C}_{\bar{x}}$. We can now write the asymptotic PDF of $\mathbf{X}$ as $\mathbf{X} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_X)$, where from (15.71) $\mathbf{C}_X = N \text{diag}(P_{\bar{x}\bar{x}}(f_0), P_{\bar{x}\bar{x}}(f_1), \dots, P_{\bar{x}\bar{x}}(f_{N-1}))$. To transform the PDF back to that of $\bar{\mathbf{x}}$ we note that $\mathbf{X} = \mathbf{E}\bar{\mathbf{x}}$, and since $\mathbf{U} = \mathbf{E}/\sqrt{N}$,

$$\begin{aligned} \bar{\mathbf{x}} &= \mathbf{E}^{-1}\mathbf{X} \\ &= \frac{1}{\sqrt{N}}\mathbf{U}^{-1}\mathbf{X} \\ &= \frac{1}{\sqrt{N}}\mathbf{U}^H\mathbf{X} \\ &= \frac{1}{N}\mathbf{E}^H\mathbf{X} \\ &\sim \mathcal{CN}(\mathbf{0}, \frac{1}{N^2}\mathbf{E}^H\mathbf{C}_X\mathbf{E}). \end{aligned}$$

In explicit form this becomes

$$\begin{aligned} \ln p(\bar{\mathbf{x}}; \boldsymbol{\xi}) &= -N \ln \pi - \ln \det \left(\frac{1}{N^2} \mathbf{E}^H \mathbf{C}_X \mathbf{E} \right) - \bar{\mathbf{x}}^H \left(\frac{1}{N^2} \mathbf{E}^H \mathbf{C}_X \mathbf{E} \right)^{-1} \bar{\mathbf{x}} \\ &= -N \ln \pi - \ln \left[\det \left(\frac{1}{N} \mathbf{E}^H \mathbf{E} \right) \det \left(\frac{1}{N} \mathbf{C}_X \right) \right] - \bar{\mathbf{x}}^H \mathbf{E}^H \mathbf{C}_X^{-1} \mathbf{E} \bar{\mathbf{x}} \\ &= -N \ln \pi - \ln \det \left(\frac{1}{N} \mathbf{C}_X \right) - \mathbf{X}^H \mathbf{C}_X^{-1} \mathbf{X} \end{aligned}$$

since $\mathbf{E}^H \mathbf{E} / N = \mathbf{I}$ and $\mathbf{E}^{-1} = \mathbf{E}^H / N$. Hence,

$$\ln p(\bar{\mathbf{x}}; \boldsymbol{\xi}) = -N \ln \pi - \ln \prod_{k=0}^{N-1} P_{\bar{x}\bar{x}}(f_k) - \sum_{k=0}^{N-1} \frac{|X(f_k)|^2}{NP_{\bar{x}\bar{x}}(f_k)}.$$

Since N is assumed to be large, we can further approximate this as

$$\begin{aligned}
 \ln p(\tilde{\mathbf{x}}; \xi) &= -N \ln \pi - N \sum_{k=0}^{N-1} [\ln P_{\tilde{x}\tilde{x}}(f_k)] \frac{1}{N} - N \sum_{k=0}^{N-1} \frac{\frac{1}{N} |X(f_k)|^2}{P_{\tilde{x}\tilde{x}}(f_k)} \frac{1}{N} \\
 &\approx -N \ln \pi - N \int_0^1 \ln P_{\tilde{x}\tilde{x}}(f) df - N \int_0^1 \frac{\frac{1}{N} |X(f)|^2}{P_{\tilde{x}\tilde{x}}(f)} df \\
 &= -N \ln \pi - N \int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\ln P_{\tilde{x}\tilde{x}}(f) + \frac{I(f)}{P_{\tilde{x}\tilde{x}}(f)} \right] df.
 \end{aligned}$$

In the last step we have used the fact that $I(f)$ and $P_{\tilde{x}\tilde{x}}(f)$ are periodic with period 1 to change the integration limits. To find the Fisher information matrix we need only compute the second partial derivatives of $\ln p(\tilde{\mathbf{x}}; \xi)$, noting that

$$\begin{aligned}
 E(I(f)) &= E\left(\frac{1}{N} |X(f)|^2\right) \\
 &= \frac{1}{N} \text{var}(X(f)) \\
 &\approx P_{\tilde{x}\tilde{x}}(f)
 \end{aligned}$$

for large data records. An example of the utility of the asymptotic results follows.

Example 15.12 - Periodogram Spectral Estimation

The periodogram spectral estimator is defined as

$$\begin{aligned}
 \hat{P}_{\tilde{x}\tilde{x}}(f_k) &= I(f_k) \\
 &= \frac{1}{N} \left| \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_k n) \right|^2 \\
 &= \frac{1}{N} |X(f_k)|^2.
 \end{aligned}$$

To determine the quality of the estimator we compute its mean and variance using the results of (15.71). The mean has already been shown to be approximately $P_{\tilde{x}\tilde{x}}(f_k)$, so that it is asymptotically unbiased. The variance is asymptotically

$$\begin{aligned}
 \text{var}(\hat{P}_{\tilde{x}\tilde{x}}(f_k)) &= E[\hat{P}_{\tilde{x}\tilde{x}}^2(f_k)] - E^2[\hat{P}_{\tilde{x}\tilde{x}}(f_k)] \\
 &= E[\hat{P}_{\tilde{x}\tilde{x}}^2(f_k)] - P_{\tilde{x}\tilde{x}}^2(f_k).
 \end{aligned}$$

But

$$E[\hat{P}_{\tilde{x}\tilde{x}}^2(f_k)] = \frac{1}{N^2} E[X(f_k)X^*(f_k)X(f_k)X^*(f_k)]$$

$$\begin{aligned}
&= \frac{1}{N^2} E[X(f_k)X^*(f_k)] E[X(f_k)X^*(f_k)] \\
&\quad + \frac{1}{N^2} E[X(f_k)X^*(f_k)] E[X(f_k)X^*(f_k)] \\
&= \frac{2}{N^2} \text{var}^2(X(f_k)) \\
&= 2P_{\tilde{x}\tilde{x}}^2(f_k)
\end{aligned}$$

making use of the fourth-order moment properties of jointly complex Gaussian random variables. Thus, we have that asymptotically

$$\text{var}(\hat{P}_{\tilde{x}\tilde{x}}(f_k)) = P_{\tilde{x}\tilde{x}}^2(f_k)$$

and does not decrease with increasing N . Unfortunately, the periodogram is inconsistent, an example of which is shown in Figure 15.4 for CWGN with $\sigma^2 = 1$. As can be seen, the average value of the periodogram is the PSD value $P_{\tilde{x}\tilde{x}}(f) = \sigma^2 = 1$. However, the variance does not decrease as the data record length increases. The increasingly “ragged” nature of the spectral estimate results from the larger number of points in the PSD that we are estimating. As previously shown, these estimates are independent (since $X(f_k)$ is independent of $X(f_l)$ for $k \neq l$), causing the more rapid fluctuations with increasing N . $\diamond$

15.10 Signal Processing Examples



We now apply some of the previously developed theory to determine estimators based on complex data. First we revisit the problem of sinusoidal parameter estimation in Examples 3.14 and 7.16 to compare the use of real to complex data. Next, we show how an adaptive beamformer results from classical estimation theory.

Example 15.13 - Sinusoidal Parameter Estimation

Consider a data model for a complex sinusoid in CWGN or



$$\tilde{x}[n] = \tilde{A} \exp(j2\pi f_0 n) + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A}$ is a complex amplitude, f_0 is the frequency, and $\tilde{w}[n]$ is a CWGN process with variance σ^2 . We wish to estimate $\tilde{A}$ and f_0 . Since the linear model does not apply because of the unknown frequency, we resort to an MLE. To determine the asymptotic performance of the MLE we also compute the CRLB. Recall from Chapter 7 that at least at high SNRs and/or large data records the MLE is usually efficient. We observe that one parameter, $\tilde{A}$, is complex and one, f_0 , is real. The simplest way to find the CRLB is to let $\tilde{A} = A \exp(j\phi)$ and find it for the real parameter vector $\boldsymbol{\xi} = [A \ f_0 \ \phi]^T$.

(a) $N = 128$ (b) $N = 256$ **Figure 15.4** Inconsistency of periodogram

Then, from (15.52) we have

$$\begin{aligned}
 [\mathbf{I}(\boldsymbol{\xi})]_{ij} &= 2 \operatorname{Re} \left[\frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \frac{1}{\sigma^2} \mathbf{I} \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j} \right] \\
 &= \frac{2}{\sigma^2} \operatorname{Re} \left[\sum_{n=0}^{N-1} \frac{\partial \tilde{s}^*[n]}{\partial \xi_i} \frac{\partial \tilde{s}[n]}{\partial \xi_j} \right]
 \end{aligned}$$

(c) $N = 512$ (d) $N = 1024$

where $\tilde{s}[n] = \tilde{A} \exp(j2\pi f_0 n) = A \exp[j(2\pi f_0 n + \phi)]$. Now the partial derivatives are

$$\begin{aligned} \frac{\partial \tilde{s}[n]}{\partial A} &= \exp[j(2\pi f_0 n + \phi)] \\ \frac{\partial \tilde{s}[n]}{\partial f_0} &= j2\pi n A \exp[j(2\pi f_0 n + \phi)] \\ \frac{\partial \tilde{s}[n]}{\partial \phi} &= jA \exp[j(2\pi f_0 n + \phi)] \end{aligned}$$

so that

$$\mathbf{I}(\boldsymbol{\xi}) = \frac{2}{\sigma^2} \begin{bmatrix} N & 0 & 0 \\ 0 & A^2 \sum_{n=0}^{N-1} (2\pi n)^2 & A^2 \sum_{n=0}^{N-1} 2\pi n \\ 0 & A^2 \sum_{n=0}^{N-1} 2\pi n & NA^2 \end{bmatrix}.$$

Upon inversion and using (3.22) we have

$$\begin{aligned} \text{var}(\hat{A}) &\geq \frac{\sigma^2}{2N} \\ \text{var}(\hat{f}_0) &\geq \frac{6\sigma^2}{(2\pi)^2 A^2 N(N-1)} \\ \text{var}(\hat{\phi}) &\geq \frac{\sigma^2(2N-1)}{A^2 N(N+1)} \end{aligned}$$

and in terms of the SNR, which is $\eta = A^2/\sigma^2$, this reduces to

$$\begin{aligned} \text{var}(\hat{A}) &\geq \frac{\sigma^2}{2N} \\ \text{var}(\hat{f}_0) &\geq \frac{6}{(2\pi)^2 \eta N(N-1)} \\ \text{var}(\hat{\phi}) &\geq \frac{2N-1}{\eta N(N+1)}. \end{aligned} \quad (15.72)$$

The bounds for the complex case are one-half those for the real case for $\hat{f}_0$ and $\hat{\phi}$, and one-fourth that for $\hat{A}$. (The reader should be warned, however, about drawing further conclusions. We are actually comparing “apples to oranges” since the data models are different.) To find the MLE we must maximize

$$p(\tilde{\mathbf{x}}; \boldsymbol{\xi}) = \frac{1}{\pi^N \det(\sigma^2 \mathbf{I})} \exp \left[-\frac{1}{\sigma^2} \sum_{n=0}^{N-1} \left| \tilde{x}[n] - \tilde{A} \exp(j2\pi f_0 n) \right|^2 \right]$$

or, equivalently, we must minimize

$$J(\tilde{A}, f_0) = \sum_{n=0}^{N-1} \left| \tilde{x}[n] - \tilde{A} \exp(j2\pi f_0 n) \right|^2.$$

The minimization can be done with respect to A and ϕ or, equivalently, with respect to $\tilde{A}$. The latter approach is easier since for f_0 fixed this is just a complex linear LS problem. In matrix form we have

$$J(\tilde{A}, f_0) = (\tilde{\mathbf{x}} - \mathbf{e}\tilde{A})^H (\tilde{\mathbf{x}} - \mathbf{e}\tilde{A})$$

where $\tilde{\mathbf{x}} = [\tilde{x}[0] \tilde{x}[1] \dots \tilde{x}[N-1]]^T$ and $\mathbf{e} = [1 \exp(j2\pi f_0) \dots \exp(j2\pi f_0(N-1))]^T$. But we have already shown in (15.49) that

$$\frac{\partial J}{\partial \hat{A}} = -[\mathbf{e}^H(\tilde{\mathbf{x}} - \mathbf{e}\hat{A})]^*.$$

Setting this equal to zero and solving produces

$$\begin{aligned}\hat{A} &= \frac{\mathbf{e}^H \tilde{\mathbf{x}}}{\mathbf{e}^H \mathbf{e}} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_0 n).\end{aligned}$$

Substituting back into J

$$\begin{aligned}J(\hat{A}, f_0) &= \tilde{\mathbf{x}}^H (\tilde{\mathbf{x}} - \mathbf{e}\hat{A}) \\ &= \tilde{\mathbf{x}}^H \tilde{\mathbf{x}} - \frac{\tilde{\mathbf{x}}^H \mathbf{e} \mathbf{e}^H \tilde{\mathbf{x}}}{\mathbf{e}^H \mathbf{e}} \\ &= \tilde{\mathbf{x}}^H \tilde{\mathbf{x}} - \frac{|\mathbf{e}^H \tilde{\mathbf{x}}|^2}{\mathbf{e}^H \mathbf{e}}.\end{aligned}$$

To minimize J over f_0 we need to maximize

$$\frac{|\mathbf{e}^H \tilde{\mathbf{x}}|^2}{\mathbf{e}^H \mathbf{e}} = \frac{1}{N} \left| \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi f_0 n) \right|^2$$

which is recognized as the periodogram. The MLE of A and ϕ are found from

$$\hat{A} = |\hat{A}| = \left| \frac{1}{N} \sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi \hat{f}_0 n) \right| \quad (15.73)$$

$$\begin{aligned}\hat{\phi} &= \arctan \frac{\text{Im}(\hat{A})}{\text{Re}(\hat{A})} \\ &= \arctan \frac{\text{Im} \left(\sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi \hat{f}_0 n) \right)}{\text{Re} \left(\sum_{n=0}^{N-1} \tilde{x}[n] \exp(-j2\pi \hat{f}_0 n) \right)}.\end{aligned} \quad (15.74)$$

Note that the MLE is exact and does not assume N is large, as was required for the real case. The need for large N in the real case so that f_0 is not near 0 or 1/2 can be explained by the inability of the periodogram to resolve complex sinusoids closer than about $1/N$ in frequency. For a real sinusoid

$$\begin{aligned}s[n] &= A \cos(2\pi f_0 n + \phi) \\ &= \frac{A}{2} \exp[j(2\pi f_0 n + \phi)] + \frac{A}{2} \exp[-j(2\pi f_0 n + \phi)]\end{aligned}$$

the peak of the periodogram (even in the absence of noise) will shift from the true frequency if the frequency difference between the complex sinusoidal components or $|f_0 - (-f_0)| = 2f_0$ is not much greater than $1/N$. As an example, in Figure 15.5 we plot the periodogram for $A = 1$, $\phi = \pi/2$, $f_0 = 0.1, 0.05, 0.025$, and $N = 20$. No noise is present. The peak locations for $f_0 = 0.1$ and $f_0 = 0.05$ are slightly shifted towards $f = 0$. However, for $f = 0.025$ the peak location is at $f = 0$. This problem does not occur with a complex sinusoid since there is only *one* complex sinusoidal component and therefore no interaction. See also Problem 15.21 for the extension to two complex sinusoids. $\diamond$

Example 15.14 - Adaptive Beamforming

We continue the array processing problem in Example 3.15 but now assume that the noise at the sensors is nonwhite. It is desired to take into account the coloration of the noise in designing a “beamformer,” which is an estimator of the transmitted signal. In many cases the nonwhite noise is the result of an intentional jammer. Recall that for an emitted sinusoidal signal $A \cos(2\pi F_0 t + \phi)$ the received signals at a line array are (see Example 3.15)

$$s_n(t) = A \cos[2\pi f_s n + 2\pi F_0(t - t_0) + \phi] \quad n = 0, 1, \dots, M-1$$

where n is the sensor number, $f_s = F_0(d/c) \cos \beta$ (d is the sensor spacing, c is the propagation speed, and β is the arrival angle) is the spatial frequency, and t_0 is the propagation time to the zeroth sensor. If we choose to process the analytic signal of the received data, then letting $\phi' = -2\pi F_0 t_0 + \phi$ (ϕ and t_0 are generally unknown), we have for the signal component

$$\begin{aligned} \tilde{s}_n(t) &= A \exp(j\phi') \exp(j2\pi f_s n) \exp(j2\pi F_0 t) \\ &= \tilde{A} \exp(j2\pi f_s n) \exp(j2\pi F_0 t). \end{aligned}$$

At each instant of time we have a signal vector “snapshot” at the M sensors


$$\begin{aligned} \tilde{\mathbf{s}}(t) &= [\tilde{s}_0(t) \ \tilde{s}_1(t) \ \dots \ \tilde{s}_{M-1}(t)]^T \\ &= \tilde{A} \exp(j2\pi F_0 t) \mathbf{e} \end{aligned}$$

where $\mathbf{e} = [1 \ \exp(j2\pi f_s) \ \dots \ \exp(j2\pi f_s(M-1))]^T$. The data model is assumed to be

$$\tilde{\mathbf{x}}(t) = \tilde{A} \exp(j2\pi F_0 t) \mathbf{e} + \tilde{\mathbf{w}}(t)$$

where $\tilde{\mathbf{w}}(t)$ is a zero mean noise vector with a covariance matrix $\mathbf{C}$. The mean and covariance are assumed not to vary with t . We wish to design a beamformer that combines the sensor outputs linearly as

$$\begin{aligned} \tilde{y}(t) &= \sum_{n=0}^{M-1} a_n^* \tilde{x}_n(t) \\ &= \mathbf{a}^H \tilde{\mathbf{x}}(t) \end{aligned}$$

**Figure 15.5** Periodogram for single real sinusoid

so that the signal is passed undistorted but the noise at the output is minimized. By undistorted we mean that if $\tilde{\mathbf{x}}(t) = \tilde{\mathbf{s}}(t) = \tilde{A} \exp(j2\pi F_0 t) \mathbf{e}$, then at the beamformer output we should have

$$\tilde{y}(t) = \tilde{A} \exp(j2\pi F_0 t).$$

This requires that

$$\mathbf{a}^H \tilde{\mathbf{s}}(t) = \tilde{A} \exp(j2\pi F_0 t)$$

or

$$\mathbf{a}^H \mathbf{e} = 1$$

or finally the constraint

$$\mathbf{e}^H \mathbf{a} = 1.$$

The noise at the beamformer output has variance

$$\begin{aligned} E[|\mathbf{a}^H \tilde{\mathbf{w}}(t)|^2] &= E[\mathbf{a}^H \tilde{\mathbf{w}}(t) \tilde{\mathbf{w}}^H(t) \mathbf{a}] \\ &= \mathbf{a}^H \mathbf{C} \mathbf{a} \end{aligned}$$

which is also the *variance* of $\tilde{y}(t)$. Hence, the optimal beamformer weights are given as the solution to the following problem: Minimize $\mathbf{a}^H \mathbf{C} \mathbf{a}$ subject to the constraint $\mathbf{e}^H \mathbf{a} = 1$. As a result, this is sometimes called a *minimum variance distortionless response* (MVDR) beamformer [Owsley 1985]. The reader will recognize that we have already solved this problem in connection with the BLUE. In fact, from (15.51) with $\mathbf{W} = \mathbf{C}$, $\mathbf{B} = \mathbf{e}^H$, and $\mathbf{b} = 1$ we have the optimum solution

$$\mathbf{a}_{\text{opt}} = \frac{\mathbf{C}^{-1} \mathbf{e}}{\mathbf{e}^H \mathbf{C}^{-1} \mathbf{e}}.$$

The beamformer output is given as

$$\tilde{y}(t) = \frac{\mathbf{e}^H \mathbf{C}^{-1} \tilde{\mathbf{x}}(t)}{\mathbf{e}^H \mathbf{C}^{-1} \mathbf{e}}. \quad (15.75)$$

As can be verified, $\tilde{y}(t)$ is just the BLUE of $\tilde{A} \exp(j2\pi F_0 t)$ for a given t . In Example 15.7 we derived this result for $F_0 = 0$, and hence for $\mathbf{e} = 1$. The name “adaptive beamformer” comes from the practical implementation of (15.75) in which the covariance matrix of the noise is usually unknown and hence must be estimated before a signal is present. The beamformer is then said to “adapt” to the noise field present. Of course, when an estimated covariance matrix is used, there is no optimality associated with the beamformer. The performance may be poor if a signal is present when the covariance is estimated [Cox 1973]. Note that if $\mathbf{C} = \sigma^2 \mathbf{I}$, so that the noise at the sensors is uncorrelated and of equal variance (“spatially white”), then (15.75) reduces to

$$\tilde{y}(t) = \frac{1}{M} \sum_{n=0}^{M-1} \tilde{x}_n(t) \exp(-j2\pi f_s n).$$

The beamformer first phases the signals at the various sensors to align them in time, because of the varying propagation delays, and then averages them. This is the so-called *conventional beamformer* [Knight, Pridham, and Kay 1981].

To illustrate the effect of a nonwhite spatial noise field assume that in addition to white noise we have an interfering plane wave at the same temporal frequency but a different spatial frequency f_i or arrival angle. Then, the model for the received data is

$$\tilde{\mathbf{x}}(t) = \tilde{A} \exp(j2\pi F_0 t) \mathbf{e} + \tilde{B} \exp(j2\pi F_0 t) \mathbf{i} + \tilde{\mathbf{u}}(t) \quad (15.76)$$

where $\mathbf{i} = [1 \exp(j2\pi f_i) \dots \exp(j2\pi f_i(M-1))]^T$ and $\tilde{\mathbf{u}}(t)$ is a spatial white noise process or $\tilde{\mathbf{u}}(t)$ has zero mean and a covariance matrix $\sigma^2 \mathbf{I}$ for a fixed t . If we model the complex amplitude of the interference $\tilde{B}$ as a complex random variable with zero mean and covariance P and independent of $\tilde{\mathbf{u}}(t)$, then the interference plus noise can be represented as

$$\tilde{\mathbf{w}}(t) = \tilde{B} \exp(j2\pi F_0 t) \mathbf{i} + \tilde{\mathbf{u}}(t)$$

where $E(\tilde{\mathbf{w}}(t)) = \mathbf{0}$ and

$$\begin{aligned} \mathbf{C} &= E(\tilde{\mathbf{w}}(t) \tilde{\mathbf{w}}^H(t)) \\ &= P \mathbf{i} \mathbf{i}^H + \sigma^2 \mathbf{I}. \end{aligned}$$

Then, using Woodbury's identity (see Appendix 1),

$$\mathbf{C}^{-1} = \frac{1}{\sigma^2} \left(\mathbf{I} - \frac{P}{MP + \sigma^2} \mathbf{i} \mathbf{i}^H \right)$$

and thus

$$\begin{aligned} \mathbf{a}_{\text{opt}} &= \frac{\mathbf{e} - \frac{P \mathbf{i} \mathbf{i}^H \mathbf{e}}{MP + \sigma^2} \mathbf{i}}{\mathbf{e}^H \mathbf{e} - \frac{P |\mathbf{e}^H \mathbf{i}|^2}{MP + \sigma^2}} \\ &= c_1 \mathbf{e} - c_2 \mathbf{i} \end{aligned}$$

where

$$\begin{aligned} c_1 &= \frac{1}{M - \frac{P |\mathbf{e}^H \mathbf{i}|^2}{MP + \sigma^2}} \\ c_2 &= \frac{P \mathbf{i}^H \mathbf{e}}{M(MP + \sigma^2) - P |\mathbf{e}^H \mathbf{i}|^2}. \end{aligned}$$

We see that the beamformer attempts to subtract out the interference, with the amount depending upon the interference-to-noise ratio (P/σ^2) as well as the separation in arrival angles between the signal and the interference ($\mathbf{e}^H \mathbf{i}$). Note that if $\mathbf{e}^H \mathbf{i} = 0$, we will have a conventional beamformer because then $\mathbf{a}_{\text{opt}} = \mathbf{e}/M$. This is because the interference




Figure 15.6 Response of adaptive beamformer to interference

is perfectly nullified since $\mathbf{a}_{\text{opt}}^H \mathbf{i} = 0$. As an example, if the input is given by (15.76), the output is

$$\begin{aligned} \tilde{\mathbf{y}}(t) &= \mathbf{a}_{\text{opt}}^H \left[\tilde{A} \exp(j2\pi F_0 t) \mathbf{e} + \tilde{B} \exp(j2\pi F_0 t) \mathbf{i} + \tilde{\mathbf{u}}(t) \right] \\ &= \tilde{A} \exp(j2\pi F_0 t) + \tilde{B} \exp(j2\pi F_0 t) (c_1 e^H \mathbf{i} - M c_2^*) + \mathbf{a}_{\text{opt}}^H \tilde{\mathbf{u}}(t). \end{aligned}$$

The signal is passed undistorted, while the interference power at the output divided by the interference power at the input is

$$|c_1 e^H \mathbf{i} - M c_2^*|^2. \quad (15.77)$$

As an example, we assume a line array with half-wavelength spacing so that the spatial frequency is

$$\begin{aligned} f_s &= F_0 \left(\frac{d}{c} \right) \cos \beta \\ &= F_0 \left(\frac{\lambda/2}{F_0 \lambda} \right) \cos \beta \\ &= \frac{1}{2} \cos \beta. \end{aligned}$$

Then, for $P = 10$, $\sigma^2 = 1$, $M = 10$, and a signal at $\beta_s = 90^\circ$ so that $f_s = 0$, we plot (15.77) in dB versus the arrival angle of the interference in Figure 15.6. Note that the response is unity when $\beta_i = 90^\circ$ or when the interference arrival angle is the same as that of the signal. However, it quickly drops to about -60 dB when the difference in arrival angles is about 6° . Recall that the adaptive beamformer is constrained to pass the signal undistorted, so that when $\beta_i = \beta_s$, the interference cannot be attenuated.

◇

References

- Brandwood, D.H., "A Complex Gradient Operator and Its Application in Adaptive Array Theory," *IEE Proc.*, Vol. 130, pp. 11-16, Feb. 1983.
- Cox, H., "Resolving Power and Sensitivity to Mismatch of Optimum Array Processors," *J. Acoust. Soc. Am.*, Vol. 54, pp. 771-785, 1973.
- Dzhaparidze, K., *Parameter Estimation and Hypothesis Testing in Spectral Analysis and Stationary Time Series*, Springer-Verlag, New York, 1986.
- Grenander, U., G. Szego, *Toeplitz Forms and Their Applications*, University of California Press, Berkeley, 1958.
- Kay, S.M., *Modern Spectral Estimation: Theory and Application*, Prentice-Hall, Englewood Cliffs, N.J., 1988.
- Knight, W.S., R.G. Pridham, S.M. Kay, "Digital Signal Processing for Sonar," *Proc. IEEE*, Vol. 69, pp. 1451-1506, Nov. 1981.
- Miller, K.S., *Complex Stochastic Processes*, Addison-Wesley, Reading, Mass., 1974, available from University Microfilms International, Ann Arbor, Mich.
- Monzingo, R.A., T.W. Miller, *Introduction to Adaptive Arrays*, J. Wiley, New York, 1980.
- Owsley, N.L., "Sonar Array Processing," in *Array Signal Processing*, S. Haykin, ed., Prentice-Hall, Englewood Cliffs, N.J., 1985.
- Papoulis, A., *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, 1965.
- Rao, C.R., *Linear Statistical Inference and Its Applications*, J. Wiley, New York, 1973.
- Van Trees, H.L., *Detection, Estimation, and Modulation Theory*, Part III, J. Wiley, New York, 1971.
- Williams, J.R., G.G. Ricker, "Signal Detectability of Optimum Fourier Receivers," *IEEE Trans. Audio Electroacoust.*, Vol. 20, pp. 264-270, Oct. 1972.

Problems

- 15.1 Two complex random variables $\tilde{x}_1 = u_1 + jv_1$ and $\tilde{x}_2 = u_2 + jv_2$ are said to be independent if the real random vectors $[u_1 \ v_1]^T$ and $[u_2 \ v_2]^T$ are independent. Prove that if two complex random variables are independent, then their covariance is zero.
- 15.2 Prove that the complex covariance matrix in (15.13) is positive semidefinite. Under what conditions will it be positive definite?
- 15.3 The real random vector $\mathbf{x} = [u_1 \ u_2 \ v_1 \ v_2]^T$ has the PDF $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_x)$, where

$$\mathbf{C}_x = \begin{bmatrix} 2 & 1 & 0 & -1 \\ 1 & 2 & 1 & 0 \\ 0 & 1 & 2 & 1 \\ -1 & 0 & 1 & 2 \end{bmatrix}.$$

Can we define a complex Gaussian random vector, and if so, what is its covariance matrix?

- 15.4 For the real covariance matrix in Problem 15.3 verify directly the following:

- a. $\mathbf{x}^T \mathbf{C}_x^{-1} \mathbf{x} = 2\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{x}}^{-1} \tilde{\mathbf{x}}$, where $\mathbf{x} = [\mathbf{u}^T \mathbf{v}^T]^T$ and $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$
 b. $\det(\mathbf{C}_x) = \det^2(\mathbf{C}_{\tilde{x}})/16$

15.5 Let M^2 denote the vector space of real 2×2 matrices of the form

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix}$$

and C^1 the vector space of scalar complex numbers. Define addition and multiplication in the vector spaces in the usual way. Then, let m_1, m_2 be two vectors in M^2 and consider the operations

- a. αm_1 , where α is a real scalar
 b. $m_1 + m_2$
 c. $m_1 m_2$

Show that these operations can be performed as follows.

- a. Transform m_1, m_2 to C^1 using the transformation

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix} \rightarrow a + jb.$$

- b. Carry out the equivalent operation in C^1 .
 c. Transform the result back to M^2 using the inverse transformation

$$e + jf \rightarrow \begin{bmatrix} e & -f \\ f & e \end{bmatrix}.$$

Because operations in M^2 are equivalent to those in C^1 , we say that M^2 is isomorphic to C^1 .

15.6 Using the ideas in Problem 15.5, compute the matrix product $\mathbf{A}^T \mathbf{A}$, where

$$\mathbf{A} = \begin{bmatrix} a_1 & -b_1 \\ b_1 & a_1 \\ a_2 & -b_2 \\ b_2 & a_2 \\ a_3 & -b_3 \\ b_3 & a_3 \end{bmatrix}$$

by manipulating complex numbers.

15.7 In this problem we prove that all third-order moments of a complex Gaussian PDF are zero. The third-order moments are $E(\tilde{x}^3), E(\tilde{x}^* \tilde{x}^3), E(\tilde{x}^* \tilde{x}^2), E(\tilde{x} \tilde{x}^* \tilde{x})$. Prove that these are all zero by using the characteristic function (see Appendix 15B)

$$\phi_{\tilde{x}}(\tilde{\omega}) = \exp\left(-\frac{1}{4}\sigma^2|\tilde{\omega}|^2\right).$$

- 15.8** Consider the zero mean complex random variable $\tilde{x}$. If we assume that $E(\tilde{x}^2) = 0$, what does this say about the variances and covariances of the real and imaginary parts of $\tilde{x}$?
- 15.9** Show that the assumption $E[(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^T] = \mathbf{0}$, where $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$ leads to the special form of the real covariance matrix for $[\mathbf{u}^T \mathbf{v}^T]^T$ given by (15.19).
- 15.10** A complex Gaussian random vector $\tilde{\mathbf{x}}$ has the PDF $\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{B})$, where $\mathbf{B}$ is a known complex covariance matrix. It is desired to estimate σ^2 using the estimator $\hat{\sigma}^2 = \tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}}$, where $\mathbf{A}$ is a Hermitian matrix. Find $\mathbf{A}$ so that $\hat{\sigma}^2$ is unbiased and has minimum variance. If $\mathbf{B} = \mathbf{I}$, what is $\hat{\sigma}^2$? Hint: Use (15.29) and (15.30) and the fact that $\text{tr}(\mathbf{D}^k) = \sum_{i=1}^N \lambda_i^k$, where the λ_i 's are the eigenvalues of the $N \times N$ matrix $\mathbf{D}$ and k is a positive integer.
- 15.11** If $\tilde{x}[n]$ is CWGN with variance σ^2 , find the mean and variance of $\sum_{n=0}^{N-1} |\tilde{x}[n]|^2$. Propose an estimator for σ^2 . Compare its variance to that of $\hat{\sigma}^2 = (1/N) \sum_{n=0}^{N-1} x^2[n]$ for the real WGN case. Explain your results.
- 15.12** In this problem we show that the random analytic signal is a complex Gaussian process. First consider the real zero mean WSS Gaussian random process $u[n]$. Assume that $u[n]$ is input to a Hilbert transformer, which is a linear time-invariant system with frequency response

$$H(f) = \begin{cases} -j & 0 \leq f < \frac{1}{2} \\ j & -\frac{1}{2} \leq f < 0 \end{cases}$$

and denote the output by $v[n]$. Show that $v[n]$ is also a real zero mean WSS Gaussian process. Then, show that $\tilde{x}[n] = u[n] + jv[n]$, termed the *analytic signal*, is a complex WSS Gaussian random process by verifying that $r_{uu}[k] = r_{vv}[k]$ and $r_{uv}[k] = -r_{vu}[k]$. What is the PDF of $\tilde{x}[n]$?

- 15.13** Prove that $\partial \theta^* / \partial \theta = 0$. What would $\partial \theta / \partial \theta$ and $\partial \theta^* / \partial \theta$ be if we used the alternative definition

$$\frac{\partial}{\partial \theta} = \frac{\partial}{\partial \alpha} + j \frac{\partial}{\partial \beta}$$

for the complex derivative?

- 15.14** Prove that

$$\frac{\partial \theta^H \mathbf{b}}{\partial \theta} = 0$$

where $\mathbf{b}$ and θ are complex.

- 15.15** Determine the complex gradient of $\theta^H \mathbf{A} \theta$ with respect to complex θ for an arbitrary complex $p \times p$ matrix $\mathbf{A}$, not necessarily Hermitian. Note that in this case $\theta^H \mathbf{A} \theta$ is not necessarily real.

15.16 If we observe the complex data

$$\tilde{x}[n] = \tilde{A}\gamma^n + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A}$ is a complex deterministic amplitude, γ is a known complex constant, and $\tilde{w}[n]$ is CWGN with variance σ^2 , find the LSE and also the MLE of $\tilde{A}$. Also, find its mean and variance. Explain what happens as $N \rightarrow \infty$ if $|\gamma| < 1$, $|\gamma| = 1$, and $|\gamma| > 1$.

15.17 We observe the complex data $\tilde{x}[n]$ for $n = 0, 1, \dots, N-1$. It is known that the $\tilde{x}[n]$'s are uncorrelated and have equal variance σ^2 . The mean is $E(\tilde{x}[n]) = A$, where A is *real*. Find the BLUE of A and compare it to the case when A is complex (see Example 15.7). Explain your results.

15.18 If we observe the complex data $\tilde{x}[n] = \tilde{s}[n; \theta] + \tilde{w}[n]$, where the deterministic signal is known to within a real parameter θ and $\tilde{w}[n]$ is CWGN with variance σ^2 , find the general CRLB for θ . Compare it to the case of real data (see Section 3.5) and explain the difference.

15.19 If $\tilde{x}[n]$ is CWGN with variance σ^2 , find the CRLB for σ^2 based on N complex samples. Can the bound be attained, and if so what is the efficient estimator?

15.20 In this problem we study the complex equivalent of the Gauss-Markov theorem (see Chapter 6). The data model is

$$\tilde{\mathbf{x}} = \mathbf{H}\boldsymbol{\theta} + \tilde{\mathbf{w}}$$

where $\mathbf{H}$ is a known complex $N \times p$ matrix with $N > p$ and full rank, $\boldsymbol{\theta}$ is a complex $p \times 1$ parameter vector to be estimated, and $\tilde{\mathbf{w}}$ is a complex $N \times 1$ noise vector with zero mean and covariance matrix $\mathbf{C}$. Show that the BLUE of $\boldsymbol{\theta}$ is given by (15.58). Hint: Let $\hat{\theta}_i = \mathbf{a}_i^H \tilde{\mathbf{x}}$ and use (15.51).

15.21 The MLE for the frequencies of two complex sinusoids in CWGN is explored in this problem, extending the results in Example 15.13. The data model is

$$\tilde{x}[n] = \tilde{A}_1 \exp(j2\pi f_1 n) + \tilde{A}_2 \exp(j2\pi f_2 n) + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A}_1, \tilde{A}_2$ are complex deterministic amplitudes that are unknown, f_1, f_2 are unknown frequencies which are of primary interest to us, and $\tilde{w}[n]$ is CWGN with variance σ^2 . We wish to estimate the frequencies, but since the amplitudes are also unknown, we will also need to estimate them as well. Show that to find the MLE of the frequencies we will need to maximize the function

$$J(f_1, f_2) = \tilde{\mathbf{x}}^H \mathbf{E}(\mathbf{E}^H \mathbf{E})^{-1} \mathbf{E}^H \tilde{\mathbf{x}}$$

where $\mathbf{E} = [\mathbf{e}_1 \ \mathbf{e}_2]$ and $\mathbf{e}_i = [1 \ \exp(j2\pi f_i) \ \dots \ \exp(j2\pi f_i(N-1))]^T$ for $i = 1, 2$. To do so first note that for known frequencies the data model is linear in the amplitudes and so the PDF can be maximized easily over $\tilde{A}_1, \tilde{A}_2$. Finally, show that under the constraint $|f_1 - f_2| \gg 1/N$, the function J decouples into the sum of two periodograms. Now determine the MLE. Hint: Show that $\mathbf{E}^H \mathbf{E}$ is approximately diagonal when the frequencies are spaced far apart.

- 15.22** A signal frequently used in radar/sonar is a chirp whose discrete-time equivalent is

$$\tilde{s}[n] = \tilde{A} \exp \left[j2\pi \left(f_0 n + \frac{1}{2} \alpha n^2 \right) \right]$$

where the parameter α is real. The instantaneous frequency of the chirp may be defined as the difference of the phase for successive samples or

$$\begin{aligned} f_i[n] &= \left[f_0 n + \frac{1}{2} \alpha n^2 \right] - \left[f_0 (n-1) + \frac{1}{2} \alpha (n-1)^2 \right] \\ &= \left(f_0 - \frac{1}{2} \alpha \right) + \alpha n. \end{aligned}$$

The parameter α is thus the *frequency sweep rate*. Assume the chirp signal is embedded in CWGN with variance σ^2 and N samples are observed. Show that the function to be maximized to yield the MLE of f_0 and α is

$$\left| \sum_{n=0}^{N-1} \tilde{x}[n] \exp \left[-j2\pi \left(f_0 n + \frac{1}{2} \alpha n^2 \right) \right] \right|^2.$$

Assume that $\tilde{A}$ is an unknown deterministic constant. Show how the function can be computed efficiently using an FFT for each assumed sweep rate.

- 15.23** Consider the complex random parameter $\theta = \alpha + j\beta$, which is to be estimated based on the complex data vector $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$. The real Bayesian MMSE estimator is to be used for each real parameter. Hence, we wish to find the estimators that minimize the Bayesian MSEs

$$\begin{aligned} \text{Bmse}(\hat{\alpha}) &= E[(\alpha - \hat{\alpha})^2] \\ \text{Bmse}(\hat{\beta}) &= E[(\beta - \hat{\beta})^2] \end{aligned}$$

where the expectations are with respect to the PDFs $p(\mathbf{u}, \mathbf{v}, \alpha)$ and $p(\mathbf{u}, \mathbf{v}, \beta)$, respectively. Show that

$$\begin{aligned} \hat{\theta} &= \hat{\alpha} + j\hat{\beta} \\ &= E(\theta | \mathbf{u}, \mathbf{v}) \\ &= E(\theta | \tilde{\mathbf{x}}) \end{aligned}$$

is the MMSE estimator. Comment on the PDF used to implement the expectation operator of $E(\theta | \tilde{\mathbf{x}})$.

- 15.24** Assume that $\tilde{x}[n]$ is a zero mean complex Gaussian WSS random process whose PSD is given as $P_{\tilde{x}\tilde{x}}(f) = P_0 Q(f)$, where P_0 is a real deterministic parameter to be estimated. The function $Q(f)$ satisfies

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} Q(f) df = 1$$

so that P_0 is the total power in $\tilde{x}[n]$. Find the CRLB and the MLE for P_0 by using the exact results as well as the asymptotic results (see Section 15.9) and compare.

- 15.25** In this problem we examine the “processing gain” of a DFT in detecting a complex sinusoid in CWGN. Assume that we observe

$$\tilde{x}[n] = \tilde{A} \exp(j2\pi f_c n) + \tilde{w}[n] \quad n = 0, 1, \dots, N-1$$

where $\tilde{A}$ is a complex deterministic amplitude and $\tilde{w}[n]$ is CWGN with variance σ^2 . If we take the DFT of $\tilde{x}[n]$ and if $f_c = l/N$ for some integer l , find the PDF of the DFT coefficients $X(f_k)$ for $k = 0, 1, \dots, N-1$. Note that

$$\sum_{n=0}^{N-1} \exp\left(j2\pi \frac{k}{N} n\right) = 0$$

for $k = 0, 1, \dots, N-1$. Then, compare the input SNR or

$$\text{SNR}(\text{input}) = \frac{|E(\tilde{x}[n])|^2}{\sigma^2}$$

to the output SNR of the DFT at the frequency of the sinusoid or

$$\text{SNR}(\text{output}) = \frac{|E(X(f_c))|^2}{\text{var}(X(f_c))}$$

and discuss. The improvement in the SNR is called the processing gain [Williams and Ricker 1972]. If the frequency was unknown, how could you detect its presence?

Appendix 15A

Derivation of Properties of Complex Covariance Matrices

We assume that the $2n \times 2n$ real covariance matrix $\mathbf{C}_x$ has the form of (15.19), where $\mathbf{A}$ is a real and symmetric $n \times n$ matrix, $\mathbf{B}$ is a real and skew-symmetric $n \times n$ matrix, and $\mathbf{C}_x$ is positive definite. We define a complex matrix as $\mathbf{C}_{\tilde{x}} = \mathbf{A} + j\mathbf{B}$. Then,

1. $\mathbf{C}_{\tilde{x}}$ is Hermitian.

$$\begin{aligned}\mathbf{C}_{\tilde{x}}^H &= (\mathbf{A} + j\mathbf{B})^H \\ &= \mathbf{A}^T - j\mathbf{B}^T \\ &= \mathbf{A} + j\mathbf{B} \\ &= \mathbf{C}_{\tilde{x}}.\end{aligned}$$

2. $\mathbf{C}_{\tilde{x}}$ is positive definite.

We first prove that if $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$, then

$$\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{x}} \tilde{\mathbf{x}} = 2\mathbf{x}^T \mathbf{C}_x \mathbf{x}$$

where $\mathbf{x} = [\mathbf{u}^T \ \mathbf{v}^T]^T$.

$$\begin{aligned}\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{x}} \tilde{\mathbf{x}} &= (\mathbf{u}^T - j\mathbf{v}^T)(\mathbf{A} + j\mathbf{B})(\mathbf{u} + j\mathbf{v}) \\ &= \mathbf{u}^T \mathbf{A} \mathbf{u} + j\mathbf{u}^T \mathbf{A} \mathbf{v} + j\mathbf{u}^T \mathbf{B} \mathbf{u} - \mathbf{u}^T \mathbf{B} \mathbf{v} \\ &\quad - j\mathbf{v}^T \mathbf{A} \mathbf{u} + \mathbf{v}^T \mathbf{A} \mathbf{v} + \mathbf{v}^T \mathbf{B} \mathbf{u} + j\mathbf{v}^T \mathbf{B} \mathbf{v}.\end{aligned}$$

Since $\mathbf{B}$ is skew-symmetric, $\mathbf{u}^T \mathbf{B} \mathbf{u} = \mathbf{v}^T \mathbf{B} \mathbf{v} = 0$, and since $\mathbf{A}$ is symmetric, we have

$$\begin{aligned}\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{x}} \tilde{\mathbf{x}} &= \mathbf{u}^T \mathbf{A} \mathbf{u} - \mathbf{u}^T \mathbf{B} \mathbf{v} + \mathbf{v}^T \mathbf{B} \mathbf{u} + \mathbf{v}^T \mathbf{A} \mathbf{v} \\ &= \begin{bmatrix} \mathbf{u}^T & \mathbf{v}^T \end{bmatrix} \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \\ &= 2\mathbf{x}^T \mathbf{C}_x \mathbf{x}.\end{aligned}$$

Thus, if $\mathbf{C}_x$ is positive definite or $\mathbf{x}^T \mathbf{C}_x \mathbf{x} > 0$ for all $\mathbf{x} \neq \mathbf{0}$, it follows that $\tilde{\mathbf{x}}^H \mathbf{C}_{\tilde{x}} \tilde{\mathbf{x}} > 0$ for all $\tilde{\mathbf{x}} \neq \mathbf{0}$ since $\mathbf{x} = \mathbf{0}$ if and only if $\tilde{\mathbf{x}} = \mathbf{0}$.

3. $\mathbf{C}_{\tilde{\mathbf{x}}}$ is the covariance matrix of $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$.

By definition $\mathbf{C}_{\tilde{\mathbf{x}}} = \mathbf{A} + j\mathbf{B}$, where

$$\begin{aligned}\frac{1}{2}\mathbf{A} &= E[(\mathbf{u} - E(\mathbf{u}))(\mathbf{u} - E(\mathbf{u}))^T] = E[(\mathbf{v} - E(\mathbf{v}))(\mathbf{v} - E(\mathbf{v}))^T] \\ \frac{1}{2}\mathbf{B} &= E[(\mathbf{v} - E(\mathbf{v}))(\mathbf{u} - E(\mathbf{u}))^T].\end{aligned}$$

But

$$\begin{aligned}\mathbf{C}_{\tilde{\mathbf{x}}} &= E[(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))(\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))^H] \\ &= E\left\{[(\mathbf{u} - E(\mathbf{u})) + j(\mathbf{v} - E(\mathbf{v}))][(\mathbf{u} - E(\mathbf{u})) - j(\mathbf{v} - E(\mathbf{v}))]^T\right\} \\ &= \frac{1}{2}\mathbf{A} - \frac{1}{2}j\mathbf{B}^T + \frac{1}{2}j\mathbf{B} + \frac{1}{2}\mathbf{A} \\ &= \mathbf{A} + j\mathbf{B}.\end{aligned}$$

4. The inverse of $\mathbf{C}_x$ has the special form

$$\mathbf{C}_x^{-1} = 2 \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix}$$

so that the inverse of the complex covariance matrix may be found directly from the inverse of the real covariance matrix as $\mathbf{C}_x^{-1} = \mathbf{E} + j\mathbf{F}$.

To show this we use the partitioned matrix formula (see Appendix 1) to yield

$$\begin{aligned}\mathbf{C}_x^{-1} &= 2 \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix}^{-1} \\ &= 2 \begin{bmatrix} (\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B})^{-1} & (\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{B}\mathbf{A}^{-1} \\ -(\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{B}\mathbf{A}^{-1} & (\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B})^{-1} \end{bmatrix} \\ &= 2 \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix}.\end{aligned}$$

Now, to show that $\mathbf{E} + j\mathbf{F}$ is the inverse of $\mathbf{C}_{\tilde{\mathbf{x}}}$

$$\begin{aligned}\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} &= \mathbf{C}_x \mathbf{C}_x^{-1} \\ &= \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{E} & -\mathbf{F} \\ \mathbf{F} & \mathbf{E} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{AE} - \mathbf{BF} & -\mathbf{AF} - \mathbf{BE} \\ \mathbf{BE} + \mathbf{AF} & -\mathbf{BF} + \mathbf{AE} \end{bmatrix}\end{aligned}$$

and thus

$$\begin{aligned}\mathbf{AE} - \mathbf{BF} &= \mathbf{I} \\ \mathbf{BE} + \mathbf{AF} &= \mathbf{0}.\end{aligned}$$

Using these relations, we have

$$\begin{aligned}
 \mathbf{C}_{\bar{x}} \mathbf{C}_{\bar{x}}^{-1} &= (\mathbf{A} + j\mathbf{B})(\mathbf{E} + j\mathbf{F}) \\
 &= \mathbf{A}\mathbf{E} - \mathbf{B}\mathbf{F} + j(\mathbf{B}\mathbf{E} + \mathbf{A}\mathbf{F}) \\
 &= \mathbf{I} \\
 \mathbf{C}_{\bar{x}}^{-1} \mathbf{C}_{\bar{x}} &= (\mathbf{C}_{\bar{x}}^H \mathbf{C}_{\bar{x}}^{-1H})^H \\
 &= (\mathbf{C}_{\bar{x}} \mathbf{C}_{\bar{x}}^{-1})^H \\
 &= \mathbf{I}^H = \mathbf{I}.
 \end{aligned}$$

$$5. (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}_x^{-1} (\mathbf{x} - \boldsymbol{\mu}) = 2(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}})^H \mathbf{C}_{\bar{x}}^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}).$$

Let $\mathbf{y} = \mathbf{x} - \boldsymbol{\mu}$ and $\tilde{\mathbf{y}} = \tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}$, so that we wish to show

$$\mathbf{y}^T \mathbf{C}_x^{-1} \mathbf{y} = 2\tilde{\mathbf{y}}^H \mathbf{C}_{\bar{x}}^{-1} \tilde{\mathbf{y}}.$$

But from property 4 we know that $\mathbf{C}_x^{-1}$ has the special form, and from property 2 we have already shown the equivalence of a quadratic and Hermitian form. Accounting for the factor of 2 because we now have an *inverse* covariance matrix, this property follows.

$$6. \det(\mathbf{C}_x) = \det^2(\mathbf{C}_{\bar{x}})/(2^{2n})$$

We first note that $\det(\mathbf{C}_{\bar{x}})$ is real. This is because $\mathbf{C}_{\bar{x}}$ is Hermitian, resulting in real eigenvalues, and the relationship $\det(\mathbf{C}_{\bar{x}}) = \prod_{i=1}^n \lambda_i$. Also, because $\mathbf{C}_{\bar{x}}$ is positive definite, we have $\lambda_i > 0$, and thus $\det(\mathbf{C}_{\bar{x}}) > 0$. The determinant is real and positive. Now since

$$\mathbf{C}_x = \frac{1}{2} \begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix}$$

we can use the determinant formula for partitioned matrices (see Appendix 1) to yield

$$\begin{aligned}
 \det(\mathbf{C}_x) &= \left(\frac{1}{2}\right)^{2n} \det \left(\begin{bmatrix} \mathbf{A} & -\mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \right) \\
 &= \left(\frac{1}{2}\right)^{2n} \det(\mathbf{A}) \det(\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B}).
 \end{aligned}$$

But $\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B} = (\mathbf{A} - j\mathbf{B})\mathbf{A}^{-1}(\mathbf{A} + j\mathbf{B})$, so that

$$\begin{aligned}
 \det(\mathbf{A} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B}) &= \frac{\det(\mathbf{A} - j\mathbf{B}) \det(\mathbf{A} + j\mathbf{B})}{\det(\mathbf{A})} \\
 &= \frac{\det(\mathbf{C}_{\bar{x}}^*) \det(\mathbf{C}_{\bar{x}})}{\det(\mathbf{A})}.
 \end{aligned}$$

Since $\det(\mathbf{C}_{\bar{x}}^*) = \det^*(\mathbf{C}_{\bar{x}}) = \det(\mathbf{C}_{\bar{x}})$, we have finally

$$\det(\mathbf{C}_x) = \frac{1}{2^{2n}} \det^2(\mathbf{C}_{\bar{x}}).$$

Appendix 15B

Derivation of Properties of Complex Gaussian PDF

1. Any subset of a complex Gaussian random vector is also complex Gaussian.

This property follows from property 4 to be proven. If we let $\tilde{\mathbf{y}}$ be a subset of the elements of $\tilde{\mathbf{x}}$ by setting $\tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}}$, where $\mathbf{A}$ is an appropriate linear transformation, then by property 4 $\tilde{\mathbf{y}}$ is also complex Gaussian. For example, if $\tilde{\mathbf{x}} = [\tilde{x}_1 \ \tilde{x}_2]^T$ is a complex Gaussian random vector, then

$$\tilde{\mathbf{y}} = \tilde{x}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tilde{\mathbf{x}}$$

will also be complex Gaussian.

2. If $\tilde{x}_1$ and $\tilde{x}_2$ are jointly complex Gaussian and uncorrelated, they are also independent.

By inserting the covariance matrix

$$\mathbf{C}_{\tilde{\mathbf{x}}} = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}$$

into (15.22), it is easy to see that the PDF factors as $p(\tilde{\mathbf{x}}) = p(\tilde{x}_1)p(\tilde{x}_2)$. Then, we note that $p(\tilde{x}_1)$ is equivalent to $p(u_1, v_1)$ and $p(\tilde{x}_2)$ is equivalent to $p(u_2, v_2)$. Hence, $[u_1 \ v_1]^T$ is independent of $[u_2 \ v_2]^T$. Of course, this property extends to any number of uncorrelated random variables.

3. If $\tilde{x}_1$ and $\tilde{x}_2$ are each complex Gaussian and also independent, then $\tilde{\mathbf{x}} = [\tilde{x}_1 \ \tilde{x}_2]^T$ is complex Gaussian.

This follows by forming $p(\tilde{x}_1)p(\tilde{x}_2)$ and noting that it is equal to $p(\tilde{\mathbf{x}})$ of (15.22). Again this extends to any number of independent complex Gaussian random variables.

4. A linear transformation of a complex Gaussian random vector is also complex Gaussian.

We first determine the characteristic function. It is well known that for a real multivariate Gaussian PDF or $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}_x)$ the characteristic function is [Rao 1973]

$$\begin{aligned}\phi_{\mathbf{x}}(\boldsymbol{\omega}) &= E[\exp(j\boldsymbol{\omega}^T \mathbf{x})] \\ &= \exp\left[j\boldsymbol{\omega}^T \boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\omega}^T \mathbf{C}_x \boldsymbol{\omega}\right]\end{aligned}$$

where $\boldsymbol{\omega} = [\omega_1 \omega_2 \dots \omega_{2n}]^T$. If $\mathbf{x} = [\mathbf{u}^T \mathbf{v}^T]^T$ and $\mathbf{C}_x$ has the special form, then letting $\tilde{\mathbf{x}} = \mathbf{u} + j\mathbf{v}$, $\tilde{\boldsymbol{\mu}} = \boldsymbol{\mu}_u + j\boldsymbol{\mu}_v$, and $\tilde{\boldsymbol{\omega}} = [\boldsymbol{\omega}_R^T \boldsymbol{\omega}_I^T]^T$, so that $\tilde{\boldsymbol{\omega}} = \boldsymbol{\omega}_R + j\boldsymbol{\omega}_I$, we have

$$\begin{aligned}\boldsymbol{\omega}^T \boldsymbol{\mu} &= [\boldsymbol{\omega}_R^T \boldsymbol{\omega}_I^T] \begin{bmatrix} \boldsymbol{\mu}_u \\ \boldsymbol{\mu}_v \end{bmatrix} \\ &= \boldsymbol{\omega}_R^T \boldsymbol{\mu}_u + \boldsymbol{\omega}_I^T \boldsymbol{\mu}_v \\ &= \text{Re}(\tilde{\boldsymbol{\omega}}^H \tilde{\boldsymbol{\mu}})\end{aligned}$$

and $2\boldsymbol{\omega}^T \mathbf{C}_x \boldsymbol{\omega} = \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{x}} \tilde{\boldsymbol{\omega}}$ from Appendix 15A. Likewise, we have that $\boldsymbol{\omega}^T \mathbf{x} = \text{Re}(\tilde{\boldsymbol{\omega}}^H \tilde{\mathbf{x}})$. Hence we can define the characteristic function of $\tilde{\mathbf{x}}$ as

$$\phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}}) = E[\exp(j \text{Re}(\tilde{\boldsymbol{\omega}}^H \tilde{\mathbf{x}}))] \quad (15B.1)$$

and for the complex Gaussian PDF this becomes

$$\phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}}) = \exp\left[j \text{Re}(\tilde{\boldsymbol{\omega}}^H \tilde{\boldsymbol{\mu}}) - \frac{1}{4}\tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{x}} \tilde{\boldsymbol{\omega}}\right]. \quad (15B.2)$$

Now, if $\tilde{\mathbf{y}} = \mathbf{A}\tilde{\mathbf{x}} + \mathbf{b}$, we have

$$\begin{aligned}\phi_{\tilde{\mathbf{y}}}(\tilde{\boldsymbol{\omega}}) &= E[\exp(j \text{Re}(\tilde{\boldsymbol{\omega}}^H \tilde{\mathbf{y}}))] \\ &= E[\exp(j \text{Re}(\tilde{\boldsymbol{\omega}}^H \mathbf{A}\tilde{\mathbf{x}} + \tilde{\boldsymbol{\omega}}^H \mathbf{b}))] \\ &= \exp[j \text{Re}(\tilde{\boldsymbol{\omega}}^H \mathbf{b})] E[\exp(j \text{Re}((\mathbf{A}^H \tilde{\boldsymbol{\omega}})^H \tilde{\mathbf{x}}))] \\ &= \exp[j \text{Re}(\tilde{\boldsymbol{\omega}}^H \mathbf{b})] \phi_{\tilde{\mathbf{x}}}(\mathbf{A}^H \tilde{\boldsymbol{\omega}}) \\ &= \exp[j \text{Re}(\tilde{\boldsymbol{\omega}}^H \mathbf{b})] \exp[j \text{Re}(\tilde{\boldsymbol{\omega}}^H \mathbf{A}\tilde{\boldsymbol{\mu}})] \exp\left[-\frac{1}{4}\tilde{\boldsymbol{\omega}}^H \mathbf{A} \mathbf{C}_{\tilde{x}} \mathbf{A}^H \tilde{\boldsymbol{\omega}}\right] \\ &= \exp\left[j \text{Re}(\tilde{\boldsymbol{\omega}}^H (\mathbf{A}\tilde{\boldsymbol{\mu}} + \mathbf{b})) - \frac{1}{4}\tilde{\boldsymbol{\omega}}^H \mathbf{A} \mathbf{C}_{\tilde{x}} \mathbf{A}^H \tilde{\boldsymbol{\omega}}\right].\end{aligned}$$

By identifying the characteristic function with that for the complex Gaussian PDF we have

$$\tilde{\mathbf{y}} \sim \mathcal{CN}(\mathbf{A}\tilde{\boldsymbol{\mu}} + \mathbf{b}, \mathbf{A} \mathbf{C}_{\tilde{x}} \mathbf{A}^H).$$

5. The sum of two independent complex Gaussian random variables is also complex Gaussian.

Using the characteristic function, we have

$$\begin{aligned}
 \phi_{\tilde{x}_1 + \tilde{x}_2}(\tilde{\omega}) &= E[\exp(j \operatorname{Re}(\tilde{\omega}^*(\tilde{x}_1 + \tilde{x}_2)))] \\
 &= E[\exp(j \operatorname{Re}(\tilde{\omega}^* \tilde{x}_1)) \exp(j \operatorname{Re}(\tilde{\omega}^* \tilde{x}_2))] \\
 &= E[\exp(j \operatorname{Re}(\tilde{\omega}^* \tilde{x}_1))] E[\exp(j \operatorname{Re}(\tilde{\omega}^* \tilde{x}_2))] \\
 &= \phi_{\tilde{x}_1}(\tilde{\omega}) \phi_{\tilde{x}_2}(\tilde{\omega}) \\
 &= \exp \left[j \operatorname{Re}(\tilde{\omega}^* \tilde{\mu}_1) - \frac{1}{4} |\tilde{\omega}|^2 \sigma_1^2 \right] \exp \left[j \operatorname{Re}(\tilde{\omega}^* \tilde{\mu}_2) - \frac{1}{4} |\tilde{\omega}|^2 \sigma_2^2 \right] \\
 &= \exp \left[j \operatorname{Re}(\tilde{\omega}^* (\tilde{\mu}_1 + \tilde{\mu}_2)) - \frac{1}{4} |\tilde{\omega}|^2 (\sigma_1^2 + \sigma_2^2) \right].
 \end{aligned}$$

Hence, we have $\tilde{x}_1 + \tilde{x}_2 \sim \mathcal{CN}(\tilde{\mu}_1 + \tilde{\mu}_2, \sigma_1^2 + \sigma_2^2)$. Of course, this property extends to any number of independent complex Gaussian random variables.

6. The fourth moment of a complex Gaussian random vector is given by (15.24).

Consider the characteristic function for $\tilde{\mathbf{x}} = [\tilde{x}_1 \tilde{x}_2 \tilde{x}_3 \tilde{x}_4]^T$, where $\tilde{\mathbf{x}}$ is complex Gaussian with zero mean. Then,

$$\begin{aligned}
 \phi_{\tilde{\mathbf{x}}}(\tilde{\omega}) &= E[\exp(j \operatorname{Re}(\tilde{\omega}^H \tilde{\mathbf{x}}))] \\
 &= \exp \left[-\frac{1}{4} \tilde{\omega}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\omega} \right].
 \end{aligned}$$

We first show that

$$\frac{\partial^4 \phi_{\tilde{\mathbf{x}}}(\tilde{\omega})}{\partial \tilde{\omega}_1 \partial \tilde{\omega}_2^* \partial \tilde{\omega}_3 \partial \tilde{\omega}_4^*} \bigg|_{\tilde{\omega}=0} = \frac{1}{2^4} E(\tilde{x}_1^* \tilde{x}_2 \tilde{x}_3^* \tilde{x}_4). \quad (15B.3)$$

Differentiating, we obtain

$$\begin{aligned}
 \frac{\partial \phi_{\tilde{\mathbf{x}}}(\tilde{\omega})}{\partial \tilde{\omega}_i} &= \frac{\partial}{\partial \tilde{\omega}_i} E \left\{ \exp \left[\frac{j}{2} (\tilde{\omega}^H \tilde{\mathbf{x}} + \tilde{\mathbf{x}}^H \tilde{\omega}) \right] \right\} \\
 &= E \left\{ \frac{\partial}{\partial \tilde{\omega}_i} \exp \left[\frac{j}{2} (\tilde{\omega}^H \tilde{\mathbf{x}} + \tilde{\mathbf{x}}^H \tilde{\omega}) \right] \right\} \\
 &= E \left\{ \frac{j}{2} \tilde{x}_i^* \exp \left[\frac{j}{2} (\tilde{\omega}^H \tilde{\mathbf{x}} + \tilde{\mathbf{x}}^H \tilde{\omega}) \right] \right\}
 \end{aligned}$$

since, using (15.41) and (15.42),

$$\begin{aligned}
 \frac{\partial \tilde{\omega}_i}{\partial \tilde{\omega}_i} &= 1 \\
 \frac{\partial \tilde{\omega}_i^*}{\partial \tilde{\omega}_i} &= 0.
 \end{aligned}$$

Likewise,

$$\frac{\partial \phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}})}{\partial \tilde{\boldsymbol{\omega}}_i^*} = E \left\{ \frac{j}{2} \tilde{x}_i \exp \left[\frac{j}{2} (\tilde{\boldsymbol{\omega}}^H \tilde{\mathbf{x}} + \tilde{\mathbf{x}}^H \tilde{\boldsymbol{\omega}}) \right] \right\}$$

since

$$\begin{aligned} \frac{\partial \tilde{\boldsymbol{\omega}}_i}{\partial \tilde{\boldsymbol{\omega}}_i^*} &= 0 \\ \frac{\partial \tilde{\boldsymbol{\omega}}_i^*}{\partial \tilde{\boldsymbol{\omega}}_i^*} &= 1 \end{aligned}$$

(similar to the results of (15.41) and (15.42)). By repeated application (15B.3) follows. We now need to evaluate the fourth partial derivative of

$$\begin{aligned} \phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}}) &= \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \\ &= \exp\left(-\frac{1}{4} \sum_{i=1}^4 \sum_{j=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{ij} \tilde{\boldsymbol{\omega}}_j\right) \end{aligned}$$

which is straightforward but tedious. Proceeding, we have

$$\begin{aligned} \frac{\partial \phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}})}{\partial \tilde{\boldsymbol{\omega}}_1} &= -\frac{1}{4} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i1} \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \\ \frac{\partial^2 \phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}})}{\partial \tilde{\boldsymbol{\omega}}_1 \partial \tilde{\boldsymbol{\omega}}_2^*} &= \left(-\frac{1}{4} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i1}\right) \left(-\frac{1}{4} \sum_{j=1}^4 [\mathbf{C}_{\tilde{\mathbf{x}}}]_{2j} \tilde{\boldsymbol{\omega}}_j \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right)\right) \\ &\quad - \frac{1}{4} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{21} \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \\ \frac{\partial^3 \phi_{\tilde{\mathbf{x}}}(\tilde{\boldsymbol{\omega}})}{\partial \tilde{\boldsymbol{\omega}}_1 \partial \tilde{\boldsymbol{\omega}}_2^* \partial \tilde{\boldsymbol{\omega}}_3} &= \left(-\frac{1}{4} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i1}\right) \left[\left(-\frac{1}{4} \sum_{j=1}^4 [\mathbf{C}_{\tilde{\mathbf{x}}}]_{2j} \tilde{\boldsymbol{\omega}}_j\right) \left(-\frac{1}{4} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i3}\right) \right. \\ &\quad \cdot \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) - \frac{1}{4} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{23} \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \Big] \\ &\quad + \frac{1}{16} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{21} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i3} \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \\ &= -\left(\frac{1}{64} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i1} \sum_{j=1}^4 [\mathbf{C}_{\tilde{\mathbf{x}}}]_{2j} \tilde{\boldsymbol{\omega}}_j \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i3}\right) \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \\ &\quad + \frac{1}{16} \sum_{i=1}^4 \tilde{\boldsymbol{\omega}}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i1} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{23} \exp\left(-\frac{1}{4} \tilde{\boldsymbol{\omega}}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\boldsymbol{\omega}}\right) \end{aligned}$$

$$+ \frac{1}{16} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{21} \sum_{i=1}^4 \tilde{\omega}_i^* [\mathbf{C}_{\tilde{\mathbf{x}}}]_{i3} \exp \left(-\frac{1}{4} \tilde{\omega}^H \mathbf{C}_{\tilde{\mathbf{x}}} \tilde{\omega} \right).$$

Only the last two terms will be nonzero after differentiation with respect to $\tilde{\omega}_4^*$ and setting $\tilde{\omega} = \mathbf{0}$. Hence,

$$\left. \frac{\partial^4 \phi_{\tilde{\mathbf{x}}}(\tilde{\omega})}{\partial \tilde{\omega}_1 \partial \tilde{\omega}_2^* \partial \tilde{\omega}_3 \partial \tilde{\omega}_4^*} \right|_{\tilde{\omega}=\mathbf{0}} = \frac{1}{2^4} ([\mathbf{C}_{\tilde{\mathbf{x}}}]_{41} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{23} + [\mathbf{C}_{\tilde{\mathbf{x}}}]_{21} [\mathbf{C}_{\tilde{\mathbf{x}}}]_{43})$$

or

$$E(\tilde{x}_1^* \tilde{x}_2 \tilde{x}_3^* \tilde{x}_4) = E(\tilde{x}_4 \tilde{x}_1^*) E(\tilde{x}_2 \tilde{x}_3^*) + E(\tilde{x}_2 \tilde{x}_1^*) E(\tilde{x}_4 \tilde{x}_3^*).$$

7. The conditional PDF of a complex Gaussian PDF is also complex Gaussian with mean and covariance given by (15.61) and (15.62).

The easiest way to verify the form of the conditional PDF is to consider the complex random vector $\tilde{\mathbf{z}} = \tilde{\mathbf{y}} - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \tilde{\mathbf{x}}$. The PDF of $\tilde{\mathbf{z}}$ is complex Gaussian, being a linear transformation of jointly complex random vectors, with mean

$$E(\tilde{\mathbf{z}}) = E(\tilde{\mathbf{y}}) - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} E(\tilde{\mathbf{x}})$$

and covariance

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{z}}\tilde{\mathbf{z}}} &= E[(\tilde{\mathbf{y}} - E(\tilde{\mathbf{y}}) - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}}))) \\ &\quad \cdot (\tilde{\mathbf{y}} - E(\tilde{\mathbf{y}}) - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}})))^H] \\ &= \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{y}}} - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}}^H - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} \\ &\quad + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}}^H \\ &= \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{y}}} - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}}. \end{aligned}$$

But conditioned on $\tilde{\mathbf{x}}$ we have $\tilde{\mathbf{y}} = \tilde{\mathbf{z}} + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \tilde{\mathbf{x}}$, where $\tilde{\mathbf{x}}$ is just a constant. Then, $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$ must be a complex Gaussian PDF since $\tilde{\mathbf{z}}$ is complex Gaussian and $\tilde{\mathbf{x}}$ is a constant. The mean and covariance are found as

$$\begin{aligned} E(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) &= E(\tilde{\mathbf{z}}) + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \tilde{\mathbf{x}} \\ &= E(\tilde{\mathbf{y}}) - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} E(\tilde{\mathbf{x}}) + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \tilde{\mathbf{x}} \\ &= E(\tilde{\mathbf{y}}) + \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} (\tilde{\mathbf{x}} - E(\tilde{\mathbf{x}})) \end{aligned}$$

and

$$\begin{aligned} \mathbf{C}_{\tilde{\mathbf{y}}|\tilde{\mathbf{x}}} &= \mathbf{C}_{\tilde{\mathbf{z}}\tilde{\mathbf{z}}} \\ &= \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{y}}} - \mathbf{C}_{\tilde{\mathbf{y}}\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1} \mathbf{C}_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}}. \end{aligned}$$

Appendix 15C

Derivation of CRLB and MLE Formulas

We first prove (15.60) using (15.47).

$$\begin{aligned}\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \xi_i} &= -\frac{\partial \ln \det(\mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi}))}{\partial \xi_i} - \frac{\partial (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))}{\partial \xi_i} \\ &= -\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) - (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \frac{\partial}{\partial \xi_i} [\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))] \\ &\quad + \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))\end{aligned}$$

But

$$\begin{aligned}\frac{\partial}{\partial \xi_i} [\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))] &= -\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} + \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi})}{\partial \xi_i} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \\ &= -\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} - \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))\end{aligned}$$

The last step follows from $\mathbf{C}_{\tilde{\mathbf{x}}} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} = \mathbf{I}$, so that

$$\frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1} + \mathbf{C}_{\tilde{\mathbf{x}}} \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}}{\partial \xi_i} = \mathbf{0}.$$

Thus,

$$\begin{aligned}\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \xi_i} &= -\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \\ &\quad + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \\ &\quad + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i}\end{aligned}$$

$$+ \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))$$

from which (15.60) follows. Next we prove (15.52). To do so we need the following lemma. If $\tilde{\mathbf{x}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C})$, then for $\mathbf{A}$ and $\mathbf{B}$ Hermitian matrices [Miller 1974]

$$E(\tilde{\mathbf{x}}^H \mathbf{A} \tilde{\mathbf{x}} \tilde{\mathbf{x}}^H \mathbf{B} \tilde{\mathbf{x}}) = \text{tr}(\mathbf{A}\mathbf{C})\text{tr}(\mathbf{B}\mathbf{C}) + \text{tr}(\mathbf{A}\mathbf{C}\mathbf{B}\mathbf{C}).$$

By definition of the Fisher information matrix and (15.60)

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\xi})]_{ij} &= E \left[\frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \xi_i} \frac{\partial \ln p(\tilde{\mathbf{x}}; \boldsymbol{\xi})}{\partial \xi_j} \right] \\ &= E \left\{ \left[-\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \right. \right. \\ &\quad \left. \left. + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} + \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \right] \right. \\ &\quad \cdot \left[-\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \right) + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \right. \\ &\quad \left. \left. + (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j} + \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \right] \right\}. \end{aligned}$$

We note that all first- and third-order moments of $\tilde{\mathbf{y}} = \tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}$ are zero. Also, all second-order moments of the form $E(\tilde{\mathbf{y}}\tilde{\mathbf{y}}^T)$ and therefore $E(\tilde{\mathbf{y}}^* \tilde{\mathbf{y}}^H) = [E(\tilde{\mathbf{y}}\tilde{\mathbf{y}}^T)]^*$ are zero (see Problem 15.7), so that the expectation becomes

$$\begin{aligned} &= \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\ &\quad - \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\ &\quad - \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\ &\quad + E \left[\tilde{\mathbf{y}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \tilde{\mathbf{y}} \tilde{\mathbf{y}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \tilde{\mathbf{y}} \right] \\ &\quad + E \left[(\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) \right] \\ &\quad + E \left[\frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})) (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi}))^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j} \right] \end{aligned}$$

where we have used

$$E \left[\tilde{\mathbf{y}}^H \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \tilde{\mathbf{y}} \right] = E \left[\text{tr} \left(\mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\tilde{\mathbf{x}}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\tilde{\mathbf{x}}}^{-1}(\boldsymbol{\xi}) \tilde{\mathbf{y}} \tilde{\mathbf{y}}^H \right) \right].$$

$$\begin{aligned}
&= \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) E(\tilde{\mathbf{y}} \tilde{\mathbf{y}}^H) \right) \\
&= \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \right).
\end{aligned}$$

Simplifying, we have

$$\begin{aligned}
[\mathbf{I}(\boldsymbol{\xi})]_{ij} &= -\operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\
&\quad + E[\tilde{\mathbf{y}}^H \mathbf{A} \tilde{\mathbf{y}} \tilde{\mathbf{y}}^H \mathbf{B} \tilde{\mathbf{y}}] + \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \mathbf{C}_{\bar{x}}(\boldsymbol{\xi}) \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_i} \\
&\quad + \frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \mathbf{C}_{\bar{x}}(\boldsymbol{\xi}) \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j}
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{A} &= \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \\
\mathbf{B} &= \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}).
\end{aligned}$$

Note that the last two terms are complex conjugates of each other. Now, using the lemma,

$$\begin{aligned}
[\mathbf{I}(\boldsymbol{\xi})]_{ij} &= -\operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\
&\quad + \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \right) \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\
&\quad + \operatorname{tr} \left(\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \right) \\
&\quad + 2 \operatorname{Re} \left[\frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j} \right] \\
&= \operatorname{tr} \left[\mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \mathbf{C}_{\bar{x}}(\boldsymbol{\xi})}{\partial \xi_j} \right] \\
&\quad + 2 \operatorname{Re} \left[\frac{\partial \tilde{\boldsymbol{\mu}}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}_{\bar{x}}^{-1}(\boldsymbol{\xi}) \frac{\partial \tilde{\boldsymbol{\mu}}(\boldsymbol{\xi})}{\partial \xi_j} \right].
\end{aligned}$$

Appendix 1

Review of Important Concepts

A1.1 Linear and Matrix Algebra

Important results from linear and matrix algebra theory are reviewed in this section. In the discussions to follow it is assumed that the reader already has some familiarity with these topics. The specific concepts to be described are used heavily throughout the book. For a more comprehensive treatment the reader is referred to the books [Noble and Daniel 1977] and [Graybill 1969]. All matrices and vectors are assumed to be real.

A1.1.1 Definitions

Consider an $m \times n$ matrix $\mathbf{A}$ with elements a_{ij} , $i = 1, 2, \dots, m$; $j = 1, 2, \dots, n$. A shorthand notation for describing $\mathbf{A}$ is

$$[\mathbf{A}]_{ij} = a_{ij}.$$

The transpose of $\mathbf{A}$, which is denoted by $\mathbf{A}^T$, is defined as the $n \times m$ matrix with elements a_{ji} or

$$[\mathbf{A}^T]_{ij} = a_{ji}.$$

A square matrix is one for which $m = n$. A square matrix is symmetric if $\mathbf{A}^T = \mathbf{A}$.

The rank of a matrix is the number of linearly independent rows or columns, whichever is less. The inverse of a square $n \times n$ matrix is the square $n \times n$ matrix $\mathbf{A}^{-1}$ for which

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{A}\mathbf{A}^{-1} = \mathbf{I}$$

where $\mathbf{I}$ is the $n \times n$ identity matrix. The inverse will exist if and only if the rank of $\mathbf{A}$ is n . If the inverse does not exist, then $\mathbf{A}$ is singular.

The determinant of a square $n \times n$ matrix is denoted by $\det(\mathbf{A})$. It is computed as

$$\det(\mathbf{A}) = \sum_{j=1}^n a_{ij}C_{ij}$$

where

$$C_{ij} = (-1)^{i+j} M_{ij}$$

M_{ij} is the determinant of the submatrix of $\mathbf{A}$ obtained by deleting the i th row and j th column and is termed the *minor* of a_{ij} . C_{ij} is the *cofactor* of a_{ij} . Note that any choice of i for $i = 1, 2, \dots, n$ will yield the same value for $\det(\mathbf{A})$.

A *quadratic form* Q is defined as

$$Q = \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i x_j$$

In defining the quadratic form it is assumed that $a_{ji} = a_{ij}$. This entails no loss in generality since any quadratic function may be expressed in this manner. Q may also be expressed as

$$Q = \mathbf{x}^T \mathbf{A} \mathbf{x}$$

where $\mathbf{x} = [x_1 \ x_2 \ \dots \ x_n]^T$ and $\mathbf{A}$ is a square $n \times n$ matrix with $a_{ji} = a_{ij}$ or $\mathbf{A}$ is a symmetric matrix.

A square $n \times n$ matrix $\mathbf{A}$ is *positive semidefinite* if $\mathbf{A}$ is symmetric and

$$\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$$

for all $\mathbf{x} \neq 0$. If the quadratic form is strictly positive, then $\mathbf{A}$ is *positive definite*. When referring to a matrix as positive definite or positive semidefinite, it is always assumed that the matrix is symmetric.

The *trace* of a square $n \times n$ matrix is the sum of its diagonal elements or

$$\text{tr}(\mathbf{A}) = \sum_{i=1}^n a_{ii}$$

A *partitioned* $m \times n$ matrix $\mathbf{A}$ is one that is expressed in terms of its submatrices. An example is the 2×2 partitioning

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}$$

Each "element" $\mathbf{A}_{ij}$ is a submatrix of $\mathbf{A}$. The dimensions of the partitions are given as

$$\begin{bmatrix} k \times l & k \times (n-l) \\ (m-k) \times l & (m-k) \times (n-l) \end{bmatrix}$$

A1.1.2 Special Matrices

A *diagonal matrix* is a square $n \times n$ matrix with $a_{ij} = 0$ for $i \neq j$ or all elements off the principal diagonal are zero. A diagonal matrix appears as

$$\mathbf{A} = \begin{bmatrix} a_{11} & 0 & \dots & 0 \\ 0 & a_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{nn} \end{bmatrix}$$

A diagonal matrix will sometimes be denoted by $\text{diag}(a_{11}, a_{22}, \dots, a_{nn})$. The inverse of a diagonal matrix is found by simply inverting each element on the principal diagonal.

A generalization of the diagonal matrix is the square $n \times n$ **block diagonal matrix**⁹

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{A}_{kk} \end{bmatrix}$$

in which all submatrices $\mathbf{A}_{ii}$ are square and the other submatrices are identically zero. The dimensions of the submatrices need not be identical. For instance, if $k = 2$, $\mathbf{A}_{11}$ might have dimension 2×2 while $\mathbf{A}_{22}$ might be a scalar. If all $\mathbf{A}_{ii}$ are nonsingular, then the inverse is easily found as

$$\mathbf{A}^{-1} = \begin{bmatrix} \mathbf{A}_{11}^{-1} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22}^{-1} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{A}_{kk}^{-1} \end{bmatrix}.$$

Also, the determinant is

$$\det(\mathbf{A}) = \prod_{i=1}^n \det(\mathbf{A}_{ii}).$$

A square $n \times n$ matrix is **orthogonal** if

$$\mathbf{A}^{-1} = \mathbf{A}^T.$$

For a matrix to be orthogonal the columns (and rows) must be orthonormal or if

$$\mathbf{A} = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \dots \quad \mathbf{a}_n]$$

where $\mathbf{a}_i$ denotes the i th column, the conditions

$$\mathbf{a}_i^T \mathbf{a}_j = \begin{cases} 0 & \text{for } i \neq j \\ 1 & \text{for } i = j \end{cases}$$

must be satisfied. An important example of an orthogonal matrix arises in modeling of data by a sum of harmonically related sinusoids or by a discrete Fourier series. As an example, for n even

$$\mathbf{A} = \frac{1}{\sqrt{\frac{n}{2}}} \begin{bmatrix} \frac{1}{\sqrt{2}} & 1 & \dots & \frac{1}{\sqrt{2}} & 0 & \dots & 0 \\ \frac{1}{\sqrt{2}} & \cos \frac{2\pi}{n} & \dots & \frac{1}{\sqrt{2}} \cos \frac{2\pi(\frac{n}{2})}{n} & \sin \frac{2\pi}{n} & \dots & \sin \frac{2\pi(\frac{n}{2}-1)}{n} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\sqrt{2}} & \cos \frac{2\pi(n-1)}{n} & \dots & \frac{1}{\sqrt{2}} \cos \frac{2\pi\frac{n}{2}(n-1)}{n} & \sin \frac{2\pi(n-1)}{n} & \dots & \sin \frac{2\pi(\frac{n}{2}-1)(n-1)}{n} \end{bmatrix}$$

is an orthogonal matrix. This follows from the orthogonality relationships for $i, j = 0, 1, \dots, n/2$

$$\sum_{k=0}^{n-1} \cos \frac{2\pi ki}{n} \cos \frac{2\pi kj}{n} = \begin{cases} 0 & i \neq j \\ \frac{n}{2} & i = j = 1, 2, \dots, \frac{n}{2} - 1 \\ n & i = j = 0, \frac{n}{2} \end{cases}$$

and for $i, j = 1, 2, \dots, n/2 - 1$

$$\sum_{k=0}^{n-1} \sin \frac{2\pi ki}{n} \sin \frac{2\pi kj}{n} = \frac{n}{2} \delta_{ij}$$

and finally for $i = 0, 1, \dots, n/2; j = 1, 2, \dots, n/2 - 1$

$$\sum_{k=0}^{n-1} \cos \frac{2\pi ki}{n} \sin \frac{2\pi kj}{n} = 0.$$

These orthogonality relationships may be proven by expressing the sines and cosines in terms of complex exponentials and using the result

$$\sum_{k=0}^{n-1} \exp \left(j \frac{2\pi}{n} kl \right) = n \delta_{l0}$$

for $l = 0, 1, \dots, n-1$ [Oppenheim and Schaffer 1975].

An **idempotent matrix** is a square $n \times n$ matrix which satisfies

$$\mathbf{A}^2 = \mathbf{A}.$$

This condition implies that $\mathbf{A}^l = \mathbf{A}$ for $l \geq 1$. An example is the **projection matrix**

$$\mathbf{A} = \mathbf{H}(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T$$

where $\mathbf{H}$ is an $m \times n$ full rank matrix with $m > n$.

A **square $n \times n$ Toeplitz matrix** is defined as

$$[\mathbf{A}]_{ij} = a_{i-j}$$

or

$$\mathbf{A} = \begin{bmatrix} a_0 & a_{-1} & a_{-2} & \dots & a_{-(n-1)} \\ a_1 & a_0 & a_{-1} & \dots & a_{-(n-2)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{n-1} & a_{n-2} & a_{n-3} & \dots & a_0 \end{bmatrix}. \quad (\text{A1.1})$$

Each element along a northwest-southeast diagonal is the same. If in addition, $a_{-k} = a_k$, then $\mathbf{A}$ is **symmetric Toeplitz**.

A1.1.3 Matrix Manipulation and Formulas

Some useful formulas for the algebraic manipulation of matrices are summarized in this section. For $n \times n$ matrices **A** and **B** the following relationships are useful.

$$\begin{aligned}
 (\mathbf{AB})^T &= \mathbf{B}^T \mathbf{A}^T \\
 (\mathbf{A}^T)^{-1} &= (\mathbf{A}^{-1})^T \\
 (\mathbf{AB})^{-1} &= \mathbf{B}^{-1} \mathbf{A}^{-1} \\
 \det(\mathbf{A}^T) &= \det(\mathbf{A}) \\
 \det(c\mathbf{A}) &= c^n \det(\mathbf{A}) \quad (c \text{ a scalar}) \\
 \det(\mathbf{AB}) &= \det(\mathbf{A}) \det(\mathbf{B}) \\
 \det(\mathbf{A}^{-1}) &= \frac{1}{\det(\mathbf{A})} \\
 \text{tr}(\mathbf{AB}) &= \text{tr}(\mathbf{BA}) \\
 \text{tr}(\mathbf{A}^T \mathbf{B}) &= \sum_{i=1}^n \sum_{j=1}^n [\mathbf{A}]_{ij} [\mathbf{B}]_{ij}.
 \end{aligned}$$

Also, for vectors **x** and **y** we have

$$\mathbf{y}^T \mathbf{x} = \text{tr}(\mathbf{xy}^T).$$

It is frequently necessary to determine the inverse of a matrix analytically. To do so one can make use of the following formula. The inverse of a square $n \times n$ matrix is

$$\mathbf{A}^{-1} = \frac{\mathbf{C}^T}{\det(\mathbf{A})},$$

where **C** is the square $n \times n$ matrix of cofactors of **A**. The cofactor matrix is defined by

$$[\mathbf{C}]_{ij} = (-1)^{i+j} M_{ij}$$

where M_{ij} is the minor of a_{ij} obtained by deleting the i th row and j th column of **A**.

Another formula which is quite useful is the **matrix inversion lemma**

$$(\mathbf{A} + \mathbf{BCD})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{B} (\mathbf{D} \mathbf{A}^{-1} \mathbf{B} + \mathbf{C}^{-1})^{-1} \mathbf{D} \mathbf{A}^{-1}$$

where it is assumed that **A** is $n \times n$, **B** is $n \times m$, **C** is $m \times m$, and **D** is $m \times n$ and that the indicated inverses exist. A special case known as **Woodbury's identity** results for **B** an $n \times 1$ column vector **u**, **C** a scalar of unity, and **D** a $1 \times n$ row vector $\mathbf{u}^T$. Then,

$$(\mathbf{A} + \mathbf{uu}^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1} \mathbf{uu}^T \mathbf{A}^{-1}}{1 + \mathbf{u}^T \mathbf{A}^{-1} \mathbf{u}}.$$

Partitioned matrices may be manipulated according to the usual rules of matrix algebra by considering each submatrix as an element. For multiplication of partitioned

matrices the submatrices which are multiplied together must be conformable. As an illustration, for 2×2 partitioned matrices

$$\begin{aligned} \mathbf{AB} &= \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} & \mathbf{A}_{11}\mathbf{B}_{12} + \mathbf{A}_{12}\mathbf{B}_{22} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} & \mathbf{A}_{21}\mathbf{B}_{12} + \mathbf{A}_{22}\mathbf{B}_{22} \end{bmatrix}. \end{aligned}$$

The transposition of a partitioned matrix is formed by transposing the submatrices of the matrix and applying T to each submatrix. For a 2×2 partitioned matrix

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}^T = \begin{bmatrix} \mathbf{A}_{11}^T & \mathbf{A}_{21}^T \\ \mathbf{A}_{12}^T & \mathbf{A}_{22}^T \end{bmatrix}.$$

The extension of these properties to arbitrary partitioning is straightforward. Determination of the inverses and determinants of partitioned matrices is facilitated by employing the following formulas. Let $\mathbf{A}$ be a square $n \times n$ matrix partitioned as

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} = \begin{bmatrix} k \times k & k \times (n-k) \\ (n-k) \times k & (n-k) \times (n-k) \end{bmatrix}.$$

Then,

$$\mathbf{A}^{-1} = \begin{bmatrix} (\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21})^{-1} & -(\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21})^{-1}\mathbf{A}_{12}\mathbf{A}_{22}^{-1} \\ -(\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12})^{-1}\mathbf{A}_{21}\mathbf{A}_{11}^{-1} & (\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12})^{-1} \end{bmatrix}$$

$$\begin{aligned} \det(\mathbf{A}) &= \det(\mathbf{A}_{22})\det(\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21}) \\ &= \det(\mathbf{A}_{11})\det(\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12}) \end{aligned}$$

where the inverses of $\mathbf{A}_{11}$ and $\mathbf{A}_{22}$ are assumed to exist.

A1.1.4 Theorems

Some important theorems used throughout the text are summarized in this section.

1. A square $n \times n$ matrix $\mathbf{A}$ is invertible (nonsingular) if and only if its columns (or rows) are linearly independent or, equivalently, if its determinant is nonzero. In such a case, $\mathbf{A}$ is full rank. Otherwise, it is singular.
2. A square $n \times n$ matrix $\mathbf{A}$ is positive definite if and only if

- a. it can be written as

$$\mathbf{A} = \mathbf{C}\mathbf{C}^T \tag{A1.2}$$

where $\mathbf{C}$ is also $n \times n$ and is full rank and hence invertible, or

- b. the principal minors are all positive. (The i th principal minor is the determinant of the submatrix formed by deleting all rows and columns with an index greater than i .) If $\mathbf{A}$ can be written as in (A1.2), but $\mathbf{C}$ is not full rank or the principal minors are only nonnegative, then $\mathbf{A}$ is positive semidefinite.
3. If $\mathbf{A}$ is positive definite, then the inverse exists and may be found from (A1.2) as $\mathbf{A}^{-1} = (\mathbf{C}^{-1})^T (\mathbf{C}^{-1})$.
4. Let $\mathbf{A}$ be positive definite. If $\mathbf{B}$ is an $m \times n$ matrix of full rank with $m \leq n$, then $\mathbf{BAB}^T$ is also positive definite.
5. If $\mathbf{A}$ is positive definite (positive semidefinite), then
- the diagonal elements are positive (nonnegative)
 - the determinant of $\mathbf{A}$, which is a principal minor, is positive (nonnegative).

A1.1.5 Eigendecomposition of Matrices

An eigenvector of a square $n \times n$ matrix $\mathbf{A}$ is an $n \times 1$ vector $\mathbf{v}$ satisfying

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v} \quad (\text{A1.3})$$

for some scalar λ , which may be complex. λ is the eigenvalue of $\mathbf{A}$ corresponding to the eigenvector $\mathbf{v}$. It is assumed that the eigenvector is normalized to have unit length or $\mathbf{v}^T \mathbf{v} = 1$. If $\mathbf{A}$ is symmetric, then one can always find n linearly independent eigenvectors, although they will not in general be unique. An example is the identity matrix for which any vector is an eigenvector with eigenvalue 1. If $\mathbf{A}$ is symmetric, then the eigenvectors corresponding to distinct eigenvalues are orthonormal or $\mathbf{v}_i^T \mathbf{v}_j = \delta_{ij}$, and the eigenvalues are real. If, furthermore, the matrix is positive definite (positive semidefinite), then the eigenvalues are positive (nonnegative). For a positive semidefinite matrix the rank is equal to the number of nonzero eigenvalues.

The defining relation of (A1.3) can also be written as

$$\mathbf{A} \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{bmatrix} = \begin{bmatrix} \lambda_1 \mathbf{v}_1 & \lambda_2 \mathbf{v}_2 & \dots & \lambda_n \mathbf{v}_n \end{bmatrix}$$

or

$$\mathbf{A}\mathbf{V} = \mathbf{V}\mathbf{\Lambda} \quad (\text{A1.4})$$

where

$$\begin{aligned} \mathbf{V} &= \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{bmatrix} \\ \mathbf{\Lambda} &= \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n). \end{aligned}$$

If $\mathbf{A}$ is symmetric so that the eigenvectors corresponding to distinct eigenvalues are orthonormal and the remaining eigenvectors are chosen to yield an orthonormal eigenvector set, then $\mathbf{V}$ is an orthogonal matrix. As such, its inverse is $\mathbf{V}^T$, so that (A1.4)

becomes

$$\begin{aligned} \mathbf{A} &= \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \\ &= \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^T \end{aligned}$$

Also, the inverse is easily determined as

$$\begin{aligned} \mathbf{A}^{-1} &\Rightarrow \mathbf{V}^T{}^{-1} \mathbf{\Lambda}^{-1} \mathbf{V}^{-1} \\ &= \mathbf{V} \mathbf{\Lambda}^{-1} \mathbf{V}^T \\ &= \sum_{i=1}^n \frac{1}{\lambda_i} \mathbf{v}_i \mathbf{v}_i^T \end{aligned}$$

A final useful relationship follows from (A1.4) as

$$\begin{aligned} \det(\mathbf{A}) &= \det(\mathbf{V}) \det(\mathbf{\Lambda}) \det(\mathbf{V}^{-1}) \\ &= \det(\mathbf{\Lambda}) \\ &= \prod_{i=1}^n \lambda_i \end{aligned}$$

A1.2 Probability, Random Processes, and Time Series Models

An assumption is made that the reader already has some familiarity with probability theory and basic random process theory. This chapter serves as a review of these topics. For those readers needing a more extensive treatment the text by Papoulis [1965] on probability and random processes is recommended. For a discussion of time series modeling see [Kay 1988].

A1.2.1 Useful Probability Density Functions

A probability density function (PDF) which is frequently used to model the statistical behavior of a random variable is the Gaussian distribution. A random variable x with mean μ_x and variance σ_x^2 is distributed according to a *Gaussian or normal distribution* if the PDF is given by

$$p(x) = \frac{1}{\sqrt{2\pi\sigma_x^2}} \exp \left[-\frac{1}{2\sigma_x^2} (x - \mu_x)^2 \right] \quad -\infty < x < \infty. \quad (\text{A1.5})$$

The shorthand notation $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$ is often used, where $\sim$ means “is distributed according to.” If $x \sim \mathcal{N}(0, \sigma_x^2)$, then the moments of x are

$$E(x^k) = \begin{cases} 1 \cdot 3 \cdots (k-1) \sigma_x^2 & k \text{ even} \\ 0 & k \text{ odd.} \end{cases}$$

The extension to a set of random variables or a random vector $\mathbf{x} = [x_1 \ x_2 \ \dots \ x_n]^T$ with mean

$$E(\mathbf{x}) = \boldsymbol{\mu}_x$$

and covariance matrix

$$E[(\mathbf{x} - \boldsymbol{\mu}_x)(\mathbf{x} - \boldsymbol{\mu}_x)^T] = \mathbf{C}_x$$

is the multivariate Gaussian PDF

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{n}{2}} \det^{\frac{1}{2}}(\mathbf{C}_x)} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_x)^T \mathbf{C}_x^{-1} (\mathbf{x} - \boldsymbol{\mu}_x) \right]. \quad (\text{A1.6})$$

Note that $\mathbf{C}_x$ is an $n \times n$ symmetric matrix with $[\mathbf{C}_x]_{ij} = E\{[x_i - E(x_i)][x_j - E(x_j)]\} = \text{cov}(x_i, x_j)$ and is assumed to be positive definite so that $\mathbf{C}_x$ is invertible. If $\mathbf{C}_x$ is a diagonal matrix, then the random variables are uncorrelated. In this case (A1.6) factors into the product of N univariate Gaussian PDFs of the form of (A1.5), and hence the random variables are also independent. If $\mathbf{x}$ is zero mean, then the higher order joint moments are easily computed. In particular the fourth-order moment is

$$E(x_i x_j x_k x_l) = E(x_i x_j) E(x_k x_l) + E(x_i x_k) E(x_j x_l) + E(x_i x_l) E(x_j x_k).$$

If $\mathbf{x}$ is linearly transformed as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{b}$$

where $\mathbf{A}$ is $m \times n$ and $\mathbf{b}$ is $m \times 1$ with $m \leq n$ and $\mathbf{A}$ full rank (so that $\mathbf{C}_y$ is nonsingular), then $\mathbf{y}$ is also distributed according to a multivariate Gaussian distribution with

$$E(\mathbf{y}) = \boldsymbol{\mu}_y = \mathbf{A}\boldsymbol{\mu}_x + \mathbf{b}$$

and

$$E[(\mathbf{y} - \boldsymbol{\mu}_y)(\mathbf{y} - \boldsymbol{\mu}_y)^T] = \mathbf{C}_y = \mathbf{A}\mathbf{C}_x\mathbf{A}^T.$$

Another useful PDF is the χ^2 distribution, which is derived from the Gaussian distribution. If $\mathbf{x}$ is composed of independent and identically distributed random variables with $x_i \sim \mathcal{N}(0, 1)$, $i = 1, 2, \dots, n$, then

$$y = \sum_{i=1}^n x_i^2 \sim \chi_n^2$$

where χ_n^2 denotes a χ^2 random variable with n degrees of freedom. The PDF is given as

$$p(y) = \begin{cases} \frac{1}{2^{\frac{n}{2}} \Gamma(\frac{n}{2})} y^{\frac{n}{2}-1} \exp(-\frac{1}{2}y) & \text{for } y \geq 0 \\ 0 & \text{for } y < 0 \end{cases}$$

where $\Gamma(u)$ is the gamma integral. The mean and variance of y are

$$\begin{aligned} E(y) &= n \\ \text{var}(y) &= 2n. \end{aligned}$$

A1.2.2 Random Process Characterization

A discrete random process $x[n]$ is a sequence of random variables defined for every integer n . If the discrete random process is wide sense stationary (WSS), then it has a mean

$$E(x[n]) = \mu_x$$

which does not depend on n , and an autocorrelation function (ACF)

$$r_{xx}[k] = E(x[n]x[n+k]) \quad (\text{A1.7})$$

which depends only on the lag k between the two samples and not their absolute positions. Also, the autocovariance function is defined as

$$c_{xx}[k] = E[(x[n] - \mu_x)(x[n+k] - \mu_x)] = r_{xx}[k] - \mu_x^2.$$

In a similar manner, two jointly WSS random processes $x[n]$ and $y[n]$ have a cross-correlation function (CCF)

$$r_{xy}[k] = E(x[n]y[n+k])$$

and a cross-covariance function

$$c_{xy}[k] = E[(x[n] - \mu_x)(y[n+k] - \mu_y)] = r_{xy}[k] - \mu_x\mu_y.$$

Some useful properties of the ACF and CCF are

$$\begin{aligned} r_{xx}[0] &\geq |r_{xx}[k]| \\ r_{xx}[-k] &= r_{xx}[k] \\ r_{xy}[-k] &= r_{yx}[k]. \end{aligned}$$

Note that $r_{xx}[0]$ is positive, which follows from (A1.7).

The z transforms of the ACF and CCF defined as

$$\begin{aligned} \mathcal{P}_{xx}(z) &= \sum_{k=-\infty}^{\infty} r_{xx}[k]z^{-k} \\ \mathcal{P}_{xy}(z) &= \sum_{k=-\infty}^{\infty} r_{xy}[k]z^{-k} \end{aligned}$$

lead to the definition of the power spectral density (PSD). When evaluated on the unit circle, $\mathcal{P}_{xx}(z)$ and $\mathcal{P}_{xy}(z)$ become the auto-PSD, $P_{xx}(f) = \mathcal{P}_{xx}(\exp[j2\pi f])$, and cross-PSD, $P_{xy}(f) = \mathcal{P}_{xy}(\exp[j2\pi f])$, or explicitly

$$P_{xx}(f) = \sum_{k=-\infty}^{\infty} r_{xx}[k] \exp(-j2\pi f k) \quad (\text{A1.8})$$

$$P_{xy}(f) = \sum_{k=-\infty}^{\infty} r_{xy}[k] \exp(-j2\pi f k). \quad (\text{A1.9})$$

It also follows from the definition of the cross-PSD and the property $r_{yx}[k] = r_{xy}[-k]$ that

$$P_{yx}(f) = P_{xy}^*(f).$$

The auto-PSD describes the distribution in frequency of the power of $x[n]$ and as such is real and nonnegative. The cross-PSD, on the other hand, is in general complex. The magnitude of the cross-PSD describes whether frequency components in $x[n]$ are associated with large or small amplitudes at the same frequency in $y[n]$, and the phase of the cross-PSD indicates the phase lag or lead of $x[n]$ with respect to $y[n]$ for a given frequency component. Note that both spectral densities are periodic with period 1. The frequency interval $-1/2 \leq f \leq 1/2$ will be considered the fundamental period. When there is no confusion, $P_{xx}(f)$ will simply be referred to as the *power spectral density*.

A process which will frequently be encountered is discrete white noise. It is defined as a zero mean process having an ACF

$$r_{xx}[k] = \sigma^2 \delta[k]$$

where $\delta[k]$ is the discrete impulse function. This says that each sample is uncorrelated with all the others. Using (A1.8), the PSD becomes

$$P_{xx}(f) = \sigma^2$$

and is observed to be completely flat with frequency. Alternatively, white noise is composed of equi-power contributions from all frequencies.

For a linear shift invariant (LSI) system with impulse response $h[n]$ and with a WSS random process input, various relationships between the correlations and spectral density functions of the input process $x[n]$ and output process $y[n]$ hold. The correlation relationships are

$$\begin{aligned} r_{xy}[k] &= h[k] \star r_{xx}[k] = \sum_{l=-\infty}^{\infty} h[l] r_{xx}[k-l] \\ r_{yx}[k] &= h[-k] \star r_{xx}[k] = \sum_{l=-\infty}^{\infty} h[-l] r_{xx}[k-l] \\ r_{yy}[k] &= h[k] \star r_{yx}[k] = h[k] \star h[-k] \star r_{xx}[k] \\ &= \sum_{m=-\infty}^{\infty} h[k-m] \sum_{l=-\infty}^{\infty} h[-l] r_{xx}[m-l] \end{aligned}$$

where $\star$ denotes convolution. Denoting the system function by $\mathcal{H}(z) = \sum_{n=-\infty}^{\infty} h[n]z^{-n}$, the following relationships for the PSDs follow from these correlation properties.

$$\begin{aligned} \mathcal{P}_{xy}(z) &= \mathcal{H}(z) \mathcal{P}_{xx}(z) \\ \mathcal{P}_{yx}(z) &= \mathcal{H}(1/z) \mathcal{P}_{xx}(z) \\ \mathcal{P}_{yy}(z) &= \mathcal{H}(z) \mathcal{H}(1/z) \mathcal{P}_{xx}(z). \end{aligned}$$

In particular, letting $H(f) = \mathcal{H}[\exp(j2\pi f)]$ be the frequency response of the LSI system results in

$$\begin{aligned}P_{xy}(f) &= H(f)P_{xx}(f) \\P_{yx}(f) &= H^*(f)P_{xx}(f) \\P_{yy}(f) &= |H(f)|^2 P_{xx}(f).\end{aligned}$$

For the special case of a white noise input process the output PSD becomes

$$P_{yy}(f) = |H(f)|^2 \sigma^2 \quad (\text{A1.10})$$

since $P_{xx}(f) = \sigma^2$. This will form the basis of the time series models to be described shortly.

A1.2.3 Gaussian Random Process

A Gaussian random process is one for which any samples $x[n_0], x[n_1], \dots, x[n_{N-1}]$ are jointly distributed according to a multivariate Gaussian PDF. If the samples are taken at successive times to generate the vector $\mathbf{x} = [x[0] x[1] \dots x[N-1]]^T$, then assuming a zero mean WSS random process, the covariance matrix takes on the form

$$\mathbf{C}_x = \begin{bmatrix} r_{xx}[0] & r_{xx}[1] & \dots & r_{xx}[N-1] \\ r_{xx}[1] & r_{xx}[0] & \dots & r_{xx}[N-2] \\ \vdots & \vdots & \ddots & \vdots \\ r_{xx}[N-1] & r_{xx}[N-2] & \dots & r_{xx}[0] \end{bmatrix} \quad \mathbf{x}$$

The covariance matrix, or more appropriately the autocorrelation matrix, has the special symmetric Toeplitz structure of (A1.1) with $a_k = a_{-k}$.

An important Gaussian random process is the white process. As discussed previously, the ACF for a white process is a discrete delta function. In light of the definition of a Gaussian random process a white Gaussian random process $x[n]$ with mean zero and variance σ^2 is one for which

$$x[n] \sim N(0, \sigma^2) \quad -\infty < n < \infty$$

$$r_{xx}[m-n] = E(x[n]x[m]) = 0 \quad m \neq n.$$

Because of the Gaussian assumption, all the samples are statistically independent.

A1.2.4 Time Series Models

A useful class of time series models consists of the rational transfer function models. The time series $x[n]$ is modeled as the output of a LSI filter with frequency response $H(f)$ excited at the input by white noise $u[n]$ with variance σ_u^2 . As such, its PSD follows from (A1.10) as

$$P_{xx}(f) = |H(f)|^2 \sigma_u^2.$$

The first time series model is termed an *autoregressive* (AR) process, which has the time domain representation

$$x[n] = - \sum_{k=1}^p a[k]x[n-k] + u[n].$$

It is said to be an AR process of order p and is denoted by $\text{AR}(p)$. The AR parameters consist of the filter coefficients $\{a[1], a[2], \dots, a[p]\}$ and the driving white noise variance σ_u^2 . Since the frequency response is

$$H(f) = \frac{1}{1 + \sum_{k=1}^p a[k] \exp(-j2\pi f k)}$$

the AR PSD is

$$P_{xx}(f) = \frac{\sigma_u^2}{\left| 1 + \sum_{k=1}^p a[k] \exp(-j2\pi f k) \right|^2}.$$

It can be shown that the ACF satisfies the recursive difference equation

$$r_{xx}[k] = \begin{cases} -\sum_{l=1}^p a[l] r_{xx}[k-l] & k \geq 1 \\ -\sum_{l=1}^p a[l] r_{xx}[l] + \sigma_u^2 & k = 0 \end{cases}$$

In matrix form this becomes for $k = 1, 2, \dots, p$

$$\begin{bmatrix} r_{xx}[0] & r_{xx}[1] & \dots & r_{xx}[p-1] \\ r_{xx}[1] & r_{xx}[0] & \dots & r_{xx}[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ r_{xx}[p-1] & r_{xx}[p-2] & \dots & r_{xx}[0] \end{bmatrix} \begin{bmatrix} a[1] \\ a[2] \\ \vdots \\ a[p] \end{bmatrix} = - \begin{bmatrix} r_{xx}[1] \\ r_{xx}[2] \\ \vdots \\ r_{xx}[p] \end{bmatrix}$$

and also

$$\sigma_u^2 = r_{xx}[0] + \sum_{k=1}^p a[k] r_{xx}[k].$$

Given the ACF samples $r_{xx}[k]$ for $k = 0, 1, \dots, p$, the AR parameters may be determined by solving the set of p linear equations. These equations are termed the *Yule-Walker equations*. As an example, for an $\text{AR}(1)$ process

$$r_{xx}[k] = -a[1]r_{xx}[k-1] \quad k \geq 1$$

which can be solved to yield

$$r_{xx}[k] = r_{xx}[0] (-a[1])^k \quad k \geq 0$$

and from

$$\sigma_u^2 = r_{xx}[0] + a[1]r_{xx}[1]$$

we can solve for $r_{xx}[0]$ to produce the ACF of an AR(1) process

$$r_{xx}[k] = \frac{\sigma_u^2}{1 - a^2[1]} (-a[1])^{|k|}.$$

The corresponding PSD is

$$P_{xx}(f) = \frac{\sigma_u^2}{|1 + a[1] \exp(-j2\pi f)|^2}.$$

The AR(1) PSD will have most of its power at lower frequencies if $a[1] < 0$, and at higher frequencies if $a[1] > 0$. The system function for the AR(1) process is

$$\mathcal{H}(z) = \frac{1}{1 + a[1]z^{-1}}$$

and has a pole at $z = -a[1]$. Hence, for a stable process we must have $|a[1]| < 1$.

While the AR process is generated as the output of a filter having only poles, the *moving average* (MA) process is formed by passing white noise through a filter whose system function has only zeros. The time domain representation of a MA(q) process is

$$x[n] = u[n] + \sum_{k=1}^q b[k]u[n-k].$$

Since the filter frequency response is

$$H(f) = 1 + \sum_{k=1}^q b[k] \exp(-j2\pi f k)$$

its PSD is

$$P_{xx}(f) = \left| 1 + \sum_{k=1}^q b[k] \exp(-j2\pi f k) \right|^2 \sigma_u^2.$$

The MA(q) process has the ACF

$$r_{xx}[k] = \begin{cases} \sigma_u^2 \sum_{l=0}^{q-|k|} b[l]b[l+|k|] & |k| \leq q \\ 0 & |k| > q. \end{cases}$$

The system function is

$$\mathcal{H}(z) = 1 + \sum_{k=1}^q b[k]z^{-k}$$

and is seen to consist of zeros only. Although not required for stability, it is usually assumed that the zeros satisfy $|z_i| < 1$. This is because the zeros z_i and $1/z_i^*$ can both result in the same PSD, resulting in a problem of identifiability of the process parameters.

References

- Graybill, F.A., *Introduction to Matrices with Application in Statistics*, Wadsworth, Belmont, Calif., 1969.
- Kay, S., *Modern Spectral Estimation: Theory and Application*, Prentice-Hall, Englewood Cliffs, N.J., 1988.
- Noble, B., J.W. Daniel, *Applied Linear Algebra*, Prentice-Hall, Englewood Cliffs, N.J., 1977.
- Oppenheim, A.V., R.W. Schaffer, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1975.
- Papoulis, A., *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, 1965.

Appendix 2

Glossary of Symbols and Abbreviations

Symbols

(Boldface characters denote vectors or matrices. All others are scalars.)

$*$	complex conjugate
$\star$	convolution
$\hat{\cdot}$	denotes estimator
$\cdot$	denotes estimator
$\sim$	denotes <i>is distributed according to</i>
$\stackrel{\sim}{\sim}$	denotes <i>is asymptotically distributed according to</i>
$\arg \max_{\theta} g(\theta)$	denotes the value of θ that maximizes $g(\theta)$
$[\mathbf{A}]_{ij}$	ij th element of $\mathbf{A}$
$[\mathbf{b}]_i$	i th element of $\mathbf{b}$
$\text{Bmse}(\hat{\theta})$	Bayesian mean square error of $\hat{\theta}$
χ_n^2	chi-squared distribution with n degrees of freedom
$\text{cov}(x, y)$	covariance of x and y
$\mathbf{C}_x$ or $\mathbf{C}_{xx}$	covariance matrix of $\mathbf{x}$
$\mathbf{C}_{xy}$	covariance matrix of $\mathbf{x}$ and $\mathbf{y}$
$\mathbf{C}_{y x}$	covariance matrix of $\mathbf{y}$ with respect to PDF of $\mathbf{y}$ conditioned on $\mathbf{x}$
$\mathcal{CN}(\tilde{\mu}, \sigma^2)$	complex normal distribution with mean $\tilde{\mu}$ and variance σ^2
$\mathcal{CN}(\tilde{\mu}, \mathbf{C})$	multivariate complex normal distribution with mean $\tilde{\mu}$ and covariance $\mathbf{C}$

$\delta(t)$	Dirac delta function
$\delta[n]$	discrete-time impulse sequence
δ_{ij}	Kronecker delta
Δ	time sampling interval
$\det(\mathbf{A})$	determinant of matrix $\mathbf{A}$
$\text{diag}(\cdots)$	diagonal matrix with elements $\cdots$ on main diagonal
$\mathbf{e}_i$	natural unit vector in i th direction
E	expected value
E_x	expected value with respect to PDF of $\mathbf{x}$
$E_{x \theta}$ or $E(x \theta)$	conditional expected value with respect to PDF of x conditioned on θ
$\mathcal{E}$	energy
η	signal-to-noise-ratio
f	discrete-time frequency
F	continuous-time frequency
$\mathcal{F}$	Fourier transform
$\mathcal{F}^{-1}$	inverse Fourier transform
H	conjugate transpose
$\mathbf{H}$	observation matrix
(x, y)	inner product of x and y
$i(\theta)$	Fisher information for single data sample and scalar θ
$I(\theta)$	Fisher information for scalar θ
$\mathbf{I}$	identity matrix
$\mathbf{I}(\theta)$	Fisher information matrix for vector θ
$I(f)$	periodogram
$\text{Im}(\)$	imaginary part of
j	$\sqrt{-1}$
$\text{mse}(\hat{\theta})$	mean square error of $\hat{\theta}$ (classical)
$\mathbf{M}_{\hat{\theta}}$	mean square error matrix of $\hat{\theta}$
μ	mean
n	sequence index
N	length of observed data set
$\mathcal{N}(\mu, \sigma^2)$	normal distribution with mean μ and variance σ^2
$\mathcal{N}(\boldsymbol{\mu}, \mathbf{C})$	multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance $\mathbf{C}$

$\ x\ $	norm of x
$\mathbf{1}$	vector of all ones
$p(x)$ or $p_x(x)$	probability density function of x
$p(\mathbf{x}; \theta)$	probability density function of $\mathbf{x}$ with θ as a parameter
$p(\mathbf{x} \theta)$	conditional probability density function of $\mathbf{x}$ conditioned on θ
$\mathbf{P}$	projection matrix
$\mathbf{P}^\perp$	orthogonal projection matrix
$\frac{\partial}{\partial \mathbf{x}}$	gradient vector with respect to $\mathbf{x}$
$\frac{\partial^2}{\partial \mathbf{x} \partial \mathbf{x}^T}$	Hessian matrix with respect to $\mathbf{x}$
$\Pr\{ \}$	probability
$P_{xx}(f)$	power spectral density of discrete-time process $x[n]$
$P_{xy}(f)$	cross-spectral density of discrete-time processes $x[n]$ and $y[n]$
$P_{xx}(F)$	power spectral density of continuous-time process $x(t)$
ρ	correlation coefficient
$r_{xx}[k]$	autocorrelation function of discrete-time process $x[n]$
$r_{xx}(\tau)$	autocorrelation function of continuous-time process $x(t)$
$r_{xy}[k]$	cross-correlation function of discrete-time processes $x[n]$ and $y[n]$
$r_{xy}(\tau)$	cross-correlation function of continuous-time processes $x(t)$ and $y(t)$
$\mathbf{R}_{xx}$	autocorrelation matrix of $\mathbf{x}$
$\text{Re}(\)$	real part
σ^2	variance
$s[n]$	discrete-time signal
$\mathbf{s}$	vector of signal samples
$s(t)$	continuous-time signal
t	continuous time
$\text{tr}(\mathbf{A})$	trace of matrix $\mathbf{A}$
$\theta(\theta)$	unknown parameter (vector)
$\hat{\theta}(\hat{\theta})$	estimator of $\theta(\theta)$
T	transpose
$\mathcal{U}[a, b]$	uniform distribution over the interval $[a, b]$
$\text{var}(x)$	variance of x
$\text{var}(x \theta)$	variance of conditional PDF or of $p(x \theta)$
$w[n]$	observation noise sequence

$\mathbf{w}$	vector of noise samples
$w(t)$	continuous-time noise
$x[n]$	observed discrete-time data
$\mathbf{x}$	vector of data samples
$x(t)$	observed continuous-time waveform
$\bar{x}$	sample mean of x
$\mathcal{Z}$	z transform
$\mathcal{Z}^{-1}$	inverse z transform
$\mathbf{0}$	vector or matrix of all zeros

Abbreviations

ACF	autocorrelation function
ANC	adaptive noise canceler
AR	autoregressive
AR(p)	autoregressive process of order p
ARMA	autoregressive moving average
BLUE	best linear unbiased estimator
CCF	cross-correlation function
CRLB	Cramer-Rao lower bound
CWGN	complex white Gaussian noise
DC	constant level (direct current)
DFT	discrete Fourier transform
EM	expectation-maximization
FFT	fast Fourier transform
FIR	finite impulse response
IID	independent and identically distributed
IIR	infinite impulse response
LMMSE	linear minimum mean square error
LPC	linear predictive coding
LS	least squares
LSE	least squares estimator
LSI	linear shift invariant

MA	moving average
MAP	maximum a posteriori
MLE	maximum likelihood estimator
MMSE	minimum mean square error
MSE	mean square error
MVDR	minimum variance distortionless response
MVU	minimum variance unbiased
OOK	on-off keyed
PDF	probability density function
PRN	pseudorandom noise
PSD	power spectral density
RBLS	Rao-Blackwell-Lehmann-Scheffe
SNR	signal-to-noise ratio
TDL	tapped delay line (same as FIR)
TDOA	time difference of arrival
WGN	white Gaussian noise
WSS	wide sense stationary

Index

- ACF (*see* Autocorrelation)
- Adaptive beamforming, 544–48
- Adaptive filters: (*see also* Least squares, sequential)
 - Kalman, 439
 - noise canceler, 268–73
- Analytic signal, 497, 551
- AR (*see* Autoregressive)
- ARMA (*see* Autoregressive moving average)
- Asymptotic:
 - Cramer-Rao lower bound, 51, 77–81
 - efficiency, 38–39, 160, 164
 - Gaussian PDF:
 - real, 80
 - complex, 535
 - mean and variance, 295, 301–2, 305–6
 - MLE, 190
 - probability density function of MLE, 164
 - unbiasedness, 38, 160
- Autocorrelation:
 - definition, 575
 - estimator, 197, 204, 267
- Autocorrelation method of linear prediction, 198
- Autoregressive: (*see also* Linear predictive coding)
 - CRLB, 59–62
 - definition, 59–60, 578
 - MLE, 196–98
 - power spectral density, complex process, 497–98
 - prediction, 414
 - Autoregressive moving average:
 - definition, 266
 - dynamic model, 468
 - estimation, 266–68
- Beamforming, conventional, 547
- Bearing estimation, 3, 57–59, 195–96
- Bernoulli trial, 123, 200
- Best linear unbiased estimator:
 - complex data, 523–24
 - covariance errors, 150
 - definition, 134, 137, 139–40
 - derivation, 151–55
 - linear model, 141
 - transformations, 135, 147, 149–50
- Bias error, 18
- Biomedical signal processing, 23
- BLUE (*see* Best linear unbiased estimator)
- CCF (*see* Cross-correlation)
- Chirp rate estimator, 553
- Communications:
 - channel equalization, 365
 - coherent demodulation, 273
 - on-off keying, 148
- Complete sufficient statistic, 109–12, 119
- Complex envelope, 494
- Conditional mean estimator (*see* Minimum mean square error estimator, Bayesian)
- Consistency, estimator, 24, 161, 200
- Correlation coefficient:
 - conditional Gaussian PDF, 323

- Correlation coefficient (*Contd.*):
 - CRLB, 66
 - definition, 64
 - least squares, 241
 - MLE, 200, 304
- Correlation time, 50, 77–78, 535
- Correlator, signal, 192
- Cost function, 342
- Covariance matrix, complex:
 - definition, 501
 - properties, 505–6, 555–57
- Cramer-Rao lower bound:
 - asymptotic, 51, 77–81
 - complex Gaussian, 525
 - definition, 22, 30, 39–40, 44
 - Gaussian PDF, 47–48
 - signals in WGN, 36, 48
 - transformed parameters, 37, 45
- CRLB (*see* Cramer-Rao lower bound)
- Cross-correlation, 514, 575
- Cross-power spectral density, 576–77
- Curve fitting:
 - CRLB, 65
 - least squares, 232–35
 - linear model, 86–88
- CWGN (*see* White Gaussian noise, complex)
- Cyclical data (*see* Sinusoidal estimation)

- DC level in noise: (*see* Examples)
 - definition, 31
- Deconvolution, 365–70
- Derivative, complex, 499–500, 517, 519–21
- Detection:
 - jump in level, 278
 - sinusoidal, 98–99, 148–49, 554
- DFT (*see* Discrete Fourier transform)
- Digital filter design:
 - equation error, 261–65
 - least squares, 280–81
- Discrete Fourier transform:
 - normalization of, 511
 - orthogonality, 89, 569–70
 - PDF for WGN, 509–11, 537
- Dispersive channel, 452

- Efficiency, estimator, 34, 38–39, 84–86, 160, 167, 187, 528
- Eigenanalysis of covariance matrix, 147–48, 537
- Eigenvalue/eigenvector, 573
- EM (*see* Expectation-maximization)
- Entropy, 336
- Equation error modeling, 266
- Error ellipse, 364
- Estimators:
 - classical vs. Bayesian, 8, 309, 312
 - combining, 17
 - definition, 9
 - performance, 9–12, 24, 295 (*see also* Monte Carlo method and Asymptotic, mean and variance)
 - selection, rationale for, 489–90
 - summary:
 - classical, 480–83
 - Bayesian, 484–85
- Examples:
 - adaptive beamformer, 544–48
 - adaptive noise canceler, 268–73
 - autoregressive parameters, CRLB, 59–62
 - autoregressive parameters, MLE, 196–98
 - autoregressive parameters in ARMA, LSE, 266–68
 - bandpass Gaussian noise, 515–17
 - bearing, CRLB, 57–59
 - bearing, MLE, 195–96
 - channel estimation, 452–56
 - covariance matrix scale factor, Bayesian estimation, 329–30
 - curve fitting, MVU estimator, 86–88
 - DC level in colored noise, complex BLUE, 523–24
 - DC level in colored noise, MVU estimator, 95–96
 - DC level and exponential in WGN, MVU estimator, 96–97
 - DC level in noise, LSE, 221
 - DC level in non-Gaussian noise, 172–73
 - DC level in uncorrelated noise, BLUE, 138–39
 - DC level in WGN, amplitude and variance, MAP estimator, 355–58
 - DC level in WGN, amplitude/variance, MLE, 158–163
 - DC level in WGN, amplitude and variance sufficient statistics, 118
 - DC level in WGN, biased estimator, 17
 - DC level in WGN, CRLB for amplitude, 31–32

Examples (*Contd.*):

- DC level in WGN, CRLB for amplitude and variance, 40-41
- DC level in WGN, CRLB for random amplitude variance, 49-50
- DC level in WGN, Gaussian prior, MMSE estimator, 317-21, 326-28, 360-61
- DC level in WGN, method of moments, 291-92
- DC level in WGN, MLE for amplitude, 163-64
- DC level in WGN, MLE for amplitude and variance, 183
- DC level in WGN, MLE Monte Carlo performance, 164-66
- DC level in WGN, MVU amplitude estimator from sufficient statistic, 107-109
- DC level in WGN, MVU amplitude and variance estimator from sufficient statistic, 119-22
- DC level in WGN, sequential LMMSE estimator, 392-93
- DC level in WGN, sequential LSE, 243-48
- DC level in WGN, sufficient statistic, 105
- DC level in WGN, transformed parameter MLE, 173-77
- DC level in WGN, unbiased estimator, 16
- DC level in WGN, uniform prior, LMMSE estimator, 383
- DC level in WGN, uniform prior, MAP estimator, 352-53
- DC level in WGN, uniform prior, MMSE estimator, 315
- DC level in white noise, BLUE, 137-38
- digital filter design, LSE, 261-65
- discrete Fourier transform, PDF of CWGN, 535-37
- discrete Fourier transform, PDF of WGN, 509-11
- exponential PDF parameter, MAP estimator, 351-52
- exponential PDF parameter, method of moments, 292, 295-97
- exponential PDF parameter transformation, MAP estimator, 358-59
- exponential signal, LSE, 257-58
- exponential signal in WGN, MLE, 178-82
- exponential signal in white noise, ad-hoc estimator, 298-99

Examples (*Contd.*):

- Fourier analysis, Bayesian, 347-49, 362-64, 399-400
- Fourier analysis, LSE, 226-27, 230-31
- Fourier analysis, MVU estimator, 88-90
- Fourier analysis, sequential LSE, 250-51
- frequencies of sinusoids, EM estimator, 187-89
- frequency of sinusoid, CRLB, 36
- frequency of sinusoid, method of moments, 299-304
- frequency of WSS process, center, CRLB, 51-53
- Gaussian mixture parameters, 290-91, 293-94
- Gauss-Markov model, 427-28
- Hermitian form, mean and variance, 512-13
- Hermitian function, minimization, 521-23
- identification of FIR system, MVU estimator, 90-94
- Kalman filter, 436-38, 443-45
- linear model, classical complex, 529-30
- line fitting, CRLB, 41-43
- line fitting, order-recursive LSE, 237-40
- localization, source, BLUE, 142-46
- mean of uniform noise, MVU estimator, 113-16
- moving average, MLE, 190-91
- MVU estimator, possible nonexistence of, 20-21
- orthogonal random variables, LMMSE estimator, 388-89
- PDF parameter dependence, 28-31
- periodogram spectral estimation, 538-39
- phase-locked loop, 273-75
- phase of complex sinusoid, MLE, 531-32
- phase of sinusoid, CRLB, 33-34
- phase of sinusoid, MLE, 167-72
- phase of sinusoid, sufficient statistic, 106-7
- power of noise, CRLB, 49
- power of noise, sufficient statistic, 105
- range, CRLB, 53-56
- range, MLE, 192
- signal, constrained LSE, 252-54
- signal amplitude estimation, complex LSE, 498-500
- signal in non-Gaussian noise, MLE, 184-85
- signal in WGN, CRLB, 48
- signal-to-noise ratio, CRLB, 46

Examples (*Contd.*):

- sinusoidal amplitude, LSE, 255–56
 - sinusoidal complex amplitude, MMSE estimator, 534–35
 - sinusoidal modeling, complex, 496–97
 - sinusoidal parameters, complex MLE, 539–44
 - sinusoidal parameters, CRLB, 56–57
 - sinusoidal parameters, LSE, 222–23
 - sinusoidal parameters, MLE, 193–95
 - sinusoidal parameters, sufficient statistics, 117–18
 - sinusoidal power, complex MVU estimator, 525–27
 - sufficient statistic, completeness of, 110–11
 - sufficient statistic, incompleteness of, 111–12
 - sufficient statistic verification, 103–4
 - vehicle tracking, 456–66
 - Wiener filtering, 365–70, 400–409, 443–45
- Expectation-maximization, 182, 187–89
- Exponential PDF family:
- definition, *see* Probability density functions
 - MLE, 200
- Exponential signals:
- estimation, 257–58, 298–99
- Fading signal, 100, 452
- Finite impulse response filter, 90–94
- FIR (*see* Finite impulse response filter)
- Fisher information:
- decoupled matrix, 41, 65
 - definition, 34, 40
 - properties, 35, 65
- Fourier analysis, 88–90, 226–27, 250–51, 347–49, 362–64, 399–400
- Frequency estimation (*see* Sinusoidal estimation and Examples)
- Gaussian random process, 467, 513, 577–78
- Gauss-Markov process:
- definition, 421, 426, 430–31
 - properties, 424, 429
- Gauss-Markov theorem, 141, 143, 552
- Gauss-Newton iteration, 260
- Gradient formulas, 73–74, 84, 519–21
- Gram-Schmidt orthogonalization, 236, 396, 411

Grid search, 177

Hermitian form:

- definition, 502
 - minimization, 521–23
 - moments, 502–3, 513
- Histogram, 10, 165, 206–7, 209

Image signal processing, 365

Innovations, 396, 433, 441

In-phase signal, 495–96

Interference suppression, 270

Interpolation, 412

Kalman filter:

- definition, 436, 446–49, 455
- derivation, 471–75
- extended, 451–52, 462, 476–77
- gain, 436, 447
- information form, 449
- steady state, 443

Least squares:

- BLUE, relationship with, 225
- constrained, 252
- definition, 220–21
- estimator, 225
- modified Yule-Walker equations, 268
- nonlinear, 222, 254
- numerical determination, 259–60
- order-recursive, 237, 282–84
- separable, 222–23, 256–57
- sequential, 249, 279, 286–88
- weighted, 150, 225–26, 244–48, 270

Levinson recursion, 198, 403

Likelihood function:

- definition, 29
- modified, 175, 185

Linear minimum mean square error estimator:

- definition, 380–82, 389
- properties, 390
- sequential, 393, 398, 415–18
- vector space interpretation, 386

Linear model (Bayesian):

- definition, 325
- Kalman filter modeling, 447
- MMSE estimator, 364–65, 533–34
- properties, 487–89

- Linear model (classical):
 - CRLB, 85
 - definition, 84, 94–95, 97, 529–30
 - efficiency, 85–86
 - estimator and properties, 85, 486–88
 - line fitting, 45
 - MLE, 186
 - reduced, 99, 254
 - sufficient statistics, 126
- Linear predictive coding, 5, 59, 198, 407
- Linear random process, 77
- Line arrays, 58, 145
- Line fitting, 41, 83–84, 237–40, 373
- LMMSE (*see* Linear minimum mean square error estimator)
- Localization, source, 142–46, 456–66
- LPC (*see* Linear predictive coding)
- LS, LSE (*see* Least squares)
- Lyapunov equation, 430
- MA (*see* Moving average)
- MAP (*see* Maximum *a posteriori* estimator)
- Matrix:
 - autocorrelation, 62, 93
 - determinant, 567
 - diagonal, 568–69
 - eigenanalysis, 573
 - Hermitian, 501
 - idempotent, 194, 570
 - ill-conditioned, 85, 98, 240–41
 - inversion:
 - definition, 567
 - lemma, 571
 - Woodbury's identity, 571
 - orthogonal, 569
 - partitioned, 571–72
 - positive definite (semidefinite), 568, 572
 - projection, 231, 242, 277, 285
 - square, 567
 - symmetric, 567
 - Toeplitz, 62, 93, 570
 - trace, 568
 - transpose, 567
- Maximum *a posteriori* estimator:
 - definition, 344, 351, 354, 372
 - properties, 358, 372
- Maximum likelihood estimator:
 - asymptotic, 190
 - Bayesian, 352
 - Maximum likelihood estimator (*Contd.*):
 - complex data, 530–31, 563–65
 - definition, 162, 182
 - efficiency, 164, 187
 - Gaussian PDF, 185
 - invariance, 174–76, 185
 - numerical determination, 177–82, 187–89
 - probability density function, asymptotic, 167, 183, 211–13
 - properties, asymptotic, 172, 201–2
- Mean square bandwidth, 55
- Mean square error:
 - Bayesian, 311, 320, 347, 533
 - classical, 19
- Mean square error matrix, 361–62, 390
- Minimal sufficient statistic, 102, 117
- Minimum mean square error estimator:
 - Bayesian:
 - definition, 313, 316, 346
 - performance, 360, 364–65, 534
 - properties, 349–50
 - classical, 19, 311
- Minimum variance distortionless response, 546
- Minimum variance unbiased estimator:
 - definition, 20
 - determination of, 109, 112–13
 - linear model, 85–86
- MLE (*see* Maximum likelihood estimator)
- MMSE (*see* Minimum mean square error estimator)
- Modeling:
 - dynamical signal, 421
 - identifiability, 85
 - least squares, 232–34
 - linearization, 143, 259, 273, 451, 461
 - speech spectrum, 5 (*see also* Autoregressive and Linear predictive coding)
- Moments, method of:
 - definition, 293
 - exponential parameter, estimator, 292, 295–97
 - Gaussian mixture, 290–91, 293–94
- Monte Carlo method, 10, 164–167, 205–10
- Moving average:
 - asymptotic MLE, 190–91
 - definition, 580
- MSE (*see* Mean square error)

- MVU (*see* Minimum variance unbiased estimator)
- Narrowband representation, 495
- Newton-Raphson iteration, 179–82, 187, 259
- Neyman-Fisher factorization, 104–5, 117, 127–29
- Normal equations, 225, 387
- Notational conventions, 13 (*see also* Appendix 2)
- Nuisance parameters, 329
- Observation equation, 446
- Observation matrix, 84, 100, 140, 224
- Order statistics, 114
- Orthogonality, 89, 385 (*see also* Projection theorem, orthogonal)
- Outliers, 170
- PDF (*see* Probability density functions)
- Periodogram, 80, 190, 195, 197, 204 (*see also* Spectral estimation)
- Phase-locked loop, 273–75
- Posterior PDF:
 - Bayesian linear model, 326, 533
 - definition, 313, 317
- Power estimation, random process, 66, 203, 553–54
- Power spectral density, 576–77
- Prediction:
 - Kalman, 440–41, 469–70
 - Wiener, 400
- Prior PDF:
 - conjugate, 335 (*see also* Reproducing PDF)
 - definition, 313
 - noninformative, 332, 336
- Probability density functions:
 - chi-squared, 122, 575
 - complex Gaussian:
 - conditional, 508–9, 562
 - definition, 503–4, 507
 - properties, 508–9, 550, 558–62
 - exponential, 122
 - exponential family, 110, 124
 - gamma, inverted, 329–30, 355
 - Gaussian, 574
 - Gaussian, conditional, 323–25, 337–39
 - Gaussian mixture, 150
 - Probability density functions (*Contd.*):
 - Laplacian, 63
 - lognormal, 147
 - Rayleigh, 122, 371
 - Processing gain, 554
 - Projection theorem, orthogonal, 228–29, 386
 - Prony method, 264
 - PSD (*see* Power spectral density)
 - Pseudorandom noise, 92, 165, 206
 - Pythagorean theorem, least squares, 276
- Quadratic form:
 - definition, 568
 - moments, 76
- Quadrature signal, 495–96
- Radar signal processing, 1
- Random number generator (*see* Pseudorandom noise)
- Random variable, complex, 500–501
- Range estimation, 1, 14, 53–56, 192
- Rao-Blackwell-Lehmann-Scheffe theorem, 22, 109, 118–19, 130–31
- Rayleigh fading, 347
- RBLs (*see* Rao-Blackwell-Lehmann-Scheffe theorem)
- Regression, nonlinear, 254
- Regularity conditions, 30, 44, 63, 67, 70
- Reproducing PDF, 321, 334–35
- Ricatti equation, 443
- Risk, Bayes, 342
- Sample mean estimator, 115, 121, 164 (*see also* DC level in noise)
- Sample variance estimator, 121, 164
- Scoring, 180, 187
- Seismic signal processing, 365
- Separability, least squares, 222–23, 256
- Signal amplitude estimator, 136, 498–500
- Sinusoidal estimation:
 - amplitudes, 88–90
 - complex data, 525–27, 531–32, 534–35, 543
 - CRLB for frequency, 36
 - CRLB for parameters, 56–57, 542
 - CRLB for phase, 33
 - EM for frequency, 187–89
 - least squares for parameters, 255–56
 - method of moments for frequency, 300, 306
 - MLE for parameters, 193–95, 203–4

- Sinusoidal estimation (*Contd.*):
 - phase estimator, 123, 167–72
 - sufficient statistics, 117–18
- Sinusoidal modeling, complex, 496
- Slutsky's theorem, 201
- Smoothing, Wiener, 400
- Sonar signal processing, 2
- Spatial frequency, 58, 195
- Spectral estimation:
 - autoregressive, 60
 - Fourier analysis, 88–90
 - periodogram, 204, 538–39, 543, 552
- Speech recognition, 4
- State transition matrix, 426
- State vector, 424
- Statistical linearization, 39, 200 (*see also* Modeling)
- Sufficient statistic, 22, 102–3, 107, 116
- System identification:
 - nonrandom FIR, 90–94, 99
 - random FIR, 452–55
- Tapped delay line (*see* FIR)
- Threshold effect, 170
- Time delay estimation, 53–56, 142–46
- Time difference of arrival, 142
- Time series, 6
- Tracking:
 - frequency, 470 (*see also* Phase-locked loop)
 - vehicle position, 456–66
- Unbiased estimator, 16, 22
- Vector spaces:
 - least squares, 227–30
 - random variables, 384
- Wavenumber (*see* Spatial frequency)
- WGN (*see* White Gaussian noise)
- White Gaussian noise:
 - real, 7
 - complex, 517
- Whitening:
 - Kalman, 441, 444
 - matrix transformation, 94–96
- White noise, 576
- Wide sense stationary, 575
- Wiener filtering, 365–70, 373–74, 379, 400–409, 443
- Wiener-Hopf equations:
 - filtering, 403
 - prediction, 406–7
- WSS (*see* Wide sense stationary)
- Yule-Walker equations:
 - AR, 198, 579
 - ARMA, 267