

Statistical Signal Processing

Collection Editor:

Clayton Scott

Statistical Signal Processing

Collection Editor:

Clayton Scott

Authors:

Richard Baraniuk
Michael Haag
Don Johnson

Robert Nowak
Clayton Scott
Melissa Selik

Online:

< <http://cnx.org/content/col10232/1.1/> >

C O N N E X I O N S

Rice University, Houston, Texas

This selection and arrangement of content as a collection is copyrighted by Clayton Scott. It is licensed under the Creative Commons Attribution 1.0 license (<http://creativecommons.org/licenses/by/1.0>).

Collection structure revised: June 14, 2004

PDF generated: February 4, 2011

For copyright and attribution information for the modules contained in this collection, see p. 131.

Table of Contents

Introduction to Statistical Signal Processing	1
1 Preliminaries	
1.1 Review of Linear Algebra	7
1.2 The Bayesian Paradigm	17
1.3 Sufficiency	27
2 Signal Representation and Modeling	
2.1 Signal Classifications and Properties	33
2.2 Signal Models	40
3 Detection Theory	
3.1 Hypothesis Testing	47
3.2 Criteria	52
3.3 Detection with Unknowns	76
4 Estimation Theory	
4.1 Minimum Variance Unbiased Estimation	77
4.2 Maximum Likelihood Estimation	90
4.3 Bayesian Estimation	96
5 Adaptive Filtering	
5.1 Introduction to Adaptive Filtering	113
5.2 LMS Algorithm Analysis	120
Glossary	125
Bibliography	126
Index	128
Attributions	131

Introduction to Statistical Signal Processing¹

Digital Signal Processing

- Digital \equiv sampled, discrete-time, quantized
- Signal \equiv waveform, sequence of measurements or observations
- Processing \equiv analyze, modify, filter, synthesize

Examples of Digital Signals

- sampled speech waveform
- "pixelized" image
- Dow-Jones Index

DSP Applications

- Filtering (noise reduction)
- Pattern recognition (speech, faces, fingerprints)
- Compression

A Major Difficulty

In many (perhaps most) DSP applications we don't have complete or perfect knowledge of the signals we wish to process. We are faced with many **unknowns** and **uncertainties**.

Examples

- noisy measurements
- unknown signal parameters
- noisy system or environmental conditions
- natural variability in the signals encountered

¹This content is available online at <<http://cnx.org/content/m11947/1.3/>>.

Functional Magnetic Resonance Imaging
Image not finished

Figure 1: Challenges are measurement noise and intrinsic uncertainties in signal behavior.

How can we design signal processing algorithms in the face of such uncertainty?
 Can we model the uncertainty and incorporate this model into the design process?
Statistical signal processing is the study of these questions.

Modeling Uncertainty

The most widely accepted and commonly used approach to modeling uncertainty is **Probability Theory** (although other alternatives exist such as Fuzzy Logic).

Probability Theory models uncertainty by specifying the **chance** of observing certain signals.

Alternatively, one can view probability as specifying the degree to which we **believe** a signal reflects the true **state of nature**.

Examples of Probabilistic Models

- errors in a measurement (due to an imprecise measuring device) modeled as realizations of a Gaussian random variable.
- uncertainty in the phase of a sinusoidal signal modeled as a uniform random variable on $[0, 2\pi)$.
- uncertainty in the number of photons striking a CCD per unit time modeled as a Poisson random variable.

Statistical Inference

A **statistic** is a function of observed data.

Example

Suppose we observe N scalar values x_1, \dots, x_N . The following are statistics:

- $\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n$ (sample mean)
- x_1, \dots, x_N (the data itself)
- $\min\{x_1, \dots, x_N\}$ (order statistic)
- $(x_1^2 - x_2 \sin(x_3), e^{-(x_1 x_3)})$

A statistic **cannot** depend on **unknown parameters**.

Probability is used to model uncertainty.

Statistics are used to draw conclusions about probability models.

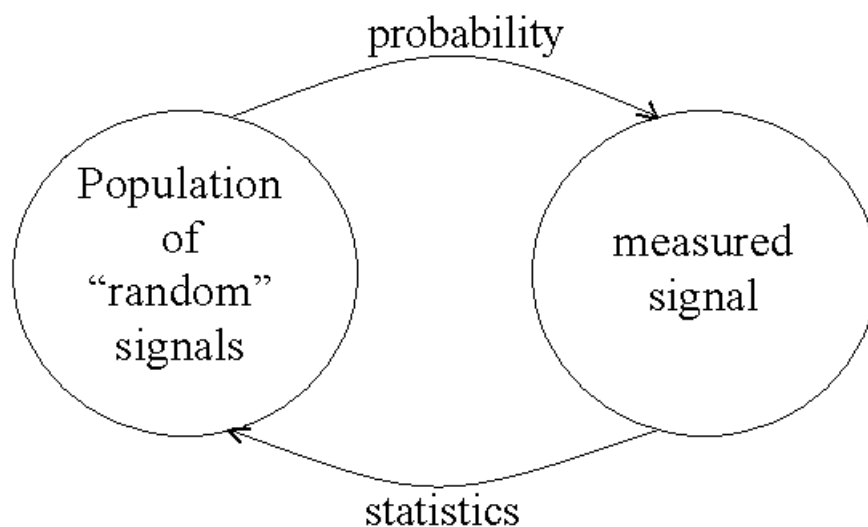


Figure 2

Probability models our uncertainty about signals we **may** observe.
Statistics reasons from the measured signal to the population of possible signals.

Statistical Signal Processing

- **Step 1** - Postulate a probability model (or models) that reasonably capture the uncertainties at hand
- **Step 2** - Collect data
- **Step 3** - Formulate statistics that allow us to interpret or understand our probability model(s)

In this class

The two major kinds of problems that we will study are **detection** and **estimation**. Most SSP problems fall under one of these two headings.

Detection Theory

Given two (or more) probability models, which one best explains the signal?

Examples

1. Decode wireless comm signal into string of 0's and 1's
2. Pattern recognition
 - voice recognition
 - face recognition
 - handwritten character recognition

3. Anomaly detection

- radar, sonar
- irregular, heartbeat
- gamma-ray burst in deep space

Estimation Theory

If our probability model has free parameters, what are the best parameter settings to describe the signal we've observed?

Examples

1. Noise reduction
2. Determine parameters of a sinusoid (phase, amplitude, frequency)
3. Adaptive filtering
 - track trajectories of space-craft
 - automatic control systems
 - channel equalization
4. Determine location of a submarine (sonar)
5. Seismology: estimate depth below ground of an oil deposit

Example : Detection Example

Suppose we observe N tosses of an unfair coin. We would like to decide which side the coin favors, heads or tails.

- **Step 1** - Assume each toss is a realization of a Bernoulli random variable.

$$Pr[\text{Heads}] = p = 1 - Pr[\text{Tails}]$$

Must decide $p = \frac{1}{4}$ vs. $p = \frac{3}{4}$.

- **Step 2** - Collect data x_1, \dots, x_N

$$x_i = 1 \equiv \text{Heads}$$

$$x_i = 0 \equiv \text{Tails}$$

- **Step 3** - Formulate a useful statistic

$$k = \sum_{n=1}^N x_n$$

If $k < \frac{N}{2}$, guess $p = \frac{1}{4}$. If $k > \frac{N}{2}$, guess $p = \frac{3}{4}$.

Example : Estimation Example

Suppose we take N measurements of a DC voltage A with a noisy voltmeter. We would like to estimate A .

- **Step 1** - Assume a Gaussian noise model

$$x_n = A + w_n$$

where $w_n \sim \mathcal{N}(0, 1)$.

- **Step 2** - Gather data x_1, \dots, x_N

- **Step 3** - Compute the sample mean,

$$A = \frac{1}{N} \sum_{n=1}^N x_n$$

and use this as an estimate.

In these examples (Example (Detection Example) and Example (Estimation Example)), we solved detection and estimation problems using intuition and heuristics (in Step 3).

This course will focus on developing principled and mathematically rigorous approaches to detection and estimation, using the theoretical framework of probability and statistics.

Summary

- DSP \equiv processing signals with computer algorithms.
- SSP \equiv statistical DSP \equiv processing in the presence of uncertainties and unknowns.

Chapter 1

Preliminaries

1.1 Review of Linear Algebra¹

Vector spaces are the principal object of study in linear algebra. A vector space is always defined with respect to a field of scalars.

1.1.1 Fields

A field is a set F equipped with two operations, addition and multiplication, and containing two special members 0 and 1 ($0 \neq 1$), such that for all $\{a, b, c\} \in F$

1.
 - a. $(a + b) \in F$
 - b. $a + b = b + a$
 - c. $(a + b) + c = a + (b + c)$
 - d. $a + 0 = a$
 - e. there exists $-a$ such that $a + -a = 0$
2.
 - a. $ab \in F$
 - b. $ab = ba$
 - c. $(ab)c = a(bc)$
 - d. $a \cdot 1 = a$
 - e. there exists a^{-1} such that $aa^{-1} = 1$
3. $a(b + c) = ab + ac$

More concisely

1. F is an abelian group under addition
2. F is an abelian group under multiplication
3. multiplication distributes over addition

1.1.1.1 Examples

$\mathbb{Q}, \mathbb{R}, \mathbb{C}$

1.1.2 Vector Spaces

Let F be a field, and V a set. We say V is a **vector space over F** if there exist two operations, defined for all $a \in F$, $u \in V$ and $v \in V$:

¹This content is available online at <http://cnx.org/content/m11948/1.2/>.

- vector addition: $(u, v) \rightarrow (u + v) \in V$
- scalar multiplication: $(a, v) \rightarrow av \in V$

and if there exists an element denoted $0 \in V$, such that the following hold for all $a \in F$, $b \in F$, and $u \in V$, $v \in V$, and $w \in V$

1.
 - a. $u + (v + w) = (u + v) + w$
 - b. $u + v = v + u$
 - c. $u + 0 = u$
 - d. there exists $-u$ such that $u + -u = 0$
2.
 - a. $a(u + v) = au + av$
 - b. $(a + b)u = au + bu$
 - c. $(ab)u = a(bu)$
 - d. $1 \cdot u = u$

More concisely,

1. V is an abelian group under plus
2. Natural properties of scalar multiplication

1.1.2.1 Examples

- \mathbb{R}^N is a vector space over R
- \mathbb{C}^N is a vector space over \mathbb{C}
- \mathbb{C}^N is a vector space over R
- \mathbb{R}^N is **not** a vector space over \mathbb{C}

The elements of V are called **vectors**.

1.1.3 Euclidean Space

Throughout this course we will think of a signal as a vector

$$x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} = \left(x_1 \quad x_2 \quad \dots \quad x_N \right)^T$$

The samples $\{x_i\}$ could be samples from a finite duration, continuous time signal, for example.

A signal will belong to one of two vector spaces:

1.1.3.1 Real Euclidean space

$$x \in \mathbb{R}^N \text{ (over } R\text{)}$$

1.1.3.2 Complex Euclidean space

$$x \in \mathbb{C}^N \text{ (over } \mathbb{C}\text{)}$$

1.1.4 Subspaces

Let V be a vector space over F .

A subset $S \subseteq V$ is called a **subspace** of V if S is a vector space over F in its own right.

Example 1.1

$V = \mathbb{R}^2$, $F = \mathbb{R}$, $S =$ any line through the origin.




Figure 1.1: S is any line through the origin.

Are there other subspaces?

Theorem 1.1:

$S \subseteq V$ is a subspace if and only if for all $a \in F$ and $b \in F$ and for all $s \in S$ and $t \in S$, $(as + bt) \in S$

1.1.5 Linear Independence

Let $u_1, \dots, u_k \in V$.

We say that these vectors are **linearly dependent** if there exist scalars $a_1, \dots, a_k \in F$ such that

$$\sum_{i=1}^k a_i u_i = 0 \quad (1.1)$$

and at least one $a_i \neq 0$.

If (1.1) only holds for the case $a_1 = \dots = a_k = 0$, we say that the vectors are **linearly independent**.

Example 1.2

$$1 \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} - 2 \begin{pmatrix} -2 \\ 3 \\ 0 \end{pmatrix} + 1 \begin{pmatrix} -5 \\ 7 \\ -2 \end{pmatrix} = 0$$

so these vectors are linearly dependent in \mathbb{R}^3 .

1.1.6 Spanning Sets

Consider the subset $S = \{v_1, v_2, \dots, v_k\}$. Define the **span** of S

$$\langle S \rangle \equiv \text{span}(S) \equiv \left\{ \sum_{i=1}^k a_i v_i \mid a_i \in F \right\}$$

Fact: $\langle S \rangle$ is a subspace of V .

Example 1.3

$$V = \mathbb{R}^3, F = \mathbb{R}, S = \{v_1, v_2\}, v_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, v_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \Rightarrow \langle S \rangle = \text{xy-plane.}$$




Figure 1.2: $\langle S \rangle$ is the xy-plane.

1.1.6.1 Aside

If S is infinite, the notions of linear independence and span are easily generalized:

We say S is linearly independent if, for every finite collection $u_1, \dots, u_k \in S$, (k arbitrary) we have

$$\left(\sum_{i=1}^k a_i u_i = 0 \right) \Rightarrow a_i = 0$$

The span of S is

$$\langle S \rangle = \left\{ \sum_{i=1}^k a_i u_i \mid a_i \in F \wedge u_i \in S \wedge (k < \infty) \right\}$$

NOTE: In both definitions, we only consider **finite** sums.

1.1.7 Bases

A set $B \subseteq V$ is called a **basis** for V over F if and only if

1. B is linearly independent
2. $\langle B \rangle = V$

Bases are of fundamental importance in signal processing. They allow us to decompose a signal into building blocks (basis vectors) that are often more easily understood.

Example 1.4

$V =$ (real or complex) Euclidean space, \mathbb{R}^N or \mathbb{C}^N .

$$B = \{e_1, \dots, e_N\} \equiv \text{standard basis}$$

$$e_i = \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}$$

where the 1 is in the i^{th} position.

Example 1.5

$V = \mathbb{C}^N$ over \mathbb{C} .

$$B = \{u_1, \dots, u_N\}$$

which is the DFT basis.

$$u_k = \begin{pmatrix} 1 \\ e^{-i2\pi \frac{k}{N}} \\ \vdots \\ e^{-i2\pi \frac{k}{N}(N-1)} \end{pmatrix}$$

where $i = \sqrt{-1}$.

1.1.7.1 Key Fact

If B is a basis for V , then every $v \in V$ can be written uniquely (up to order of terms) in the form

$$v = \sum_{i=1}^N a_i v_i$$

where $a_i \in F$ and $v_i \in B$.

1.1.7.2 Other Facts

- If S is a linearly independent set, then S can be extended to a basis.
- If $\langle S \rangle = V$, then S contains a basis.

1.1.8 Dimension

Let V be a vector space with basis B . The dimension of V , denoted $\dim(V)$, is the cardinality of B .

Theorem 1.2:

Every vector space has a basis.

Theorem 1.3:

Every basis for a vector space has the same cardinality.

$\Rightarrow \dim(V)$ is **well-defined**.

If $\dim(V) < \infty$, we say V is **finite dimensional**.

1.1.8.1 Examples

vector space	field of scalars	dimension
\mathbb{R}^N	\mathbb{R}	
\mathbb{C}^N	\mathbb{C}	
\mathbb{C}^N	\mathbb{R}	

Table 1.1

Every subspace is a vector space, and therefore has its own dimension.

Example 1.6

Suppose $(S = \{u_1, \dots, u_k\}) \subseteq V$ is a linearly independent set. Then

$$\dim(\langle S \rangle) =$$

Facts

- If S is a subspace of V , then $\dim(S) \leq \dim(V)$.
- If $\dim(S) = \dim(V) < \infty$, then $S = V$.

1.1.9 Direct Sums

Let V be a vector space, and let $S \subseteq V$ and $T \subseteq V$ be subspaces.

We say V is the **direct sum** of S and T , written $V = S \oplus T$, if and only if for every $v \in V$, there exist unique $s \in S$ and $t \in T$ such that $v = s + t$.

If $V = S \oplus T$, then T is called a **complement** of S .

Example 1.7

$$V = C' = \{f : \mathbb{R} \rightarrow \mathbb{R} \mid f \text{ is continuous}\}$$

$$S = \text{even functions in } C'$$

$$T = \text{odd functions in } C'$$

$$f(t) = \frac{1}{2}(f(t) + f(-t)) + \frac{1}{2}(f(t) - f(-t))$$

If $f = g + h = g' + h'$, $g \in S$ and $g' \in S$, $h \in T$ and $h' \in T$, then $g - g' = h' - h$ is odd and even, which implies $g = g'$ and $h = h'$.

1.1.9.1 Facts

1. Every subspace has a complement
2. $V = S \oplus T$ if and only if
 - a. $S \cap T = \{0\}$
 - b. $\langle S, T \rangle = V$
3. If $V = S \oplus T$, and $\dim(V) < \infty$, then $\dim(V) = \dim(S) + \dim(T)$

1.1.9.2 Proofs

Invoke a basis.

1.1.10 Norms

Let V be a vector space over F . A norm is a mapping $V \rightarrow F$, denoted by $\| \cdot \|$, such that for all $u \in V$, $v \in V$, and $\lambda \in F$

1. $\| u \| > 0$ if $u \neq 0$
2. $\| \lambda u \| = |\lambda| \| u \|$
3. $\| u + v \| \leq \| u \| + \| v \|$

1.1.10.1 Examples

Euclidean norms:

$$x \in \mathbb{R}^N:$$

$$\| x \| = \left(\sum_{i=1}^N x_i^2 \right)^{\frac{1}{2}}$$

$$x \in \mathbb{C}^N:$$

$$\| x \| = \left(\sum_{i=1}^N (|x_i|)^2 \right)^{\frac{1}{2}}$$

1.1.10.2 Induced Metric

Every norm induces a metric on V

$$d(u, v) \equiv \| u - v \|$$

which leads to a notion of "distance" between vectors.

1.1.11 Inner products

Let V be a vector space over F , $F = \mathbb{R}$ or \mathbb{C} . An inner product is a mapping $V \times V \rightarrow F$, denoted $\langle \cdot, \cdot \rangle$, such that

1. $\langle v, v \rangle \geq 0$, and $\langle v, v \rangle = 0 \Leftrightarrow v = 0$
2. $\langle u, v \rangle = \overline{\langle v, u \rangle}$
3. $\langle au + bv, w \rangle = a \langle u, w \rangle + b \langle v, w \rangle$

1.1.11.1 Examples

\mathbb{R}^N over \mathbb{R} :

$$\langle x, y \rangle = (x^T y) = \sum_{i=1}^N x_i y_i$$

\mathbb{C}^N over \mathbb{C} :

$$\langle x, y \rangle = (x^H y) = \sum_{i=1}^N \bar{x}_i y_i$$

If $(x = (x_1, \dots, x_N)^T) \in \mathbb{C}$, then

$$x^H \equiv \begin{pmatrix} \overline{x_1} \\ \vdots \\ \overline{x_N} \end{pmatrix}^T$$

is called the "Hermitian," or "conjugate transpose" of x .

1.1.12 Triangle Inequality

If we define $\|u\| = \langle u, u \rangle$, then

$$\|u + v\| \leq \|u\| + \|v\|$$

Hence, every inner product induces a norm.

1.1.13 Cauchy-Schwarz Inequality

For all $u \in V, v \in V$,

$$|\langle u, v \rangle| \leq \|u\| \|v\|$$

In inner product spaces, we have a notion of the angle between two vectors:

$$\left(\angle(u, v) = \arccos \left(\frac{\langle u, v \rangle}{\|u\| \|v\|} \right) \right) \in [0, 2\pi)$$

1.1.14 Orthogonality

u and v are **orthogonal** if

$$\langle u, v \rangle = 0$$

Notation: $u \perp v$.

If in addition $\|u\| = \|v\| = 1$, we say u and v are **orthonormal**.

In an orthogonal (orthonormal) **set**, each pair of vectors is orthogonal (orthonormal).




Figure 1.3: Orthogonal vectors in \mathbb{R}^2 .

1.1.15 Orthonormal Bases

An Orthonormal basis is a basis $\{v_i\}$ such that

$$\langle v_i, v_j \rangle = \delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Example 1.8

The standard basis for \mathbb{R}^N or \mathbb{C}^N

Example 1.9

The normalized DFT basis

$$u_k = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 \\ e^{-i2\pi \frac{k}{N}} \\ \vdots \\ e^{-i2\pi \frac{k}{N}(N-1)} \end{pmatrix}$$

1.1.16 Expansion Coefficients

If the representation of v with respect to $\{v_i\}$ is

$$v = \sum_i a_i v_i$$

then

$$a_i = \langle v_i, v \rangle$$

1.1.17 Gram-Schmidt

Every inner product space has an orthonormal basis. Any (countable) basis can be made orthogonal by the Gram-Schmidt orthogonalization process.

1.1.18 Orthogonal Compliments

Let $S \subseteq V$ be a subspace. The **orthogonal compliment** S^\perp is

$$S^\perp = \{ u \mid u \in V \wedge \langle u, v \rangle = 0 \wedge v \in S \}$$

S^\perp is easily seen to be a subspace.

If $\dim(v) < \infty$, then $V = S \oplus S^\perp$.

ASIDE: If $\dim(v) = \infty$, then in order to have $V = S \oplus S^\perp$ we require V to be a **Hilbert Space**.

1.1.19 Linear Transformations

Loosely speaking, a linear transformation is a mapping from one vector space to another that **preserves** vector space operations.

More precisely, let V, W be vector spaces over the same field F . A **linear transformation** is a mapping $T: V \rightarrow W$ such that

$$T(au + bv) = aT(u) + bT(v)$$

for all $a \in F, b \in F$ and $u \in V, v \in V$.

In this class we will be concerned with linear transformations between (real or complex) **Euclidean spaces**, or subspaces thereof.

1.1.20 Image

$$\text{image}(T) = \{ w \mid w \in W \wedge T(v) = w \text{ for some } v \}$$

1.1.21 Nullspace

Also known as the kernel:

$$\ker(T) = \{v \mid v \in V \wedge (T(v) = 0)\}$$

Both the image and the nullspace are easily seen to be subspaces.

1.1.22 Rank

$$\text{rank}(T) = \dim(\text{image}(T))$$

1.1.23 Nullity

$$\text{null}(T) = \dim(\ker(T))$$

1.1.24 Rank plus nullity theorem

$$\text{rank}(T) + \text{null}(T) = \dim(V)$$

1.1.25 Matrices

Every linear transformation T has a **matrix representation**. If $T : E^N \rightarrow E^M$, $E = \mathbb{R}$ or \mathbb{C} , then T is represented by an $M \times N$ matrix

$$A = \begin{pmatrix} a_{11} & \dots & a_{1N} \\ \vdots & \ddots & \vdots \\ a_{M1} & \dots & a_{MN} \end{pmatrix}$$

where $(a_{1i}, \dots, a_{Mi})^T = T(e_i)$ and $e_i = (0, \dots, 1, \dots, 0)^T$ is the i^{th} **standard basis** vector.

ASIDE: A linear transformation can be represented with respect to any bases of E^N and E^M , leading to a different A . We will always represent a linear transformation using the standard bases.

1.1.26 Column span

$$\text{colspan}(A) = \langle A \rangle = \text{image}(A)$$

1.1.27 Duality

If $A : \mathbb{R}^N \rightarrow \mathbb{R}^M$, then

$$\ker^\perp(A) = \text{image}(A^T)$$




Figure 1.4

If $A : \mathbb{C}^N \rightarrow \mathbb{C}^M$, then

$$\ker^\perp(A) = \text{image}(A^H)$$

1.1.28 Inverses

The linear transformation/matrix A is **invertible** if and only if there exists a matrix B such that $AB = BA = I$ (identity).

Only **square** matrices can be invertible.

Theorem 1.4:

Let $A : F^N \rightarrow F^N$ be linear, $F = \mathbb{R}$ or \mathbb{C} . The following are equivalent:

1. A is invertible (nonsingular)
2. $\text{rank}(A) = N$
3. $\text{null}(A) = 0$
4. $\det A \neq 0$
5. The columns of A form a basis.

If $A^{-1} = A^T$ (or A^H in the complex case), we say A is **orthogonal** (or **unitary**).

1.2 The Bayesian Paradigm²

Statistical analysis is fundamentally an **inversion** process. The objective is to the "causes"—parameters of the probabilistic data generation model—from the "effects"—observations. This can be seen in our interpretation of the likelihood function.

Given a parameter θ , observations are generated according to

$$p(x | \theta)$$

The likelihood function has the same form as the conditional density function above

$$l(\theta|x) \equiv p(x | \theta)$$

²This content is available online at <<http://cnx.org/content/m11444/1.7/>>.

except now x is given (we take measurements) and θ is the variable. The likelihood function essentially inverts the role of observation (effect) and parameter (cause).

Unfortunately, the likelihood function does not provide a formal framework for the desired inversion.

One problem is that the parameter θ is supposed to be a fixed and deterministic quantity while the observation x is the realization of a random process. So their role aren't really interchangeable in this setting.

Moreover, while it is tempting to interpret the likelihood $l(\theta|x)$ as a density function for θ , this is not always possible; for example, often

$$\int l(\theta|x) d\theta \rightarrow \infty$$

Another problematic issue is the mathematical formalization of statements like: "Based on the measurements x , I am 95% confident that θ falls in a certain range."

Example 1.10

Suppose you toss a coin 10 times and each time it comes up "heads." It might be reasonable to say that we are 99% sure that the coin is unfair, biased towards heads.

Formally:

$$H_0 : \theta \equiv \text{prob heads} > 0.5$$

$$x \sim \binom{N}{\sum x} \theta^{\sum x} (1 - \theta)^{N - \sum x}$$

which is the binomial likelihood.

$$p(\theta > 0.5 | x) = ?$$

The problem with this is that

$$p(\theta \in H_0 | x)$$

implies that θ is a **random**, not deterministic, quantity. So, while "confidence" statements are very reasonable and in fact a normal part of "everyday thinking," this idea can not be supported from the classical perspective.

All of these "deficiencies" can be circumvented by a change in how we view the parameter θ .

If we view θ as the realization of a random variable with density $p(\theta)$, then **Bayes Rule** (Bayes, 1763) shows that

$$p(\theta | x) = \frac{p(x | \theta) p(\theta)}{\int p(x | \theta) p(\theta) d\theta}$$

Thus, from this perspective we obtain a well-defined inversion: Given x , the parameter θ is distributed according to $p(\theta | x)$.

From here, confidence measures such as $p(\theta \in H_0 | x)$ are perfectly legitimate quantities to ask for.

Definition 1.1: Bayesian statistical model

A statistical model compose of a data generation model, $p(x | \theta)$, and a **prior** distribution on the parameters, $p(\theta)$.

The **prior distribution** (or **prior** for short) models the uncertainty in the parameter. More specifically, $p(\theta)$ models our knowledge—or lack thereof—prior to collecting data.

Notice that

$$p(\theta | x) = \frac{p(x | \theta) p(\theta)}{p(x)} \propto p(x | \theta) p(\theta)$$

since the data x are **known**, $p(x)$ is just a constant. Hence, $p(\theta | x)$ is proportional to the likelihood function multiplied by the prior.

Bayesian analysis has some significant advantages over classical statistical analysis:

1. properly inverts the relationship between causes and effects
2. permits meaningful assessments in confidence regions
3. enables the incorporation of prior knowledge into the analysis (which could come from previous experiments, for example)
4. leads to more accurate estimators (provided the prior knowledge is accurate)
5. obeys the Likelihood and Sufficiency principles

Example 1.11

$$x_n = A + W_n \quad , \quad n = \{1, \dots, N\}$$

$$W_n \sim \mathcal{N}(0, \sigma^2)$$

iid.

$$A = \frac{1}{N} \sum_{n=1}^N x_n$$

MVUB and MLE estimator. Now suppose that we have prior knowledge that $-A_0 \leq A \leq A_0$. We might incorporate this by forming a new estimator

$$\tilde{A} = \begin{cases} -A_0 & \text{if } A < -A_0 \\ A & \text{if } -A_0 \leq A \leq A_0 \\ A_0 & \text{if } A > A_0 \end{cases} \quad (1.2)$$

This is called a **truncated** sample mean estimator of A . Is \tilde{A} a better estimator of A than the sample mean A ?

Let $p(a)$ denote the density of A . Since $A = \frac{1}{N} \sum x_n$, $p(a) = \mathcal{N}\left(A, \frac{\sigma^2}{N}\right)$. The density of \tilde{A} is given by

$$\tilde{p}(a) = Pr[A \leq -A_0] \delta(a + A_0) + p(a) I_{\{-A_0 \leq a \leq A_0\}} + Pr[A \geq A_0] \delta(a - A_0) \quad (1.3)$$




Figure 1.5

Now consider the MSE of the sample mean A .

$$\text{MSE}(A) = \int_{-\infty}^{\infty} (a - A)^2 p(a) da \quad (1.4)$$

1.2.1 Note

1. \tilde{A} is biased (Figure 1.6).
2. Although A is MVUB, \tilde{A} is better in the MSE sense.
3. Prior information is aptly described by regarding A as a random variable with a prior distribution $U(-A_0, A_0)$, which implies that we know $-A_0 \leq A \leq A_0$, but otherwise A is arbitrary.




Figure 1.6: (a) Mean of $A = A$. (b) Mean of $\tilde{A} \neq A$.

1.2.2 The Bayesian Approach to Statistical Modeling

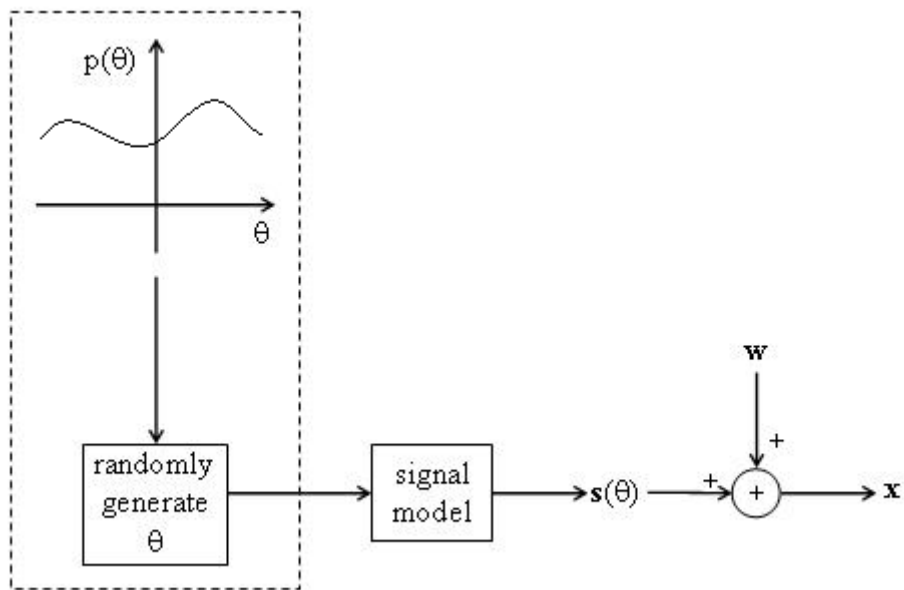


Figure 1.7: Where w is the noise and x is the observation.

Example 1.12

$$x_n = A + W_n, \quad n = \{1, \dots, N\}$$

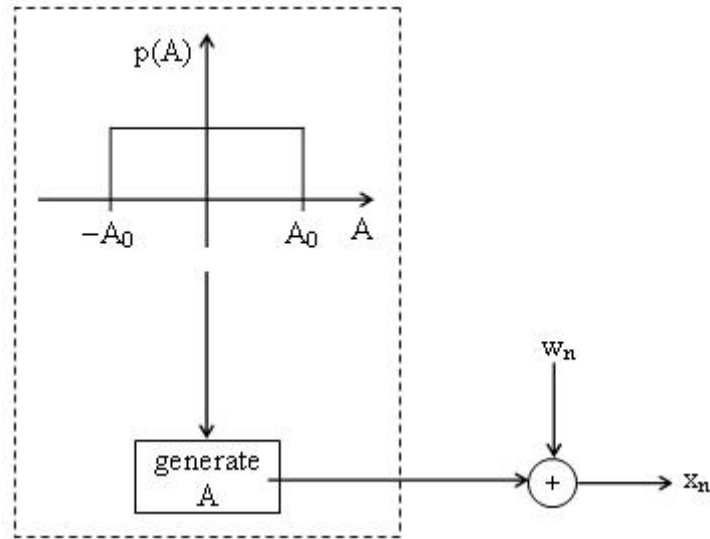


Figure 1.8

Prior distribution allows us to incorporate prior information regarding unknown parameter—probable values of parameter are supported by prior. Basically, the prior reflects what we believe "Nature" will probably throw at us.

1.2.3 Elements of Bayesian Analysis

(a) - joint distribution

$$p(x, \theta) = p(x | \theta) p(\theta)$$

(b) - marginal distributions

$$p(x) = \int p(x | \theta) p(\theta) d\theta$$

$$p(\theta) = \int p(x | \theta) p(\theta) dx$$

where $p(\theta)$ is a **prior**.

(c) - posterior distribution

$$p(\theta | x) = \frac{p(x, \theta)}{p(x)} = \frac{p(x | \theta) p(\theta)}{\int p(x | \theta) p(\theta) dx}$$

Example 1.13

$$p(x | \theta) = \binom{n}{x} \theta^x (1 - \theta)^{n-x}, \quad \theta \in [0, 1]$$

which is the Binomial likelihood.

$$p(\theta) = \frac{1}{B(\alpha, \beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1}$$

which is the Beta prior distribution and $B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$




Figure 1.9: This reflects prior knowledge that most probable values of θ are close to $\frac{\alpha}{\alpha+\beta}$.

1.2.3.1 Joint Density

$$p(x, \theta) = \frac{\binom{n}{x}}{B(\alpha, \beta)} \theta^{\alpha+x-1} (1-\theta)^{n-x+\beta-1}$$

1.2.3.2 marginal density

$$p(x) = \binom{n}{x} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(\alpha+x)\Gamma(n-x+\beta)}{\Gamma(\alpha+\beta+n)}$$

1.2.3.3 posterior density

$$p(\theta | x) = \frac{\theta^{\alpha+x-1} \theta^{\beta+n-x-1}}{B(\alpha+x, \beta+n-x)}$$

where $B(\alpha+x, \beta+n-x)$ is the Beta density with parameters $\alpha' = \alpha+x$ and $\beta' = \beta+n-x$

1.2.4 Selecting an Informative Prior

Clearly, the most important objective is to choose the prior $p(\theta)$ that best reflects the prior knowledge available to us. In general, however, our prior knowledge is imprecise and any number of prior densities may aptly capture this information. Moreover, usually the optimal estimator can't be obtained in closed-form.

Therefore, sometimes it is desirable to choose a prior density that models prior knowledge **and** is nicely matched in functional form to $p(x|\theta)$ so that the optimal estimator (and posterior density) can be expressed in a simple fashion.

1.2.5 Choosing a Prior

1. Informative Priors

- design/choose priors that are compatible with prior knowledge of unknown parameters

2. Non-informative Priors

- attempt to remove subjectiveness from Bayesian procedures
- designs are often based on invariance arguments

Example 1.14

Suppose we want to estimate the variance of a process, incorporating a prior that is amplitude-scale invariant (so that we are invariant to arbitrary amplitude rescaling of data). $p(s) = \frac{1}{s}$ satisfies this condition. $\sigma^2 \sim p(s) \Rightarrow A\sigma^2 \sim p(s)$ where $p(s)$ is non-informative since it is invariant to amplitude-scale.

1.2.6 Conjugate Priors

1.2.6.1 Idea

Given $p(x|\theta)$, choose $p(\theta)$ so that $p(\theta|x) \propto p(x|\theta)p(\theta)$ has a simple functional form.

1.2.6.2 Conjugate Priors

Choose $p(\theta) \in \mathcal{P}$, where \mathcal{P} is a family of densities (e.g., Gaussian family) so that the posterior density also belongs to that family.

Definition 1.2: conjugate prior

$p(\theta)$ is a **conjugate prior** for $p(x|\theta)$ if $p(\theta) \in \mathcal{P} \Rightarrow p(\theta|x) \in \mathcal{P}$

Example 1.15

$$x_n = A + W_n \quad , \quad n = \{1, \dots, N\}$$

$$W_n \sim \mathcal{N}(0, \sigma^2)$$

iid. Rather than modeling $A \sim U(-A_0, A_0)$ (which did not yield a closed-form estimator) consider

$$p(A) = \frac{1}{\sqrt{2\pi\sigma_A^2}} e^{-\frac{1}{2\sigma_A^2}(A-\mu)^2}$$




Figure 1.10

With $\mu = 0$ and $\sigma_A = \frac{1}{3}A_0$ this Gaussian prior also reflects prior knowledge that it is unlikely for $|A| \geq A_0$.

The Gaussian prior is also conjugate to the Gaussian likelihood

$$p(x | A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2}$$

so that the resulting posterior density is also a simple Gaussian, as shown next.

First note that

$$p(x | A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N x_n} e^{-\frac{1}{2\sigma^2} (NA^2 - 2NA\bar{x})}$$

where $\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n$.

$$\begin{aligned} p(A | x) &= \frac{p(x | A)p(A)}{\int p(x | A)p(A)dA} \\ &= \frac{e^{-\frac{1}{2} \left(\frac{1}{\sigma^2} (NA^2 - 2NA\bar{x}) + \frac{1}{\sigma_A^2} (A - \mu)^2 \right)}}{\int_{-\infty}^{\infty} e^{-\frac{1}{2} \left(\frac{1}{\sigma^2} (NA^2 - 2NA\bar{x}) + \frac{1}{\sigma_A^2} (A - \mu)^2 \right)} dA} \\ &= \frac{e^{-\frac{1}{2} Q(A)}}{\int_{-\infty}^{\infty} e^{-\frac{1}{2} Q(A)} dA} \end{aligned} \quad (1.5)$$

where $Q(A) = \frac{N}{\sigma^2} A^2 - \frac{2NA\bar{x}}{\sigma^2} + \frac{A^2}{\sigma_A^2} - \frac{2\mu A}{\sigma_A^2} + \frac{\mu^2}{\sigma_A^2}$. Now let

$$\sigma_{A|x}^2 \equiv \frac{1}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}}$$

$$\mu_{A|x} \equiv \left(\frac{N}{\sigma^2} \bar{x} + \frac{\mu}{\sigma_A^2} \right) \sigma_{A|x}^2$$

Then by "completing the square" we have

$$\begin{aligned} Q(A) &= \frac{1}{\sigma_{A|x}^2} (A^2 - 2\mu_{A|x}A + \mu_{A|x}^2) - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} + \frac{\mu^2}{\sigma_A^2} \\ &= \frac{1}{\sigma_{A|x}^2} (A - \mu_{A|x})^2 - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2} + \frac{\mu^2}{\sigma_A^2} \end{aligned} \quad (1.6)$$

Hence,

$$p(A|x) = \frac{e^{\frac{-1}{2\sigma_{A|x}^2}(A-\mu_{A|x})^2} e^{\frac{-1}{2}\left(\frac{\mu^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2}\right)}}{\int_{-\infty}^{\infty} e^{\frac{-1}{2\sigma_{A|x}^2}(A-\mu_{A|x})^2} e^{\frac{-1}{2}\left(\frac{\mu^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2}\right)} dA}$$

where $\frac{-1}{2\sigma_{A|x}^2}(A-\mu_{A|x})^2$ is the "unnormalized" Gaussian density and $\frac{-1}{2}\left(\frac{\mu^2}{\sigma_A^2} - \frac{\mu_{A|x}^2}{\sigma_{A|x}^2}\right)$ is a constant, independent of A . This implies that

$$p(A|x) = \frac{1}{\sqrt{2\pi\sigma_{A|x}^2}} e^{\frac{-1}{2\sigma_{A|x}^2}(A-\mu_{A|x})^2}$$

where $A|x \sim \mathcal{N}(\mu_{A|x}, \sigma_{A|x}^2)$. Now

$$\begin{aligned} A &= E[A|x] \\ &= \int Ap(A|x) dA \\ &= \mu_{A|x} \\ &= \frac{\frac{N}{\sigma^2}\bar{x} + \frac{\mu}{\sigma_A^2}}{\frac{N}{\sigma^2} + \frac{1}{\sigma_A^2}} \\ &= \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} \bar{x} + \frac{\frac{\sigma^2}{N}}{\sigma_A^2 + \frac{\sigma^2}{N}} \mu \\ &= \alpha \bar{x} - 1\mu \end{aligned} \tag{1.7}$$

Where $0 < \alpha = \frac{\sigma_A^2}{\sigma_A^2 + \frac{\sigma^2}{N}} < 1$

Interpretation

1. When there is little data $\sigma_A^2 \ll \frac{\sigma^2}{N}$ α is small and $A = \mu$.
2. When there is a lot of data $\sigma_A^2 \gg \frac{\sigma^2}{N}$, $\alpha \simeq 1$ and $A = \bar{x}$.

1.2.7 Interplay Between Data and Prior Knowledge

Small $N \rightarrow A$ favors prior.

Large $N \rightarrow A$ favors data.

1.2.8 The Multivariate Gaussian Model

The multivariate Gaussian model is the most important Bayesian tool in signal processing. It leads directly to the celebrated Wiener and Kalman filters.

Assume that we are dealing with random vectors x and y . We will regard y as a signal vector that is to be estimated from an observation vector x .

y plays the same role as θ did in earlier discussions. We will assume that y is $p \times 1$ and x is $N \times 1$. Furthermore, assume that x and y are **jointly** Gaussian distributed

$$\begin{pmatrix} x \\ y \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix} \right)$$

$$E[x] = 0, E[y] = 0, E[xx^T] = R_{xx}, E[xy^T] = R_{xy}, E[yx^T] = R_{yx}, E[yy^T] = R_{yy}.$$

$$R \equiv \begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix}$$

Example 1.16

$x = y + W$, $W \sim \mathcal{N}(0, \sigma^2 I)$

$$p(y) = \mathcal{N}(0, R_{yy})$$

which is independent of W . $E[x] = E[y] + E[W] = 0$, $E[xx^T] = E[yy^T] + E[yW^T] + E[Wy^T] + E[WW^T] = R_{yy} + \sigma^2 I$, $E[xy^T] = E[yy^T] + E[Wy^T] = R_{yy} = E[yx^T]$.

$$\begin{pmatrix} x \\ y \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} R_{yy} + \sigma^2 I & R_{yy} \\ R_{yy} & R_{yy} \end{pmatrix} \right)$$

From our Bayesian perspective, we are interested in $p(y|x)$.

$$\begin{aligned} p(y|x) &= \frac{p(x,y)}{p(x)} \\ &= \frac{(2\pi)^{-\frac{N}{2}} (2\pi)^{-\frac{p}{2}} (\det R)^{-\frac{1}{2}} e^{-\frac{1}{2} \begin{pmatrix} x^T & y^T \end{pmatrix} R^{-1} \begin{pmatrix} x \\ y \end{pmatrix}}}{(2\pi)^{-\frac{N}{2}} (\det R_{xx})^{-\frac{1}{2}} e^{-\frac{1}{2} x^T R_{xx}^{-1} x}} \end{aligned} \quad (1.8)$$

In this formula we are faced with

$$R^{-1} = \begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix}^{-1}$$

The inverse of this covariance matrix can be written as

$$\begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix}^{-1} = \begin{pmatrix} R_{xx}^{-1} & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} (-R_{xx}^{-1}) R_{xy} \\ I \end{pmatrix} Q^{-1} \begin{pmatrix} (-R_{yx}) R_{xx} & I \end{pmatrix}$$

where $Q \equiv R_{yy} - R_{yx} R_{xx}^{-1} R_{xy}$. (Verify this formula by applying the right hand side above to R to get I .)

1.3 Sufficiency

1.3.1 Sufficient Statistics³

1.3.1.1 Introduction

Sufficient statistics arise in nearly every aspect of statistical inference. It is important to understand them before progressing to areas such as hypothesis testing and parameter estimation.

Suppose we observe an N -dimensional random vector X , characterized by the density or mass function $f_\theta(x)$, where θ is a p -dimensional vector of parameters to be estimated. The functional form of $f(x)$ is assumed known. The parameter θ completely determines the distribution of X . Conversely, a measurement x of X provides information about θ through the probability law $f_\theta(x)$.

³This content is available online at <http://cnx.org/content/m11481/1.7/>.

Example 1.17

Suppose $X = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$, where $X_i \sim \mathcal{N}(\theta, 1)$ are IID. Here θ is a scalar parameter specifying the mean. The distribution of X is determined by θ through the density

$$f_{\theta}(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(x_1 - \theta)^2}{2}} \frac{1}{\sqrt{2\pi}} e^{-\frac{(x_2 - \theta)^2}{2}}$$

On the other hand, if we observe $x = \begin{pmatrix} 100 \\ 102 \end{pmatrix}$, then we may safely assume $\theta = 0$ is highly unlikely.

The N -dimensional observation X carries information about the p -dimensional parameter vector θ . If $p < N$, one may ask the following question: Can we compress x into a low-dimensional statistic without any loss of information? Does there exist some function $t = T(x)$, where the dimension of t is $M < N$, such that t carries all the useful information about θ ?

If so, for the purpose of studying θ we could discard the raw measurements x and retain only the low-dimensional statistic t . We call t a **sufficient statistic**. The following definition captures this notion precisely:

Definition 1.3:

Let X_1, \dots, X_M be a random sample, governed by the density or probability mass function $f(x|\theta)$.

The statistic $T(x)$ is **sufficient** for θ if the conditional distribution of x , given $T(x) = t$, is independent of θ . Equivalently, the functional form of $f_{\theta|t}(x)$ does not involve θ .

How should we interpret this definition? Here are some possibilities:

1.3.1.1.1

1. Let $f_{\theta}(x, t)$ denote the joint density or probability mass function on $(X, T(X))$. If $T(X)$ is a sufficient statistic for θ , then

$$\begin{aligned} f_{\theta}(x) &= f_{\theta}(x, T(x)) \\ &= f_{\theta|t}(x) f_{\theta}(t) \\ &= f(x|t) f_{\theta}(t) \end{aligned} \tag{1.9}$$

Therefore, the parametrization of the probability law for the measurement x is manifested in the parametrization of the probability law for the statistic $T(x)$.

1.3.1.1.2

2. Given $t = T(x)$, full knowledge of the measurement x brings no additional information about θ . Thus, we may discard x and retain on the compressed statistic t .

1.3.1.1.3

3. Any inference strategy based on $f_{\theta}(x)$ may be replaced by a strategy based on $f_{\theta}(t)$.

Example 1.18**1.3.1.1.1 Binary Information Source**

(Scharf, pp.78[6]) Suppose a binary information source emits a sequence of binary (0 or 1) valued, independent variables x_1, \dots, x_N . Each binary symbol may be viewed as a realization of a Bernoulli trial: $x_n \sim \text{Bernoulli}(\theta)$, iid. The parameter $\theta \in [0, 1]$ is to be estimated.

The probability mass function for the random sample $x = (x_1, \dots, x_N)^T$ is

$$f_{\theta}(x) = \prod_{n=1}^N f_{\theta}(x_n) \prod_{n=1}^N \theta^{x_n} (1-\theta)^{1-x_n} \quad (1.10)$$

where $k = \sum_{n=1}^N x_n$ is the number of 1's in the sample.

We will show that k is a sufficient statistic for x . This will entail showing that the conditional probability mass function $f_{\theta|k}(x)$ does not depend on θ .

The distribution of the number of ones in N independent Bernoulli trials is binomial:

$$f_{\theta}(k) = \binom{N}{k} \theta^k (1-\theta)^{N-k}$$

Next, consider the joint distribution of $(x, \sum x_n)$. We have

$$f_{\theta}(x) = f_{\theta}\left(x, \sum x_n\right)$$

Thus, the conditional probability may be written

$$\begin{aligned} f_{\theta|k}(x) &= \frac{f_{\theta}(x, k)}{f_{\theta}(k)} \\ &= \frac{f_{\theta}(x)}{f_{\theta}(k)} \\ &= \frac{\theta^k (1-\theta)^{N-k}}{\binom{N}{k} \theta^k (1-\theta)^{N-k}} \\ &= \frac{1}{\binom{N}{k}} \end{aligned} \quad (1.11)$$

This shows that k is indeed a sufficient statistic for θ . The N values x_1, \dots, x_N can be replaced by the quantity k without losing information about θ .

Exercise 1.1

In the previous example (Example 1.18), suppose we wish to store in memory the information we possess about θ . Compare the savings, in terms of bits, we gain by storing the sufficient statistic k instead of the full sample x_1, \dots, x_N .

1.3.1.2 Determining Sufficient Statistics

In the example above (Example 1.18), we had to guess the sufficient statistic, and work out the conditional probability by hand. In general, this will be a tedious way to go about finding sufficient statistics. Fortunately, spotting sufficient statistics can be made easier by the Fisher-Neyman Factorization Theorem (Section 1.3.2).

1.3.1.3 Uses of Sufficient Statistics

Sufficient statistics have many uses in statistical inference problems. In hypothesis testing, the Likelihood Ratio Test⁴ can often be reduced to a sufficient statistic of the data. In parameter estimation, the Minimum Variance Unbiased Estimator (Section 4.1.1) of a parameter θ can be characterized by sufficient statistics and the Rao-Blackwell Theorem⁵.

⁴<http://workshop.molecularevolution.org/resources/lrt.php>

⁵http://en.wikipedia.org/wiki/Rao-Blackwell_theorem

1.3.1.4 Minimality and Completeness

Minimal sufficient statistics are, roughly speaking, sufficient statistics that cannot be compressed any more without losing information about the unknown parameter. **Completeness** is a technical characterization of sufficient statistics that allows one to prove minimality. These topics are covered in detail in this⁶ module.

Further examples of sufficient statistics may be found in the module on the Fisher-Neyman Factorization Theorem (Section 1.3.2).

1.3.2 The Fisher-Neyman Factorization Theorem⁷

Determining a sufficient statistic (Section 1.3.1) directly from the definition can be a tedious process. The following result can simplify this process by allowing one to spot a sufficient statistic directly from the functional form of the density or mass function.

Theorem 1.5: Fisher-Neyman Factorization Theorem

Let $f_\theta(x)$ be the density or mass function for the random vector x , parametrized by the vector θ . The statistic $t = T(x)$ is sufficient for θ if and only if there exist functions $a(x)$ (not depending on θ) and $b_\theta(t)$ such that

$$f_\theta(x) = a(x)b_\theta(t)$$

for all possible values of x .

In an earlier example (Example 1.18) we computed a sufficient statistic for a binary communication source (independent Bernoulli trials) from the definition. Using the above result, this task becomes substantially easier.

Example 1.19

1.3.2.1 Bernoulli Trials Revisited

Suppose $x_n \sim \text{Bernoulli}(\theta)$ are IID, $n = 1, \dots, N$. Denote $x = (x_1, \dots, x_N)^T$. Then

$$\begin{aligned} f_\theta(x) &= \prod_{n=1}^N \theta^{x_n} (1-\theta)^{1-x_n} \\ &= \theta^k (1-\theta)^{N-k} \\ &= a(x)b_\theta(k) \end{aligned} \tag{1.12}$$

where $k = \sum_{n=1}^N x_n$, $a(x) = 1$, and $b_\theta(k) = \theta^k (1-\theta)^{N-k}$. By the Fisher-Neyman factorization theorem, k is sufficient for θ .

The next example illustrates the application of the theorem to a continuous random variable.

Example 1.20

1.3.2.1 Normal Data with Unknown Mean

Consider a normally distributed random sample $x_1, \dots, x_N \sim \mathcal{N}(\theta, 1)$, IID, where θ is unknown. The joint pdf of $x = (x_1, \dots, x_N)^T$ is

$$f_\theta(x) = \prod_{n=1}^N f_\theta(x_n) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} e^{-\frac{1}{2} \sum_{n=1}^N (x_n - \theta)^2}$$

⁶<http://en.wikipedia.org/wiki/Completeness>

⁷This content is available online at <http://cnx.org/content/m11480/1.6/>.

We would like to rewrite $f_\theta(x)$ in the form of $a(x)b_\theta(t)$, where $\dim(t) < N$. At this point we require a trick-one that is commonly used when manipulating normal densities, and worth remembering. Define $\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n$, the sample mean. Then

$$\begin{aligned} f_\theta(x) &= \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} e^{-\frac{1}{2} \sum_{n=1}^N (x_n - \bar{x} + \bar{x} - \theta)^2} \\ &= \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} e^{-\frac{1}{2} \sum_{n=1}^N (x_n - \bar{x})^2 - 2(x_n - \bar{x})(\bar{x} - \theta) + (\bar{x} - \theta)^2} \end{aligned} \quad (1.13)$$

Now observe

$$\begin{aligned} \sum_{n=1}^N (x_n - \bar{x})(\bar{x} - \theta) &= (\bar{x} - \theta) \sum_{n=1}^N x_n - \bar{x} \\ &= (\bar{x} - \theta) (N\bar{x} - N\bar{x}) \\ &= 0 \end{aligned} \quad (1.14)$$

so the middle term vanishes. We are left with

$$f_\theta(x) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} e^{-\frac{1}{2} \sum_{n=1}^N (x_n - \bar{x})^2} e^{-\frac{1}{2} \sum_{n=1}^N (\bar{x} - \theta)^2}$$

where $a(x) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} e^{-\frac{1}{2} \sum_{n=1}^N (x_n - \bar{x})^2}$, $b_\theta(t) = e^{-\frac{1}{2} \sum_{n=1}^N (\bar{x} - \theta)^2}$, and $t = x$. Thus, the sample mean is a one-dimensional sufficient statistic for the mean.

1.3.2.1 Proof of Theorem

First, suppose $t = T(x)$ is sufficient for θ . By definition, $f_{\theta|T(x)=t}(x)$ is independent of θ . Let $f_\theta(x, t)$ denote the joint density or mass function for $(X, T(X))$. Observe $f_\theta(x) = f_\theta(x, t)$. Then

$$\begin{aligned} f_\theta(x) &= f_\theta(x, t) \\ &= f_{\theta|t}(x) f_\theta(t) \\ &= a(x) b_\theta(t) \end{aligned} \quad (1.15)$$

where $a(x) = f_{\theta|t}(x)$ and $b_\theta(t) = f_\theta(t)$. We prove the reverse implication for the discrete case only. The continuous case follows a similar argument, but requires a bit more technical work (*Scharf, pp.82[7]; Kay, pp.127[2]*).

Suppose the probability mass function for x can be written

$$f_\theta(x) = a(x) b_\theta(x)$$

where $t = T(x)$. The probability mass function for t is obtained by summing $f_\theta(x, t)$ over all x such that $T(x) = t$:

$$\begin{aligned} f_\theta(t) &= \sum_{T(x)=t} f_\theta(x, t) \\ &= \sum_{T(x)=t} a(x) b_\theta(t) \\ &= \sum_{T(x)=t} a(x) b_\theta(t) \end{aligned} \quad (1.16)$$

Therefore, the conditional mass function of x , given t , is

$$\begin{aligned} f_{\theta|t}(x) &= \frac{f_\theta(x, t)}{f_\theta(t)} \\ &= \frac{a(x) b_\theta(t)}{b_\theta(t)} \\ &= \frac{a(x)}{\sum_{T(x)=t} a(x)} \end{aligned} \quad (1.17)$$

This last expression does not depend on θ , so t is a sufficient statistic for θ . This completes the proof.

NOTE: From the proof, the Fisher-Neyman factorization gives us a formula for the conditional probability of x given t . In the discrete case we have

$$f(x | t) = \frac{a(x)}{\sum_{T(x)=t} a(x)}$$

An analogous formula holds for continuous random variables (*Scharf, pp.82[7]*).

1.3.2.2 Further Examples

The following exercises provide additional examples where the Fisher-Neyman factorization may be used to identify sufficient statistics.

Exercise 1.2

Uniform Measurements

Suppose x_1, \dots, x_N are independent and uniformly distributed on the interval $[\theta_1, \theta_2]$. Find a sufficient statistic for $\theta = (\theta_1, \theta_2)^T$.

NOTE: Express the likelihood $f_\theta(x)$ in terms of indicator functions.

Exercise 1.3

Poisson

Suppose x_1, \dots, x_N are independent measurements of a Poisson random variable with intensity parameter θ :

$$f_\theta(x) = \frac{e^{-\theta} \theta^x}{x!}, \quad x = 0, 1, 2, \dots$$

1.3.2.2.1

Find a sufficient statistic t for θ .

1.3.2.2.2

What is the conditional probability mass function of x , given t , where $x = (x_1, \dots, x_N)^T$?

Exercise 1.4

Normal with Unknown Mean and Variance

Consider $x_1, \dots, x_N \sim \mathcal{N}(\mu, \sigma^2)$, IID, where $\theta_1 = \mu$ and $\theta_2 = \sigma^2$ are both unknown. Find a sufficient statistic for $\theta = (\theta_1, \theta_2)^T$.

NOTE: Use the same trick as in Example 1.20.

Chapter 2

Signal Representation and Modeling

2.1 Signal Classifications and Properties¹

2.1.1 Introduction

This module will begin our study of signals and systems by laying out some of the fundamentals of signal classification. It is essentially an introduction to the important definitions and properties that are fundamental to the discussion of signals and systems, with a brief discussion of each.

2.1.2 Classifications of Signals

2.1.2.1 Continuous-Time vs. Discrete-Time

As the names suggest, this classification is determined by whether or not the time axis is **discrete** (countable) or **continuous** (Figure 2.1). A continuous-time signal will contain a value for all real numbers along the time axis. In contrast to this, a discrete-time signal², often created by sampling a continuous signal, will only have values at equally spaced intervals along the time axis.




Figure 2.1

¹This content is available online at <http://cnx.org/content/m10057/2.21/>.

²"Discrete-Time Signals" <http://cnx.org/content/m0009/latest/>

2.1.2.2 Analog vs. Digital

The difference between **analog** and **digital** is similar to the difference between continuous-time and discrete-time. However, in this case the difference involves the values of the function. Analog corresponds to a continuous set of possible function values, while digital corresponds to a discrete set of possible function values. An common example of a digital signal is a binary sequence, where the values of the function can only be one or zero.




Figure 2.2

2.1.2.3 Periodic vs. Aperiodic

Periodic signals³ repeat with some **period** T , while aperiodic, or nonperiodic, signals do not (Figure 2.3). We can define a periodic function through the following mathematical expression, where t can be any number and T is a positive constant:

$$f(t) = f(T + t) \quad (2.1)$$

The **fundamental period** of our function, $f(t)$, is the smallest value of T that the still allows (2.1) to be true.

³"Continuous Time Periodic Signals" <<http://cnx.org/content/m10744/latest/>>




Figure 2.3: (a) A periodic signal with period T_0 (b) An aperiodic signal

2.1.2.4 Finite vs. Infinite Length

As the name implies, signals can be characterized as to whether they have a finite or infinite length set of values. Most finite length signals are used when dealing with discrete-time signals or a given sequence of values. Mathematically speaking, $f(t)$ is a **finite-length signal** if it is **nonzero** over a finite interval

$$t_1 < f(t) < t_2$$

where $t_1 > -\infty$ and $t_2 < \infty$. An example can be seen in Figure 2.4. Similarly, an **infinite-length signal**, $f(t)$, is defined as nonzero over all real numbers:

$$\infty \leq f(t) \leq -\infty$$






Figure 2.4: Finite-Length Signal. Note that it only has nonzero values on a set, finite interval.

2.1.2.5 Causal vs. Anticausal vs. Noncausal


Causal signals are signals that are zero for all negative time, while **anticausal** are signals that are zero for all positive time. **Noncausal** signals are signals that have nonzero values in both positive and negative time (Figure 2.5).



(a)



(b)



(c)

Figure 2.5: (a) A causal signal (b) An anticausal signal (c) A noncausal signal

2.1.2.6 Even vs. Odd

An **even signal** is any signal f such that $f(t) = f(-t)$. Even signals can be easily spotted as they are **symmetric** around the vertical axis. An **odd signal**, on the other hand, is a signal f such that $f(t) = -f(-t)$ (Figure 2.6).




Figure 2.6: (a) An even signal (b) An odd signal

Using the definitions of even and odd signals, we can show that any signal can be written as a combination of an even and odd signal. That is, every signal has an odd-even decomposition. To demonstrate this, we have to look no further than a single equation.

$$f(t) = \frac{1}{2}(f(t) + f(-t)) + \frac{1}{2}(f(t) - f(-t)) \quad (2.2)$$

By multiplying and adding this expression out, it can be shown to be true. Also, it can be shown that $f(t) + f(-t)$ fulfills the requirement of an even function, while $f(t) - f(-t)$ fulfills the requirement of an odd function (Figure 2.7).

Example 2.1




Figure 2.7: (a) The signal we will decompose using odd-even decomposition (b) Even part: $e(t) = \frac{1}{2}(f(t) + f(-t))$ (c) Odd part: $o(t) = \frac{1}{2}(f(t) - f(-t))$ (d) Check: $e(t) + o(t) = f(t)$

2.1.2.7 Deterministic vs. Random

A **deterministic signal** is a signal in which each value of the signal is fixed and can be determined by a mathematical expression, rule, or table. Because of this the future values of the signal can be calculated from past values with complete confidence. On the other hand, a **random signal**⁴ has a lot of uncertainty about its behavior. The future values of a random signal cannot be accurately predicted and can usually only be guessed based on the averages⁵ of sets of signals (Figure 2.8).




Figure 2.8: (a) Deterministic Signal (b) Random Signal

Example 2.2

Consider the signal defined for all real t described by

$$f(t) = \begin{cases} \sin(2\pi t)/t & t \geq 1 \\ 0 & t < 1 \end{cases} \quad (2.3)$$

This signal is continuous time, analog, aperiodic, infinite length, causal, neither even nor odd, and, by definition, deterministic.

2.1.3 Signal Classifications Summary

This module describes just some of the many ways in which signals can be classified. They can be continuous time or discrete time, analog or digital, periodic or aperiodic, finite or infinite, and deterministic or random. We can also divide them based on their causality and symmetry properties. There are other ways to classify signals, such as boundedness, handedness, and continuity, that are not discussed here but will be described in subsequent modules.

⁴"Introduction to Random Signals and Processes" <<http://cnx.org/content/m10649/latest/>>

⁵"Random Processes: Mean and Variance" <<http://cnx.org/content/m10656/latest/>>

2.2 Signal Models

2.2.1 Linear Models⁶

Finding an MVUB estimator is a very difficult task, in general. However, a large number of signal processing problems can be represented by a **linear model** of the data.

2.2.1.1 Importance of Class of Linear Models

1. MVUB estimator within this class is immediately evident
2. Statistical performance analysis of linear models is very straightforward

2.2.1.2 General Form of Linear Model (LM)

$$x = H\theta + w$$

where x is the **observation vector**, H is the known matrix (**observation** or **system matrix**), θ is the unknown **parameter vector**, and w is the vector of White Gaussian noise $w \sim \mathcal{N}(0, \sigma^2 I)$.

Example 2.3

$$x_n = A + B_n + w_n, \quad n \in \{1, \dots, N\}$$

$$x = \begin{pmatrix} x_1 \\ \vdots \\ x_N \end{pmatrix}$$

$$w = \begin{pmatrix} w_1 \\ \vdots \\ w_N \end{pmatrix}$$

$$\theta = \begin{pmatrix} A \\ B \end{pmatrix}$$

$$H = \begin{pmatrix} 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ 1 & N \end{pmatrix}$$

2.2.1.3 Probability Model for LM

$$x = H\theta + w$$

$$x \sim p(x | \theta) = \mathcal{N}(H\theta, \sigma^2 I)$$

⁶This content is available online at <<http://cnx.org/content/m11432/1.2/>>.

2.2.1.4 CRLB and NVUB Estimator

$\theta = g(x)$ the MVUB estimator iff

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = I(\theta) (g(x) - \theta)$$

In the case of the LM,

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = \left(-\frac{1}{2\sigma^2} \right) \frac{\partial (x^T x - 2x^T H\theta + \theta^T H^T H\theta)}{\partial \theta}$$

Now using identities

$$\frac{\partial (b^T \theta)}{\partial \theta} = b$$

$$\frac{\partial (\theta^T A \theta)}{\partial \theta} = 2A\theta$$

for A symmetric.

We have

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = \frac{1}{\sigma^2} (H^T x - H^T H\theta)$$

Assuming $H^T H$ is invertible

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = \frac{H^T H}{\sigma^2} \left((H^T H)^{-1} H^T x - \theta \right)$$

which leads to

MVUB Estimator

$$\theta = (H^T H)^{-1} H^T x \quad (2.4)$$

Fisher Information Matrix

$$I(\theta) = \frac{H^T H}{\sigma^2} \quad (2.5)$$

$$\langle \theta \rangle = C_\theta = (I(\theta))^{-1} = \sigma^2 (H^T H)^{-1}$$

Theorem 2.1: MVUB Estimator for the LM

If the observed data can be modeled as

$$x = H\theta + w$$

where $w \sim \mathcal{N}(0, \sigma^2 I)$ and H is invertible. Then, the MVUB estimator is

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

and the covariance of θ is

$$C_\theta = \sigma^2 (H^T H)^{-1}$$

and θ attains the CRLB.

NOTE: $\theta \sim \mathcal{N}(\theta, \sigma^2 (H^T H)^{-1})$

2.2.1.5 Linear Model Examples

Example 2.4: Curve Fitting

Image not finished

Figure 2.9

Model:

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

where $\theta_1 + \theta_2 t_n + \dots + \theta_p t_n^{p-1}$ is a $(p-1)^{\text{st}}$ -order polynomial and $w(t_n) \sim \mathcal{N}(0, \sigma^2)$ iid. Therefore,

$$x = H\theta + w$$

$$x = \begin{pmatrix} x(t) \\ \vdots \\ x(t_n) \end{pmatrix}$$

$$\theta = \begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_p \end{pmatrix}$$

$$H = \begin{pmatrix} 1 & t_1 & \dots & t_1^{p-1} \\ 1 & t_2 & \dots & t_2^{p-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & t_N & \dots & t_N^{p-1} \end{pmatrix}$$

where H is the **Vandermonde matrix**. The MVUB estimator for θ is

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

Example 2.5: System Identification

Image not finished

Figure 2.10

$$H(z) = \sum_{k=0}^{m-1} h[k] z^{-k}$$

$$x[n] = \sum_{k=0}^{m-1} h[k] u[n-k] + w[n] \quad , \quad n \in \{0, \dots, N-1\}$$

Where $w[n] \sim \mathcal{N}(0, \sigma^2)$ i.i.d. Given x and u , estimate h .

In matrix form

$$x = \begin{pmatrix} u[0] & 0 & \dots & 0 \\ u[1] & u[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u[N-1] & u[N-2] & \dots & u[N-m] \end{pmatrix} \begin{pmatrix} h[0] \\ \vdots \\ \vdots \\ h[N-m] \end{pmatrix} + w$$

where

$$\begin{pmatrix} u[0] & 0 & \dots & 0 \\ u[1] & u[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u[N-1] & u[N-2] & \dots & u[N-m] \end{pmatrix} = H$$

and

$$\begin{pmatrix} h[0] \\ \vdots \\ \vdots \\ h[N-m] \end{pmatrix} = \theta$$

MVUB estimator

$$\theta = (H^T H)^{-1} H^T x \quad (2.6)$$

$$\langle \theta \rangle = \sigma^2 (H^T H)^{-1} = C_{\theta}$$

An important question in system identification is how to choose the input $u[n]$ to "probe" the system most efficiently.

First note that

$$\sigma(\theta_i)^2 = e_i^T C_{\theta} e_i$$

where $e_i = (0, \dots, 0, 1, 0, \dots, 0)^T$. Also, since C_{θ}^{-1} is symmetric positive definite, we can factor it by

$$C_{\theta}^{-1} = D^T D$$

where D is invertible.⁷ Note that

$$\left(e_i^T D^T (D^T)^{-1} e_i \right)^2 = 1 \quad (2.7)$$

The Schwarz inequality shows that (2.7) can become

$$1 \leq (e_i^T D^T D e_i) (e_i^T D^{-1} (D^T)^{-1} e_i) \quad (2.8)$$

⁷Cholesky factorization

$$1 = \left(e_i^T C_{\theta}^{-1} e_i \right) \left(e_i^T C_{\theta} e_i \right)$$

which leads to

$$\sigma(\theta_i)^2 \geq \frac{1}{e_i^T C_{\theta}^{-1} e_i} = \frac{\sigma^2}{(H^T H)_{i,i}}$$

The minimum variance is achieved when equality is attained in (2.8). This happens only if $\eta_1 = D e_i$ is proportional to $\eta_2 = D^T e_i$. That is, $\eta_1 = C \eta_2$ for some constant C . Equivalently,

$$D^T D e_i = c_i e_u \quad , \quad i \in \{1, 2, \dots, m\}$$

$$D^T D = C_{\theta}^{-1} = \frac{H^T H}{\sigma^2}$$

which leads to

$$\frac{H^T H}{\sigma^2} e_i = c_i e_i$$

Combining these equations in matrix form

$$H^T H = \sigma^2 \begin{pmatrix} c_1 & 0 & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ 0 & 0 & \dots & c_m \end{pmatrix}$$

Therefore, in order to minimize the variance of the MVUB estimator, $u[n]$ should be chosen to make $H^T H$ diagonal.

$$(H^T H)_{i,j} = \sum_{n=1}^N u[n-i] u[n-j] \quad , \quad i \wedge j \in \{1, \dots, m\}$$

For large N , this can be approximated to

$$(H^T H)_{i,j} \simeq \sum_{n=0}^{N-1-|i-j|} u[n] u[n+|i-j|]$$

using the autocorrelation of seq. $u[n]$.

NOTE: $u[n] = 0$ for $n < 0$ and $n > N - 1$, letting limit of sum $-\infty, \infty$ gives approx.

These steps lead to

$$H^T H \simeq N \begin{pmatrix} r_{uu}[0] & r_{uu}[1] & \dots & r_{uu}[m-1] \\ r_{uu}[1] & r_{uu}[0] & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ r_{uu}[m-1] & r_{uu}[m-2] & \dots & r_{uu}[0] \end{pmatrix}$$

where $\begin{pmatrix} r_{uu}[0] & r_{uu}[1] & \dots & r_{uu}[m-1] \\ r_{uu}[1] & r_{uu}[0] & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ r_{uu}[m-1] & r_{uu}[m-2] & \dots & r_{uu}[0] \end{pmatrix}$ is the Toeplitz autocorrelation matrix and

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

For $H^T H$ to be diagonal, we require $r[k] = 0$, $k \neq 0$. This condition is approximately realized if we take $u[n]$ to be a **pseudorandom noise sequence** (PRN)⁸. Furthermore, the PRN sequence simplifies the estimator computation:

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

$$\theta \simeq \frac{I}{N r_{uu}[0]} H^T x$$

which leads to

$$h[i] \simeq \frac{1}{N r_{uu}[0]} \sum_{n=0}^{N-1-i} u[n-i] x[n]$$

where $\sum_{n=0}^{N-1-i} u[n-i] x[n] = N r_{ux}[i]$. $r_{ux}[i]$ is the cross-correlation between input and output sequences.

Hence, the approximate MVUB estimator for large N with a PRN input is

$$h[i] = \frac{r_{ux}[i]}{r_{uu}[0]}, \quad i \in \{0, 1, \dots, m-1\}$$

$$r_{ux}[i] = \frac{1}{N} \sum_{n=0}^{N-1-i} u[n] x[n+i]$$

$$r_{uu}[0] = \frac{1}{N} \sum_{n=0}^{N-1} u^2[n]$$

2.2.1.6 CRLB for Signal in White Gaussian Noise

$$x_n = s_n(\theta) + w_n, \quad n \in \{1, \dots, N\}$$

$$p(x|\theta) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\left(\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - s_n(\theta))^2\right)}$$

$$\frac{\partial \log p(x|\theta)}{\partial \theta} = \frac{1}{\sigma^2} \sum_{n=1}^N (x_n - s_n(\theta)) \frac{\partial s_n(\theta)}{\partial \theta}$$

$$\frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} = \frac{1}{\sigma^2} \sum_{n=1}^N \left((x_n - s_n(\theta)) \frac{\partial^2 s_n(\theta)}{\partial \theta^2} - \left(\frac{\partial s_n(\theta)}{\partial \theta} \right)^2 \right)$$

⁸maximal length sequences

$$E \left[\frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} \right] = - \left(\frac{1}{\sigma^2} \sum_{n=1}^N \left(\frac{\partial s_n(\theta)}{\partial \theta} \right)^2 \right)$$
$$\sigma(\theta)^2 \geq \frac{\sigma^2}{\sum_{n=1}^N \left(\frac{\partial s_n(\theta)}{\partial \theta} \right)^2}$$

Chapter 3

Detection Theory

3.1 Hypothesis Testing¹

Suppose you measure a collection of scalars x_1, \dots, x_N . You believe the data is distributed in one of two ways. Your first model, call it H_0 , postulates the data to be governed by the density $f_0(x)$ (some fixed density). Your second model, H_1 , postulates a different density $f_1(x)$. These models, termed **hypotheses**, are denoted as follows:

$$H_0 : x_n \sim f_0(x), n = 1 \dots N$$

$$H_1 : x_n \sim f_1(x), n = 1 \dots N$$

A **hypothesis test** is a rule that, given a measurement x , makes a decision as to which hypothesis best "explains" the data.

Example 3.1

Suppose you are confident that your data is normally distributed with variance 1, but you are uncertain about the sign of the mean. You might postulate

$$H_0 : x_n \sim \mathcal{N}(-1, 1)$$

$$H_1 : x_n \sim \mathcal{N}(1, 1)$$

These densities are depicted in Figure 3.1.

¹This content is available online at <http://cnx.org/content/m11531/1.6/>.




Figure 3.1

Assuming each hypothesis is *a priori* equally likely, an intuitively appealing hypothesis test is to compute the sample mean $\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n$, and choose H_0 if $\bar{x} \leq 0$, and H_1 if $\bar{x} > 0$. As we will see later, this test is in fact optimal under certain assumptions.

3.1.1 Generalizations and Nomenclature

The concepts introduced above can be extended in several ways. In what follows we provide more rigorous definitions, describe different kinds of hypothesis testing, and introduce terminology.

3.1.1.1 Data

In the most general setup, the observation is a collection x_1, \dots, x_N of random vectors. A common assumption, which facilitates analysis, is that the data are independent and identically distributed (IID). The random vectors may be continuous, discrete, or in some cases mixed. It is generally assumed that all of the data is available at once, although for some applications, such as Sequential Hypothesis Testing², the data is a never ending stream.

²"Sequential Hypothesis Testing" <<http://cnx.org/content/m11242/latest/>>

3.1.1.2 Binary Versus M-ary Tests

When there are two competing hypotheses, we refer to a **binary** hypothesis test. When the number of hypotheses is $M \geq 2$, we refer to an **M-ary** hypothesis test. Clearly, binary is a special case of M -ary, but binary tests are accorded a special status for certain reasons. These include their simplicity, their prevalence in applications, and theoretical results that do not carry over to the M -ary case.

Example 3.2

3.1.1.2.1 Phase-Shift Keying

Suppose we wish to transmit a binary string of length r over a noisy communication channel. We assign each of the $M = 2^r$ possible bit sequences to a signal s^k , $k = \{1, \dots, M\}$ where

$$s_n^k = \cos\left(2\pi f_0 n + \frac{2\pi(k-1)}{M}\right)$$

This symboling scheme is known as **phase-shift keying** (PSK). After transmitting a signal across the noisy channel, the receiver faces an M -ary hypothesis testing problem:

$$H_0 : x = s^1 + w$$

$$\vdots$$

$$H_M : x = s^M + w$$

where $w \sim \mathcal{N}(0, \sigma^2 I)$.

In many binary hypothesis tests, one hypothesis represents the absence of a certain feature. In such cases, the hypothesis is usually labelled H_0 and called the **null** hypothesis. The other hypothesis is labelled H_1 and called the **alternative** hypothesis.

Example 3.3

3.1.1.2.1 Waveform Detection

Consider the problem of detecting a known signal $s = (s_1, \dots, s_N)^T$ in additive white Gaussian noise (AWGN). This scenario is common in sonar and radar systems. Denoting the data as $x = (x_1, \dots, x_N)^T$, our hypothesis testing problem is

$$H_0 : x = w$$

$$H_1 : x = s + w$$

where $w \sim \mathcal{N}(0, \sigma^2 I)$. H_0 is the null hypothesis, corresponding to the absence of a signal.

3.1.1.3 Tests and Decision Regions

Consider the general hypothesis testing problem where we have N d -dimensional observations x_1, \dots, x_N and M hypotheses. If the data are real-valued, for example, then a hypothesis test is a mapping

$$\phi : (\mathbb{R}^d)^N \rightarrow \{1, \dots, M\}$$

For every possible realization of the input, the test outputs a hypothesis. The test ϕ partitions the input space into a disjoint collection R_1, \dots, R_M , where

$$R_k = \{(x_1, \dots, x_N) \mid \phi(x_1, \dots, x_N) = k\}$$

The sets R_k are called **decision regions**. The boundary between two decision regions is a **decision boundary**. Figure 3.2 depicts these concepts when $d = 2$, $N = 1$, and $M = 3$.



Figure 3.2

3.1.1.4 Simple Versus Composite Hypotheses

If the distribution of the data under a certain hypothesis is fully known, we call it a **simple** hypothesis. All of the hypotheses in the examples above are simple. In many cases, however, we only know the distribution up to certain unknown parameters. For example, in a Gaussian noise model we may not know the variance of the noise. In this case, a hypothesis is said to be **composite**.

Example 3.4

Consider the problem of detecting the signal

$$s_n = \cos(2\pi f_0(n - k)) \quad n = \{1, \dots, N\}$$

where k is an unknown delay parameter. Then

$$H_0 : x = w$$

$$H_1 : x = s + w$$

is a binary test of a simple hypothesis (H_0) versus a composite alternative. Here we are assuming $w_n \sim \mathcal{N}(0, \sigma^2)$, with σ^2 known.

Often a test involving a composite hypothesis has the form

$$H_0 : \theta = \theta_0$$

$$H_1 : \theta \neq \theta_0$$

where θ_0 is fixed. Such problems are called **two-sided** because the composite alternative "lies on both sides of H_0 ." When θ is a scalar, the test

$$H_0 : \theta \leq \theta_0$$

$$H_1 : \theta > \theta_0$$

is called **one-sided**. Here, both hypotheses are composite.

Example 3.5

Suppose a coin turns up heads with probability p . We want to assess whether the coin is fair ($p = \frac{1}{2}$). We toss the coin N times and record x_1, \dots, x_N ($x_n = 1$ means heads and $x_n = 0$ means tails). Then

$$H_0 : p = \frac{1}{2}$$

$$H_1 : p \neq \frac{1}{2}$$

is a binary test of a simple hypothesis (H_0) versus a composite alternative. This is also a two-sided test.

3.1.2 Errors and Probabilities

In binary hypothesis testing, assuming at least one of the two models does indeed correspond to reality, there are four possible scenarios:

Case 1 - H_0 is true, and we declare H_0 to be true

Case 2 - H_0 is true, but we declare H_1 to be true

Case 3 - H_1 is true, and we declare H_1 to be true

Case 4 - H_1 is true, but we declare H_0 to be true

In cases 2 and 4, errors occur. The names given to these errors depend on the area of application. In statistics, they are called **type I** and **type II errors** respectively, while in signal processing they are known as a **false alarm** or a **miss**.

Consider the general binary hypothesis testing problem

$$H_0 : x \sim f_{\theta}(x), \theta \in \Theta_0$$

$$H_1 : x \sim f_{\theta}(x), \theta \in \Theta_1$$

If H_0 is simple, that is, $\Theta_0 = \{\theta_0\}$, then the **size** (denoted α), also called the **false-alarm probability** (P_F), is defined to be

$$\alpha = P_F = Pr[\theta_0, \text{declare } H_1]$$

When Θ_0 is composite, we define

$$\alpha = P_F = \sup_{\theta \in \Theta_0} (Pr[\theta, \text{declare } H_1])$$

For $\theta \in \Theta_1$, the **power** (denoted β), or **detection probability** (P_D), is defined to be

$$\beta = P_D = Pr[\theta, \text{declare } H_1]$$

The probability of a type II error, also called the **miss probability**, is

$$P_M = 1 - P_D$$

If H_1 is composite, then $\beta = \beta(\theta)$ is viewed as a function of θ .

3.1.3 Criteria in Hypothesis Testing

The design of a hypothesis test/detector often involves constructing the solution to an optimization problem. The optimality criteria used fall into two classes: Bayesian and frequent.

Representing the former approach is the Bayes Risk Criterion (Section 3.2.2). Representing the latter is the Neyman-Pearson Criterion (Section 3.2.4.1). These two approaches are developed at length in separate modules.

3.1.4 Statistics Versus Engineering Lingo

The following table, adapted from Kay, p.65[3], summarizes the different terminology for hypothesis testing from statistics and signal processing:

Statistics	Signal Processing
Hypothesis Test	Detector
Null Hypothesis	Noise Only Hypothesis
Alternate Hypothesis	Signal + Noise Hypothesis
Critical Region	Signal Present Decision Region
Type I Error	False Alarm
Type II Error	Miss
Size of Test (α)	Probability of False Alarm (P_F)
Power of Test (β)	Probability of Detection (P_D)

Table 3.1

3.2 Criteria

3.2.1 Criteria in Hypothesis Testing³

The criterion used in the previous section - minimize the average cost of an incorrect decision - may seem to be a contrived way of quantifying decisions. Well, often it is. For example, the Bayesian decision rule depends explicitly on the *a priori* probabilities; a rational method of assigning values to these - either by experiment or through true knowledge of the relative likelihood of each model - may be unreasonable. In this section, we develop alternative decision rules that try to answer such objections. One essential point will emerge from these considerations: **the fundamental nature of the decision rule does not change with choice of optimization criterion**. Even criteria remote from error measures can result in the likelihood ratio test (see this problem⁴). Such results do not occur often in signal processing and underline the likelihood ratio test's significance.

3.2.1.1 Maximum Probability of a Correct Decision

As only one model can describe any given set of data (the models are mutually exclusive), the probability of being correct P_c for distinguishing two models is given by

$$P_c = Pr[\text{say } \mathcal{M}_0 \text{ when } \mathcal{M}_0 \text{ true}] + Pr[\text{say } \mathcal{M}_1 \text{ when } \mathcal{M}_1 \text{ true}]$$

³This content is available online at <<http://cnx.org/content/m11228/1.4/>>.

⁴"Statistical Hypothesis Testing: Problems", Exercise 3 <<http://cnx.org/content/m11271/latest/#problem3>>

We wish to determine the optimum decision region placement. Expressing the probability correct in terms of the likelihood functions $p_{r|\mathcal{M}_i}(r)$, the *a priori* probabilities, and the decision regions,

$$P_c = \int \pi_0 p_{r|\mathcal{M}_0}(r) dr + \int \pi_1 p_{r|\mathcal{M}_1}(r) dr$$

We want to maximize P_c by selecting the decision regions \mathfrak{R}_0 and \mathfrak{R}_1 . The probability correct is maximized by associating each value of r with the largest term in the expression for P_c . Decision region \mathfrak{R}_0 , for example, is defined by the collection of values of r for which the first term is largest. As all of the quantities involved are non-negative, the decision rule maximizing the probability of a correct decision is

NOTE: Given r , choose \mathcal{M}_i for which the product $\pi_i p_{r|\mathcal{M}_i}(r)$ is largest.

Simple manipulations lead to the likelihood ratio test.

$$\frac{p_{r|\mathcal{M}_1}(r)}{p_{r|\mathcal{M}_0}(r)} \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\gtrless}} \frac{\pi_0}{\pi_1}$$

Note that if the Bayes' costs were chosen so that $C_{ii} = 0$ and $C_{ij} = C$, ($i \neq j$), we would have the same threshold as in the previous section.

To evaluate the quality of the decision rule, we usually compute the **probability of error** P_e rather than the probability of being correct. This quantity can be expressed in terms of the observations, the likelihood ratio, and the sufficient statistic.

$$\begin{aligned} P_e &= \pi_0 \int p_{r|\mathcal{M}_0}(r) dr + \pi_1 \int p_{r|\mathcal{M}_1}(r) dr \\ &= \pi_0 \int p_{\Lambda|\mathcal{M}_0}(\Lambda) d\Lambda + \pi_1 \int p_{\Lambda|\mathcal{M}_1}(\Lambda) d\Lambda \\ &= \pi_0 \int p_{\Upsilon|\mathcal{M}_0}(\Upsilon) d\Upsilon + \pi_1 \int p_{\Upsilon|\mathcal{M}_1}(\Upsilon) d\Upsilon \end{aligned} \quad (3.1)$$

When the likelihood ratio is non-monotonic, the first expression is most difficult to evaluate. When monotonic, the middle expression proves the most difficult. Furthermore, these expressions point out that the likelihood ratio and the sufficient statistic can be considered a function of the observations r ; hence, they are random variables and have probability densities for each model. Another aspect of the resulting probability of error is that **no other decision rule can yield a lower probability of error**. This statement is obvious as we minimized the probability of error in deriving the likelihood ratio test. The point is that these expressions represent a lower bound on performance (as assessed by the probability of error). This probability will be non-zero if the conditional densities overlap over some range of values of r , such as occurred in the previous example. In this region of overlap, the observed values are ambiguous: either model is consistent with the observations. Our "optimum" decision rule operates in such regions by selecting that model which is most likely (has the highest probability) of generating any particular value.

3.2.1.2 Neyman-Pearson Criterion

Situations occur frequently where assigning or measuring the *a priori* probabilities P_i is unreasonable. For example, just what is the *a priori* probability of a supernova occurring in any particular region of the sky? We clearly need a model evaluation procedure which can function without *a priori* probabilities. This kind of test results when the so-called Neyman-Pearson criterion is used to derive the decision rule. The ideas behind and decision rules derived with the Neyman-Pearson criterion (*Neyman and Pearson*[5]) will serve us well in sequel; their result is important!

Using nomenclature from radar, where model \mathcal{M}_1 represents the presence of a target and \mathcal{M}_0 its absence, the various types of correct and incorrect decisions have the following names (*Woodward, pp. 127-129*[8]).⁵

Detection - we say it's there when it is; $P_D = Pr(\text{say } \mathcal{M}_1 | \mathcal{M}_1 \text{ true})$

⁵In hypothesis testing, a false-alarm is known as a **type I error** and a miss a **type II error**.

False-alarm - we say it's there when it's not; $P_F = Pr(\text{say } \mathcal{M}_1 | \mathcal{M}_0 \text{ true})$

Miss - we say it's not there when it is; $P_M = Pr(\text{say } \mathcal{M}_0 | \mathcal{M}_1 \text{ true})$

The remaining probability $Pr[\text{say } \mathcal{M}_0 | \mathcal{M}_0 \text{ true}]$ has historically been left nameless and equals $1 - P_F$. We should also note that the detection and miss probabilities are related by $P_M = 1 - P_D$. As these are conditional probabilities, they do not depend on the *a priori* probabilities and the two probabilities P_F and P_D characterize the errors when **any** decision rule is used.

These two probabilities are related to each other in an interesting way. Expressing these quantities in terms of the decision regions and the likelihood functions, we have

$$P_F = \int p_{r|\mathcal{M}_0}(r) dr$$

$$P_D = \int p_{r|\mathcal{M}_1}(r) dr$$

As the region \mathfrak{R}_1 shrinks, **both** of these probabilities tend toward zero; as \mathfrak{R}_1 expands to engulf the entire range of observation values, they both tend toward unity. This rather direct relationship between P_D and P_F does not mean that they equal each other; in most cases, as \mathfrak{R}_1 expands, P_D increases more rapidly than P_F (we had better be right more often than we are wrong!). However, the "ultimate" situation where a rule is always right and never wrong ($P_D = 1$, $P_F = 0$) cannot occur when the conditional distributions overlap. Thus, to increase the detection probability we must also allow the false-alarm probability to increase. This behavior represents the fundamental tradeoff in hypothesis testing and detection theory.

One can attempt to impose a performance criterion that depends only on these probabilities with the consequent decision rule not depending on the *a priori* probabilities. The Neyman-Pearson criterion assumes that the false-alarm probability is constrained to be less than or equal to a specified value α while we attempt to maximize the detection probability P_D .

$$\max_{\mathfrak{R}_1} \{P_D\} \quad , \quad P_F \leq \alpha$$

A subtlety of the succeeding solution is that the underlying probability distribution functions may not be continuous, with the result that P_F can never equal the constraining value α . Furthermore, an (unlikely) possibility is that the optimum value for the false-alarm probability is somewhat less than the criterion value. Assume, therefore, that we rephrase the optimization problem by requiring that the false-alarm probability equal a value α' that is less than or equal to α .

This optimization problem can be solved using Lagrange multipliers (see Constrained Optimization⁶); we seek to find the decision rule that maximizes

$$F = P_D + \lambda(P_F - \alpha')$$

where λ is the Lagrange multiplier. This optimization technique amounts to finding the decision rule that maximizes F , then finding the value of the multiplier that allows the criterion to be satisfied. As is usual in the derivation of optimum decision rules, we maximize these quantities with respect to the decision regions. Expressing P_D and P_F in terms of them, we have

$$\begin{aligned} F &= \int p_{r|\mathcal{M}_1}(r) dr + \lambda \left(\int p_{r|\mathcal{M}_0}(r) dr - \alpha' \right) \\ &= -(\lambda\alpha') + \int p_{r|\mathcal{M}_1}(r) + \lambda p_{r|\mathcal{M}_0}(r) dr \end{aligned} \quad (3.2)$$

To maximize this quantity with respect to \mathfrak{R}_1 , we need only to integrate over those regions of r where the integrand is positive. The region \mathfrak{R}_1 thus corresponds to those values of r where $p_{r|\mathcal{M}_1}(r) > -(\lambda p_{r|\mathcal{M}_0}(r))$ and the resulting decision rule is

$$\frac{p_{r|\mathcal{M}_1}(r)}{p_{r|\mathcal{M}_0}(r)} \underset{\mathfrak{M}_0}{\overset{\mathfrak{M}_1}{\gtrless}} (-\lambda)$$

⁶"Constrained Optimization" <<http://cnx.org/content/m11223/latest/>>

The ubiquitous likelihood ratio test again appears; it **is** indeed the fundamental quantity in hypothesis testing. Using the logarithm of the likelihood ratio or the sufficient statistic, this result can be expressed as either

$$\ln(\Lambda(r)) \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\geq}} \ln(-\lambda)$$

or

$$\Upsilon(r) \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\geq}} \gamma$$

We have not as yet found a value for the threshold. The false-alarm probability can be expressed in terms of the Neyman-Pearson threshold in two (useful) ways.

$$\begin{aligned} P_F &= \int_{-\lambda}^{\infty} p_{\Lambda|\mathcal{M}_0}(\Lambda) d\Lambda \\ &= \int_{\gamma}^{\infty} p_{\Upsilon|\mathcal{M}_0}(\Upsilon) d\Upsilon \end{aligned} \quad (3.3)$$

One of these implicit equations must be solved for the threshold by setting P_F equal to α' . The selection of which to use is usually based on pragmatic considerations: the easiest to compute. From the previous discussion of the relationship between the detection and false-alarm probabilities, we find that to maximize P_D we must allow α' to be as large as possible while remaining less than α . Thus, we want to find the **smallest** value of $-\lambda$ (note the minus sign) consistent with the constraint. Computation of the threshold is problem-dependent, but a solution always exists.

Example 3.6

An important application of the likelihood ratio test occurs when r is a Gaussian random vector for each model. Suppose the models correspond to Gaussian random vectors having different mean values but sharing the same identity covariance.

- $\mathcal{M}_0: r \sim \mathcal{N}(0, \sigma^2 I)$
- $\mathcal{M}_1: r \sim \mathcal{N}(m, \sigma^2 I)$

Thus, r is of dimension L and has statistically independent, equal variance components. The vector of means $m = (m_0, \dots, m_{L-1})^T$ distinguishes the two models. The likelihood functions associated this problem are

$$\begin{aligned} p_{r|\mathcal{M}_0}(r) &= \prod_{l=0}^{L-1} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\left(1/2\left(\frac{r_l}{\sigma}\right)^2\right)} \\ p_{r|\mathcal{M}_1}(r) &= \prod_{l=0}^{L-1} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\left(1/2\left(\frac{r_l - m_l}{\sigma}\right)^2\right)} \end{aligned}$$

The likelihood ratio $\Lambda(r)$ becomes

$$\Lambda(r) = \frac{\prod_{l=0}^{L-1} e^{-\left(1/2\left(\frac{r_l - m_l}{\sigma}\right)^2\right)}}{\prod_{l=0}^{L-1} e^{-\left(1/2\left(\frac{r_l}{\sigma}\right)^2\right)}}$$

This expression for the likelihood ratio is complicated. In the Gaussian case (and many others), we use the logarithm to reduce the complexity of the likelihood ratio and form a sufficient statistic.

$$\begin{aligned} \ln(\Lambda(r)) &= \sum_{l=0}^{L-1} -1/2 \frac{(r_l - m_l)^2}{\sigma^2} + 1/2 \frac{r_l^2}{\sigma^2} \\ &= \frac{1}{\sigma^2} \sum_{l=0}^{L-1} m_l r_l - \frac{1}{2\sigma^2} \sum_{l=0}^{L-1} m_l^2 \end{aligned} \quad (3.4)$$

The likelihood ratio test then has the much simpler, but equivalent form

$$\sum_{l=0}^{L-1} (m_l r_l) \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\gtrless}} (\sigma^2 \ln(\eta)) + 1/2 \sum_{l=0}^{L-1} m_l^2$$

To focus on the model evaluation aspects of this problem, let's assume means be equal to a positive constant: $m_l = m$ (0).⁷

$$\sum_{l=0}^{L-1} r_l \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\gtrless}} \left(\frac{\sigma^2}{m} \ln(\eta) \right) + \frac{Lm}{2}$$

Note that all that need be known about the observations $\{r_l\}$ is their sum. This quantity is the sufficient statistic for the Gaussian problem: $\Upsilon(r) = \sum r_l$ and $\gamma = \sigma^2 \ln(\frac{\eta}{m}) + \frac{Lm}{2}$.

When trying to compute the probability of error or the threshold in the Neyman-Pearson criterion, we must find the conditional probability density of one of the decision statistics: the likelihood ratio, the log-likelihood, or the sufficient statistic. The log-likelihood and the sufficient statistic are quite similar in this problem, but clearly we should use the latter. One practical property of the sufficient statistic is that it usually simplifies computations. For this Gaussian example, the sufficient statistic is a Gaussian random variable under each model.

- \mathcal{M}_0 : $\Upsilon(r) \sim \mathcal{N}(0, L\sigma^2)$
- \mathcal{M}_1 : $\Upsilon(r) \sim \mathcal{N}(Lm, L\sigma^2)$

To find the probability of error from (3.1), we must evaluate the area under a Gaussian probability density function. These integrals are succinctly expressed in terms of $Q(x)$, which denotes the probability that a unit-variance, zero-mean Gaussian random variable exceeds x (see Probability and Stochastic Processes⁸). As $1 - Q(x) = Q(-x)$, the probability of error can be written as

$$P_e = \pi_1 Q\left(\frac{Lm - \gamma}{\sqrt{L}\sigma}\right) + \pi_0 Q\left(\frac{\gamma}{\sqrt{L}\sigma}\right)$$

An interesting special case occurs when $\pi_0 = 1/2 = \pi_1$. In this case, $\gamma = \frac{Lm}{2}$ and the probability of error becomes

$$P_e = Q\left(\frac{\sqrt{L}m}{2\sigma}\right)$$

As $Q(\cdot)$ is a monotonically decreasing function, the probability of error decreases with increasing values of the ratio $\frac{\sqrt{L}m}{2\sigma}$. However, as shown in this figure⁹, $Q(\cdot)$ decreases in a nonlinear fashion. Thus, increasing m by a factor of two may decrease the probability of error by a larger or a smaller factor; the amount of change depends on the initial value of the ratio.

To find the threshold for the Neyman-Pearson test from the expressions given on (3.3), we need the area under a Gaussian density.

$$\begin{aligned} P_F &= Q\left(\frac{\gamma}{\sqrt{L}\sigma}\right) \\ &= \alpha' \end{aligned} \tag{3.5}$$

As $Q(\cdot)$ is a monotonic and continuous function, we can now set α' equal to the criterion value α with the result

$$\gamma = \sqrt{L}\sigma Q^{-1}(\alpha)$$

⁷Why did the authors assume that the mean was positive? What would happen if it were negative?

⁸"Foundations of Probability Theory: Basic Definitions" <<http://cnx.org/content/m11245/latest/>>

⁹"The Gaussian Random Variable", Figure 1 <<http://cnx.org/content/m11250/latest/#fig1>>

where $Q^{-1}(\cdot)$ denotes the inverse function of $Q(\cdot)$. The solution of this equation cannot be performed analytically as no closed form expression exists for $Q(\cdot)$ (much less its inverse function); the criterion value must be found from tables or numerical routines. Because Gaussian problems arise frequently, the Table 3.2 accompanying table provides numeric values for this quantity at the decade points.

x	$Q^{-1}(x)$
10^{-1}	1.281
10^{-2}	2.396
10^{-3}	3.090
10^{-4}	3.719
10^{-5}	4.265
10^{-6}	4.754

Table 3.2: The table displays interesting values for $Q^{-1}(\cdot)$ that can be used to determine thresholds in the Neyman-Pearson variant of the likelihood ratio test. Note how little the inverse function changes for decade changes in its argument; $Q(\cdot)$ is indeed **very** nonlinear.

The detection probability is given by

$$P_D = Q\left(Q^{-1}(\alpha) - \frac{\sqrt{L}m}{\sigma}\right)$$

3.2.2 The Bayes Risk Criterion in Hypothesis Testing¹⁰

The design of a hypothesis test/detector often involves constructing the solution to an optimization problem. The optimality criteria used fall into two classes: Bayesian and frequent.

In the Bayesian setup, it is assumed that the *a priori* probability of each hypothesis occurring (π_i) is known. A cost C_{ij} is assigned to each possible outcome:

$$C_{ij} = Pr[\text{say } H_i \text{ when } H_j \text{ true}]$$

The optimal test/detector is the one that minimizes the Bayes risk, which is defined to be the expected cost of an experiment:

$$\bar{C} = \sum_{i,j} C_{ij} \pi_i Pr[\text{say } H_i \text{ when } H_j \text{ true}]$$

In the event that we have a binary problem, and both hypotheses are simple (Section 3.1.1.4: Simple Versus Composite Hypotheses), the decision rule that minimizes the Bayes risk can be constructed explicitly. Let us assume that the data is continuous (*i.e.*, it has a density) under each hypothesis:

$$H_0 : x \sim f_0(x)$$

$$H_1 : x \sim f_1(x)$$

Let R_0 and R_1 denote the decision regions (Section 3.1.1.3: Tests and Decision Regions) corresponding to the optimal test. Clearly, the optimal test is specified once we specify R_0 and $R_1 = R_0'$.

¹⁰This content is available online at <<http://cnx.org/content/m11533/1.6/>>.

The Bayes risk may be written

$$\begin{aligned}\bar{C} &= \sum_{ij=0}^1 C_{ij} \pi_i \int f_j(x) dx \\ &= \int C_{00} \pi_0 f_0(x) + C_{01} \pi_1 f_1(x) dx + \int C_{10} \pi_0 f_0(x) + C_{11} \pi_1 f_1(x) dx\end{aligned}\quad (3.6)$$

Recall that R_0 and R_1 **partition** the input space: they are disjoint and their union is the full input space. Thus, every possible input x belongs to precisely one of these regions. In order to minimize the Bayes risk, a measurement x should belong to the decision region R_i for which the corresponding integrand in the preceding equation is smaller. Therefore, the Bayes risk is minimized by assigning x to R_0 whenever

$$\pi_0 C_{00} f_0(x) + \pi_1 C_{01} f_1(x) < \pi_0 C_{10} f_0(x) + \pi_1 C_{11} f_1(x)$$

and assigning x to R_1 whenever this inequality is reversed. The resulting rule may be expressed concisely as

$$\Lambda(x) \equiv \frac{f_1(x)}{f_0(x)} \underset{H_0}{\overset{H_1}{\gtrless}} \frac{\pi_0 (C_{10} - C_{00})}{\pi_1 (C_{01} - C_{11})} \equiv \eta$$

Here, $\Lambda(x)$ is called the **likelihood ratio**, η is called the threshold, and the overall decision rule is called the Likelihood Ratio Test¹¹ (LRT). The expression on the right is called a **threshold**.

Example 3.7

An instructor in a course in detection theory wants to determine if a particular student studied for his last test. The observed quantity is the student's grade, which we denote by r . Failure may not indicate studiousness: conscientious students may fail the test. Define the models as

- \mathcal{M}_0 : did not study
- \mathcal{M}_1 : did study

The conditional densities of the grade are shown in Figure 3.3.




Figure 3.3: Conditional densities for the grade distributions assuming that a student did not study (\mathcal{M}_0) or did (\mathcal{M}_1) are shown in the top row. The lower portion depicts the likelihood ratio formed from these densities.

¹¹<http://workshop.molecularevolution.org/resources/lrt.php>

Based on knowledge of student behavior, the instructor assigns *a priori* probabilities of $\pi_0 = \frac{1}{4}$ and $\pi_1 = \frac{3}{4}$. The costs C_{ij} are chosen to reflect the instructor's sensitivity to student feelings: $C_{01} = 1 = C_{10}$ (an erroneous decision either way is given the same cost) and $C_{00} = 0 = C_{11}$. The likelihood ratio is plotted in Figure 3.3 and the threshold value η , which is computed from the *a priori* probabilities and the costs to be $\frac{1}{3}$, is indicated. The calculations of this comparison can be simplified in an obvious way.

$$\frac{r}{50} \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\gtrless}} \frac{1}{3}$$

or

$$r \underset{\mathcal{M}_0}{\overset{\mathcal{M}_1}{\gtrless}} \frac{50}{3} = 16.7$$

The multiplication by the factor of 50 is a simple illustration of the reduction of the likelihood ratio to a sufficient statistic. Based on the assigned costs and *a priori* probabilities, the optimum decision rule says the instructor must assume that the student did not study if the student's grade is less than 16.7; if greater, the student is assumed to have studied despite receiving an abysmally low grade such as 20. Note that as the densities given by each model overlap entirely: the possibility of making the wrong interpretation **always** haunts the instructor. However, no other procedure will be better!

A special case of the minimum Bayes risk rule, the minimum probability of error rule (Section 3.2.3), is used extensively in practice, and is discussed at length in another module.

3.2.2.1 Problems

Exercise 3.1

Denote $\alpha = Pr[\text{declare } H_1 \text{ when } H_0 \text{ true}]$ and $\beta = Pr[\text{declare } H_1 \text{ when } H_1 \text{ true}]$. Express the Bayes risk \bar{C} in terms of α and β , C_{ij} , and π_i . Argue that the optimal decision rule is not altered by setting $C_{00} = C_{11} = 0$.

Exercise 3.2

Suppose we observe x such that $\Lambda(x) = \eta$. Argue that it doesn't matter whether we assign x to R_0 or R_1 .

3.2.3 Minimum Probability of Error Decision Rule¹²

Consider the binary hypothesis test

$$\mathcal{H}_0 : x \sim f_0(x)$$

$$\mathcal{H}_1 : x \sim f_1(x)$$

Let π_i , denote the *a priori* probability of hypothesis \mathcal{H}_i . Suppose our decision rule declares " \mathcal{H}_0 is the true model" when $x \in R_0$, and it selects \mathcal{H}_1 when $x \in R_1$, where $R_1 = R_0'$. The probability of making an error, denoted P_e , is

$$\begin{aligned} P_e &= Pr[\text{declare } \mathcal{H}_0 \text{ and } \mathcal{H}_1 \text{ true}] + Pr[\text{declare } \mathcal{H}_1 \text{ and } \mathcal{H}_0 \text{ true}] \\ &= Pr[\mathcal{H}_1] Pr[\mathcal{H}_0 \mid \mathcal{H}_1] + Pr[\mathcal{H}_0] Pr[\mathcal{H}_1 \mid \mathcal{H}_0] \\ &= \int \pi_1 f_1(x) dx + \int \pi_0 f_0(x) dx \end{aligned} \tag{3.7}$$

In this module, we study the minimum probability of error decision rule, which selects R_0 and R_1 so as to minimize the above expression.

¹²This content is available online at <http://cnx.org/content/m11534/1.11/>.

Since an observation x falls into one and only one of the decision regions R_i , in order to minimize P_e , we assign x to the region for which the corresponding integrand in (3.7) is smaller. Thus, we select $x \in R_0$ if $\pi_1 f_1(x) < \pi_0 f_0(x)$, and $x \in R_1$ if the inequality is reversed. This decision rule may be summarized concisely as

$$\Lambda(x) \equiv \frac{f_1(x)}{f_0(x)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \frac{\pi_0}{\pi_1} \equiv \eta$$

Here, $\Lambda(x)$ is called the **likelihood ratio**, η is called a **threshold**, and the overall decision rule is called the Likelihood Ratio Test¹³.

Example 3.8

3.2.3.1 Normal with Common Variance, Uncommon Means

Consider the binary hypothesis test of a scalar x

$$\mathcal{H}_0 : x \sim \mathcal{N}(0, \sigma^2)$$

$$\mathcal{H}_1 : x \sim \mathcal{N}(\mu, \sigma^2)$$

where μ and σ^2 are known, positive quantities. Suppose we observe a single measurement x . The likelihood ratio is

$$\begin{aligned} \Lambda(x) &= \frac{\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}}{\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}} \\ &= e^{\frac{1}{\sigma^2} \left(\mu x - \frac{\mu^2}{2} \right)} \end{aligned} \quad (3.8)$$

and so the minimum probability of error decision rule is

$$e^{\frac{1}{\sigma^2} \left(\mu x - \frac{\mu^2}{2} \right)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \frac{\pi_0}{\pi_1} = \eta$$


The expression for $\Lambda(x)$ is somewhat complicated. By applying a sequence of monotonically increasing functions to both sides, we can obtain a simplified expression for the optimal decision rule without changing the rule. In this example, we apply the natural logarithm and rearrange terms to arrive at

$$x \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \frac{\sigma^2}{\mu} \ln(\eta) + \frac{\mu}{2} \equiv \gamma$$


Here we have used the assumption $\mu > 0$. If $\mu < 0$, then dividing by μ would reverse the inequalities.

This form of the decision rule is much simpler: we just compare the observed value x to a threshold γ . Figure 3.4 depicts the two candidate densities and a possible value of γ . If each hypothesis is *a priori* equally likely ($\pi_0 = \pi_1 = \frac{1}{2}$), then $\gamma = \frac{\mu}{2}$. Figure 3.4 illustrates the case where $\pi_0 > \pi_1$ ($\gamma > \frac{\mu}{2}$).

¹³<http://workshop.molecularevolution.org/resources/lrt.php>



(a)



(b)

Figure 3.4: The two candidate densities, and a threshold corresponding to $\pi_0 > \pi_1$

If we plot the two densities so that each is weighted by its *a priori* probability of occurring, the two curves will intersect at the threshold γ (see Figure 3.5). (Can you explain why this is? Think back to our derivation of the LRT). This plot also offers a way to visualize the probability of error. Recall

$$\begin{aligned}
 P_e &= \int \pi_1 f_1(x) dx + \int \pi_0 f_0(x) dx \\
 &= \int \pi_1 f_1(x) dx + \int \pi_0 f_0(x) dx \\
 &= \pi_1 P_M + \pi_0 P_F
 \end{aligned} \tag{3.9}$$

where P_M and P_F denote the miss and false alarm probabilities, respectively. These quantities are depicted in Figure 3.5.




Figure 3.5: The candidate densities weighted by their *a priori* probabilities. The shaded region is the probability of error for the optimal decision rule.

We can express P_M and P_F in terms of the Q-function¹⁴ as

$$P_e = \pi_1 Q\left(\frac{\mu - \gamma}{\sigma}\right) + \pi_0 Q\left(\frac{\gamma}{\sigma}\right)$$

When $\pi_0 = \pi_1 = \frac{1}{2}$, we have $\gamma = \frac{\mu}{2}$, and the error probability is

$$P_e = Q\left(\frac{\mu}{2\sigma}\right)$$

Since $Q(x)$ is monotonically decreasing, this says that the "difficulty" of the detection problem decreases with decreasing σ and increasing μ .

In the preceding example, computation of the probability of error involved a one-dimensional integral. If we had multiple observations, or vector-valued data, generalizing this procedure would involve multi-dimensional integrals over potentially complicated decision regions. Fortunately, in many cases, we can avoid this problem through the use of sufficient statistics (Section 1.3.1).

Example 3.9

Suppose we have the same test as in the previous example (Example 3.8), but now we have N independent observations:

$$\mathcal{H}_0 : x_n \sim \mathcal{N}(0, \sigma^2), n = 1, \dots, N$$

$$\mathcal{H}_1 : x_n \sim \mathcal{N}(\mu, \sigma^2), n = 1, \dots, N$$

¹⁴"The Q-function" <<http://cnx.org/content/m11537/latest/>>

where $\mu > 0$ and $\sigma^2 > 0$ and both are known. The likelihood ratio is

$$\begin{aligned}
 \Lambda(x) &= \frac{\prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_n-\mu)^2}{2\sigma^2}}}{\prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x_n^2}{2\sigma^2}}} \\
 &= \frac{e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n-\mu)^2}}{e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N x_n^2}} \\
 &= e^{\frac{1}{2\sigma^2} \sum_{n=1}^N 2x_n\mu - \mu^2} \\
 &= e^{\frac{1}{\sigma^2} \left(\mu \sum_{n=1}^N x_n - \frac{N\mu^2}{2} \right)}
 \end{aligned} \tag{3.10}$$

As in the previous example (Example 3.8), we may apply the natural logarithm and rearrange terms to obtain an equivalent form of the LRT:

$$t \equiv \sum_{n=1}^N x_n \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \frac{\sigma^2}{\mu} \ln(\eta) + \frac{N\mu}{2} \equiv \gamma$$

The scalar quantity t is a sufficient statistic for the mean. In order to evaluate the probability of error without resorting to a multi-dimensional integral, we can express P_e in terms of t as

$$P_e = \pi_1 Pr[t < \gamma \mid \mathcal{H}_1 \text{ true}] + \pi_0 Pr[t > \gamma \mid \mathcal{H}_0 \text{ true}]$$

Now t is a linear combination of normal variates, so it is itself normal. In particular, we have $t = Ax$, where $\begin{pmatrix} 1 & \dots & 1 \end{pmatrix}$ is an N -dimensional row vector of 1's, and x is multivariate normal with mean 0 or $\mu = (\mu, \dots, \mu)^T$, and covariance $\sigma^2 I$. Thus we have

$$t \mid \mathcal{H}_0 \sim \mathcal{N}(A0, A\sigma^2 I A^T) = \mathcal{N}(0, N\sigma^2)$$

$$t \mid \mathcal{H}_1 \sim \mathcal{N}(A\mu, A\sigma^2 I A^T) = \mathcal{N}(N\mu, N\sigma^2)$$

Therefore, we may write P_e in terms of the Q-function¹⁵ as

$$P_e = \pi_1 Q\left(\frac{N\mu - \gamma}{\sqrt{N}\sigma}\right) + \pi_0 Q\left(\frac{\gamma}{\sqrt{N}\sigma}\right)$$

In the special case $\pi_0 = \pi_1 = \frac{1}{2}$,

$$P_e = Q\left(\frac{\sqrt{N}\mu}{\sigma}\right)$$

Since Q is monotonically decreasing, this result provides mathematical support for something that is intuitively obvious: The performance of our decision rule improves with increasing N and μ , and decreasing σ .

NOTE: In the context of signal processing, the foregoing problem may be viewed as the problem of detecting a constant (DC) signal in additive white Gaussian noise:

$$\mathcal{H}_0 : x_n = w_n, n = 1, \dots, N$$

$$\mathcal{H}_1 : x_n = A + w_n, n = 1, \dots, N$$

where A is a known, fixed amplitude, and $w_n \sim \mathcal{N}(0, \sigma^2)$. Here A corresponds to the mean μ in the example.

The next example explores the minimum probability of error decision rule in a **discrete** setting.

Example 3.10

¹⁵"The Q-function" <<http://cnx.org/content/m11537/latest/>>

3.2.3.1 Repetition Code

Suppose we have a friend who is trying to transmit a bit (0 or 1) to us over a noisy channel. The channel causes an error in the transmission (that is, the bit is flipped) with probability p , where $0 \leq p < \frac{1}{2}$, and p is known. In order to increase the chance of a successful transmission, our friend sends the same bit N times. Assume the N transmissions are statistically independent. Under these assumptions, the bits you receive are Bernoulli random variables: $x_n \sim \text{Bernoulli}(\theta)$. We are faced with the following hypothesis test:

\mathcal{H}_0	$\theta = p$	0 sent
\mathcal{H}_1	$\theta = 1 - p$	1 sent

Table 3.3

We decide to decode the received sequence $x = (x_1, \dots, x_N)^T$ by minimizing the probability of error. The likelihood ratio is

$$\begin{aligned} \Lambda(x) &= \frac{\prod_{n=1}^N (1-p)^{x_n} p^{1-x_n}}{\prod_{n=1}^N p^{x_n} (1-p)^{1-x_n}} \\ &= \frac{(1-p)^k p^{N-k}}{p^k (1-p)^{N-k}} \\ &= \left(\frac{1-p}{p}\right)^{2k-N} \end{aligned} \quad (3.11)$$

where $k = \sum_{n=1}^N x_n$ is the number of 1s received.

NOTE: k is a sufficient statistic (Section 1.3.1) for θ .

The LRT is

$$\left(\frac{1-p}{p}\right)^{2k-N} \frac{\mathcal{H}_1}{\mathcal{H}_0} \frac{\pi_0}{\pi_1} = \eta$$

Taking the natural logarithm of both sides and rearranging, we have

$$k \frac{\mathcal{H}_1}{\mathcal{H}_0} \frac{N}{2} + \frac{1}{2} \frac{\ln(\eta)}{\ln\left(\frac{1-p}{p}\right)} = \gamma$$

In the case that both hypotheses are equally likely, the minimum probability of error decision is the "majority-vote" rule: Declare \mathcal{H}_1 if there are more 1s than 0s, declare \mathcal{H}_0 otherwise. In the event $k = \gamma$, we may decide arbitrarily; the probability of error is the same either way. Let's adopt the convention that \mathcal{H}_0 is declared in this case.

To compute the probability of error of the optimal rule, write

$$\begin{aligned} P_e &= \pi_0 Pr[\text{declare } \mathcal{H}_1 \mid \mathcal{H}_0 \text{ true}] + \pi_1 Pr[\text{declare } \mathcal{H}_0 \mid \mathcal{H}_1 \text{ true}] \\ &= \pi_0 Pr[k > \gamma \mid \mathcal{H}_0 \text{ true}] + \pi_1 Pr[k \leq \gamma \mid \mathcal{H}_1 \text{ true}] \end{aligned} \quad (3.12)$$

Now k is a binomial random variable, $k \sim \text{Binomial}(N, \theta)$, where θ depends on which hypothesis is true. We have

$$\begin{aligned} Pr[k > \gamma \mid \mathcal{H}_0] &= \sum_{k=\lfloor \gamma \rfloor + 1}^N f_0(k) \\ &= \sum_{k=\lfloor \gamma \rfloor + 1}^N \binom{N}{k} p^k (1-p)^{N-k} \end{aligned} \quad (3.13)$$

and

$$Pr[k \leq \gamma \mid \mathcal{H}_1] = \sum_{k=0}^{\lfloor \gamma \rfloor} \binom{N}{k} (1-p)^k p^{N-k}$$

Using these formulae, we may compute P_e explicitly for given values of N , p , π_0 and π_1 .

3.2.3.1 MAP Interpretation

The likelihood ratio test is one way of expressing the minimum probability of error decision rule. Another way is

Rule 3.1:

Declare hypothesis i such that $\pi_i f_i(x)$ is maximal.

This rule is referred to as the **maximum a posteriori**, or **MAP** rule, because the quantity $\pi_i f_i(x)$ is proportional to the posterior probability of hypothesis i . This becomes clear when we write $\pi_i = Pr[\mathcal{H}_i]$ and $f_i(x) = f(x|\mathcal{H}_i)$. Then, by Bayes rule, the posterior probability of \mathcal{H}_i given the data is

$$Pr[\mathcal{H}_i | x] = \frac{Pr[\mathcal{H}_i] f(x|\mathcal{H}_i)}{f(x)}$$

Here $f(x)$ is the unconditional density or mass function for x , which is effectively a constant when trying to maximize with respect to i .

According to the MAP interpretation, the optimal decision boundary is the locus of points where the weighted densities (in the continuous case) $\pi_i f_i(x)$ intersect one another. This idea is illustrated in Example 3.9.

3.2.3.2 Multiple Hypotheses

One advantage the MAP formulation of the minimum probability of error decision rule has over the LRT is that it generalizes easily to M -ary hypothesis testing. If we are to choose between hypotheses \mathcal{H}_i , $i = \{1, \dots, M\}$, the optimal rule is still the MAP rule (Rule 3.1, p. 65)

3.2.3.3 Special Case of Bayes Risk

The Bayes risk criterion (Section 3.2.2) for constructing decision rules assigns a cost C_{ij} to the outcome of declaring \mathcal{H}_i when \mathcal{H}_j is in effect. The probability of error is simply a special case of the Bayes risk corresponding to $C_{00} = C_{11} = 0$ and $C_{01} = C_{10} = 1$. Therefore, the form of the minimum probability of error decision rule is a specialization of the minimum Bayes risk decision rule: both are likelihood ratio tests. The different costs in the Bayes risk formulation simply shift the threshold to favor one hypothesis over the other.

3.2.3.4 Problems

Exercise 3.3

Generally speaking, when is the probability of error **zero** for the optimal rule? Phrase your answer in terms of the distributions underlying each hypothesis. Does the LRT agree with your answer in this case?

Exercise 3.4

Suppose we measure N independent values x_1, \dots, x_N . We know the variance of our measurements ($\sigma^2 = 1$), but are unsure whether the data obeys a Laplacian or Gaussian probability law:

$$\mathcal{H}_0 : f_0(x) = \frac{1}{\sqrt{2}} e^{-(\sqrt{2}|x|)}$$

$$\mathcal{H}_1 : f_1(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$$

3.2.3.4.1

Show that the two densities have the same mean and variance, and plot the densities on the same graph.

3.2.3.4.2

Find the likelihood ratio.

3.2.3.4.3

Determine the decision regions for different values of the threshold η . Consider all possible values of $\eta > 0$

NOTE: There are three distinct cases.

3.2.3.4.4

Draw the decision regions and decision boundaries for $\eta = \{\frac{1}{2}, 1, 2\}$.

3.2.3.4.5

Assuming the two hypotheses are equally likely, compute the probability of error. Your answer should be a number.

Exercise 3.5**3.2.3.4.1 Arbitrary Means and Covariances**

Consider the hypothesis testing problem

$$\mathcal{H}_0 : x \sim \mathcal{N}(\mu_0, \Sigma_0)$$

$$\mathcal{H}_1 : x \sim \mathcal{N}(\mu_1, \Sigma_1)$$

where $\mu_0 \in \mathbb{R}^d$ and $\mu_1 \in \mathbb{R}^d$, and Σ_0, Σ_1 are positive definite, symmetric $d \times d$ matrices. Write down the likelihood ratio test, and simplify, for the following cases. In each case, provide a geometric description of the decision boundary.

3.2.3.4.1.1

$\Sigma_0 = \Sigma_1$, but $\mu_0 \neq \mu_1$.

3.2.3.4.1.2

$\mu_0 = \mu_1$, but $\Sigma_0 \neq \Sigma_1$.

3.2.3.4.1.3

$\mu_0 \neq \mu_1$ and $\Sigma_0 \neq \Sigma_1$.

Exercise 3.6

Suppose we observe N independent realizations of a Poisson random variable k with intensity parameter λ :

$$f(k) = \frac{e^{-\lambda} \lambda^k}{k!}$$

We must decide which of two intensities is in effect:

$$\mathcal{H}_0 : \lambda = \lambda_0$$

$$\mathcal{H}_1 : \lambda = \lambda_1$$

where $\lambda_0 < \lambda_1$.

3.2.3.4.1

Give the minimum probability of error decision rule.

3.2.3.4.2

Simplify the LRT to a test statistic involving only a sufficient statistic. Apply a monotonically increasing transformation to simplify further.

3.2.3.4.3

Determine the distribution of the sufficient statistic under both hypotheses.

NOTE: Use the characteristic function to show that a sum of IID Poisson variates is again Poisson distributed.

3.2.3.4.4

Derive an expression for the probability of error.

3.2.3.4.5

Assuming the two hypotheses are equally likely, and $\lambda_0 = 5$ and $\lambda_1 = 6$, what is the minimum number N of observations needed to attain a probability of error no greater than 0.01?

NOTE: If you have numerical trouble, try rewriting the log-factorial so as to avoid evaluating the factorial of large integers.

Exercise 3.7

In Example 3.10, suppose $\pi_0 = \pi_1 = \frac{1}{2}$, and $p = 0.1$. What is the smallest value of N needed to ensure $P_e \leq 0.01$?

3.2.4 Neyman-Person

3.2.4.1 The Neyman-Pearson Criterion¹⁶

In hypothesis testing (Section 3.1), as in all other areas of statistical inference, there are two major schools of thought on designing good tests: Bayesian and frequentist (or classical). Consider the simple binary hypothesis testing problem

$$\mathcal{H}_0 : x \sim f_0(x)$$

$$\mathcal{H}_1 : x \sim f_1(x)$$

In the Bayesian setup, the prior probability $\pi_i = Pr[\mathcal{H}_i]$ of each hypothesis occurring is assumed known. This approach to hypothesis testing is represented by the minimum Bayes risk criterion (Section 3.2.2) and the minimum probability of error criterion (Section 3.2.3).

In some applications, however, it may not be reasonable to assign an *a priori* probability to a hypothesis. For example, what is the *a priori* probability of a supernova occurring in any particular region of the sky? What is the prior probability of being attacked by a ballistic missile? In such cases we need a decision rule that does not depend on making assumptions about the *a priori* probability of each hypothesis. Here the Neyman-Pearson criterion offers an alternative to the Bayesian framework.

The Neyman-Pearson criterion is stated in terms of certain probabilities (Section 3.1.2: Errors and Probabilities) associated with a particular hypothesis test. The relevant quantities are summarized in Table 3.4. Depending on the setting, different terminology is used.

Statistics		Signal Processing		
Probability	Name	Notation	Name	Notation
P_0 (declare \mathcal{H}_1)	size	α	false-alarm probability	P_F
P_1 (declare \mathcal{H}_1)	power	β	detection probability	P_D

Table 3.4

Here P_i (declare \mathcal{H}_j) denotes the probability that we declare hypothesis \mathcal{H}_j to be in effect when \mathcal{H}_i is actually in effect. The probabilities P_0 (declare \mathcal{H}_0) and P_1 (declare \mathcal{H}_0) (sometimes called the **miss** probability), are equal to $1 - P_F$ and $1 - P_D$, respectively. Thus, P_F and P_D represent the two degrees of freedom in a binary hypothesis test. Note that P_F and P_D do not involve *a priori* probabilities of the hypotheses.

These two probabilities are related to each other through the decision regions (Section 3.1.1.3: Tests and Decision Regions). If R_1 is the decision region for \mathcal{H}_1 , we have

$$P_F = \int f_0(x) dx$$

$$P_D = \int f_1(x) dx$$

The densities $f_i(x)$ are nonnegative, so as R_1 shrinks, both probabilities tend to zero. As R_1 expands, both tend to one. The ideal case, where $P_D = 1$ and $P_F = 0$, cannot occur unless the distributions do not overlap (*i.e.*, $\int f_0(x) f_1(x) dx = 0$). Therefore, in order to increase P_D , we must also increase P_F . This represents the fundamental tradeoff in hypothesis testing and detection theory.

Example 3.11

Consider the simple binary hypothesis test of a scalar measurement x :

$$\mathcal{H}_0 : x \sim \mathcal{N}(0, 1)$$

¹⁶This content is available online at <<http://cnx.org/content/m11548/1.2/>>.

$$\mathcal{H}_1 : x \sim \mathcal{N}(1, 1)$$

Suppose we use a threshold test


$$\begin{array}{l} \mathcal{H}_1 \\ x \gtrless \gamma \\ \mathcal{H}_0 \end{array}$$

where $\gamma \in \mathbb{R}$ is a free parameter. Then the false alarm and detection probabilities are


$$P_F = \int_{\gamma}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt = Q(\gamma)$$

$$P_D = \int_{\gamma}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{(t-1)^2}{2}} dt = Q(\gamma - 1)$$

where Q denotes the Q-function¹⁷. These quantities are depicted in Figure 3.6.



(a)



(b)

Figure 3.6: False alarm and detection values for a certain threshold.

Since the Q -function is monotonically decreasing, it is evident that both P_D and P_F decay to zero as γ increases. There is also an explicit relationship

$$P_D = Q(Q^{-1}(P_F) - 1)$$

¹⁷"The Q-function" <<http://cnx.org/content/m11537/latest/>>

A common means of displaying this relationship is with a **receiver operating characteristic** (ROC) curve, which is nothing more than a plot of P_D versus P_F (Figure 3.7).




Figure 3.7: ROC curve for this example.

3.2.4.1.1 The Neyman-Pearson Lemma: A First Look

The Neyman-Pearson criterion says that we should construct our decision rule to have maximum probability of detection while not allowing the probability of false alarm to exceed a certain value α . In other words, the optimal detector according to the Neyman-Pearson criterion is the solution to the following constrained optimization problem:

3.2.4.1.1.1 Neyman-Pearson Criterion

$$\max \{P_D\}, \text{ such that } P_F \leq \alpha \quad (3.14)$$

The maximization is over all decision rules (equivalently, over all decision regions R_0, R_1). Using different terminology, the Neyman-Pearson criterion selects the **most powerful test of size (not exceeding) α** .

Fortunately, the above optimization problem has an explicit solution. This is given by the celebrated **Neyman-Pearson lemma**, which we now state. To ease the exposition, our initial statement of this result only applies to continuous random variables, and places a technical condition on the densities. A more general statement is given later in the module.

Theorem 3.1: Neyman-Pearson Lemma: initial statement

Consider the test

$$\mathcal{H}_0 : x \sim f_0(x)$$

$$\mathcal{H}_1 : x \sim f_1(x)$$

where $f_i(x)$ is a density. Define $\Lambda(x) = \frac{f_1(x)}{f_0(x)}$, and assume that $\Lambda(x)$ satisfies the condition that for each $\gamma \in \mathbb{R}$, $\Lambda(x)$ takes on the value γ with probability zero under hypothesis \mathcal{H}_0 . The solution to the optimization problem in (3.14) is given by

$$\Lambda(x) = \frac{f_1(x)}{f_0(x)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta$$

where η is such that

$$P_F = \int f_0(x) dx = \alpha$$

If $\alpha = 0$, then $\eta = \infty$. The optimal test is unique up to a set of probability zero under \mathcal{H}_0 and \mathcal{H}_1 .

The optimal decision rule is called the **likelihood ratio test**. $\Lambda(x)$ is the **likelihood ratio**, and η is a **threshold**. Observe that neither the likelihood ratio nor the threshold depends on the *a priori* probabilities $Pr[\mathcal{H}_i]$. they depend only on the conditional densities f_i and the size constraint α . The threshold can often be solved for as a function of α , as the next example shows.

Example 3.12

Continuing with Example 3.11, suppose we wish to design a Neyman-Pearson decision rule with size constraint α . We have

$$\begin{aligned} \Lambda(x) &= \frac{\frac{1}{\sqrt{2\pi}} e^{-\frac{(x-1)^2}{2}}}{\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}} \\ &= e^{x - \frac{1}{2}} \end{aligned} \tag{3.15}$$

By taking the natural logarithm of both sides of the LRT and rearranging terms, the decision rule is not changed, and we obtain

$$x \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \ln(\eta) + \frac{1}{2} \equiv \gamma$$

Thus, the optimal rule is in fact a thresholding rule like we considered in Example 3.11. The false-alarm probability was seen to be

$$P_F = Q(\gamma)$$

Thus, we may express the value of γ required by the Neyman-Pearson lemma in terms of α :

$$\gamma = Q^{-1}(\alpha)$$

3.2.4.1.2 Sufficient Statistics and Monotonic Transformations

For hypothesis testing involving multiple or vector-valued data, direct evaluation of the size (P_F) and power (P_D) of a Neyman-Pearson decision rule would require integration over multi-dimensional, and potentially complicated decision regions. In many cases, however, this can be avoided by simplifying the LRT to a test of the form

$$t \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma$$

where the test statistic $t = T(x)$ is a sufficient statistic (Section 1.3.1) for the data. Such a simplified form is arrived at by modifying both sides of the LRT with monotonically increasing transformations, and by algebraic simplifications. Since the modifications do not change the decision rule, we may calculate P_F and P_D in terms of the sufficient statistic. For example, the false-alarm probability may be written

$$\begin{aligned} P_F &= Pr[\text{declare } \mathcal{H}_1] \\ &= \int f_0(t) dt \end{aligned} \tag{3.16}$$

where $f_0(t)$ denotes the density of t under \mathcal{H}_0 . Since t is typically of lower dimension than x , evaluation of P_F and P_D can be greatly simplified. The key is being able to reduce the LRT to a threshold test involving a sufficient statistic **for which we know the distribution**.

Example 3.13

3.2.4.1.2.1 Common Variances, Uncommon Means

Let's design a Neyman-Pearson decision rule of size α for the problem

$$\mathcal{H}_0 : x \sim \mathcal{N}(0, \sigma^2 I)$$

$$\mathcal{H}_1 : x \sim \mathcal{N}(\mu 1, \sigma^2 I)$$

where $\mu > 0$, $\sigma^2 > 0$ are known, $0 = (0, \dots, 0)^T$, $1 = (1, \dots, 1)^T$ are N -dimensional vectors, and I is the $N \times N$ identity matrix. The likelihood ratio is

$$\begin{aligned} \Lambda(x) &= \frac{\prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_n - \mu)^2}{2\sigma^2}}}{\prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x_n^2}{2\sigma^2}}} \\ &= \frac{e^{-\sum_{n=1}^N \frac{(x_n - \mu)^2}{2\sigma^2}}}{e^{-\sum_{n=1}^N \frac{x_n^2}{2\sigma^2}}} \\ &= \frac{e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n^2 - 2x_n\mu + \mu^2)}}{e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N x_n^2}} \\ &= e^{\frac{1}{\sigma^2} \left(-\frac{N\mu^2}{2} + \mu \sum_{n=1}^N x_n \right)} \end{aligned} \quad (3.17)$$

To simplify the test further we may apply the natural logarithm and rearrange terms to obtain

$$t \equiv \sum_{n=1}^N x_n \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \frac{\sigma^2}{\mu} \ln(\eta) + \frac{N\mu}{2} \equiv \gamma$$

NOTE: We have used the assumption $\mu > 0$. If $\mu < 0$, then division by μ is not a monotonically increasing operation, and the inequalities would be reversed.

The test statistic t is sufficient (Section 1.3.1) for the unknown mean. To set the threshold γ , we write the false-alarm probability (size) as

$$P_F = \Pr[t > \gamma] = \int f_0(t) dt$$

To evaluate P_F , we need to know the density of t under \mathcal{H}_0 . Fortunately, t is the sum of normal variates, so it is again normally distributed. In particular, we have $t = Ax$, where $A = 1^T$, so

$$t \sim \mathcal{N}(A0, A(\sigma^2 I)A^T) = \mathcal{N}(0, N\sigma^2)$$

under \mathcal{H}_0 . Therefore, we may write P_F in terms of the Q-function¹⁸ as

$$P_F = Q\left(\frac{\gamma}{\sqrt{N}\sigma}\right)$$

The threshold is thus determined by

$$\gamma = \sqrt{N}\sigma Q^{-1}(\alpha)$$

¹⁸"The Q-function" <<http://cnx.org/content/m11537/latest/>>

Under \mathcal{H}_1 , we have

$$t \sim \mathcal{N}(A1, A(\sigma^2 I)A^T) = \mathcal{N}(N\mu, N\sigma^2)$$

and so the detection probability (power) is

$$P_D = \Pr[t > \gamma] = Q\left(\frac{\gamma - N\mu}{\sqrt{N}\sigma}\right)$$

Writing P_D as a function of P_F , the ROC curve is given by

$$P_D = Q\left(Q^{-1}(P_F) - \frac{\sqrt{N}\mu}{\sigma}\right)$$

The quantity $\frac{\sqrt{N}\mu}{\sigma}$ is called the **signal-to-noise ratio**. As its name suggests, a larger SNR corresponds to improved performance of the Neyman-Pearson decision rule.

NOTE: In the context of signal processing, the foregoing problem may be viewed as the problem of detecting a constant (DC) signal in additive white Gaussian noise:

$$\mathcal{H}_0 : x_n = w_n, n = 1, \dots, N$$

$$\mathcal{H}_1 : x_n = A + w_n, n = 1, \dots, N$$

where A is a known, fixed amplitude, and $w_n \sim \mathcal{N}(0, \sigma^2)$. Here A corresponds to the mean μ in the example.

3.2.4.1.3 The Neyman-Pearson Lemma: General Case

In our initial statement of the Neyman-Pearson Lemma, we assumed that for all η , the set $\{x, x \mid \Lambda(x) = \eta\}$ had probability zero under \mathcal{H}_0 . This eliminated many important problems from consideration, including tests of discrete data. In this section we remove this restriction.

It is helpful to introduce a more general way of writing decision rules. Let ϕ be a function of the data x with $\phi(x) \in [0, 1]$. ϕ defines the decision rule "declare \mathcal{H}_1 with probability $\phi(x)$." In other words, upon observing x , we flip a " $\phi(x)$ coin." If it turns up heads, we declare \mathcal{H}_1 ; otherwise we declare \mathcal{H}_0 . Thus far, we have only considered rules with $\phi(x) \in \{0, 1\}$

Theorem 3.2: Neyman-Pearson Lemma

Consider the hypothesis testing problem

$$\mathcal{H}_0 : x \sim f_0(x)$$

$$\mathcal{H}_1 : x \sim f_1(x)$$

where f_0 and f_1 are both pdfs or both pmfs. Let $\alpha \in [0, 1)$ be the size (false-alarm probability) constraint. The decision rule

$$\phi(x) = \begin{cases} 1 & \text{if } \Lambda(x) > \eta \\ \rho & \text{if } \Lambda(x) = \eta \\ 0 & \text{if } \Lambda(x) < \eta \end{cases}$$

is the most powerful test of size α , where η and ρ are uniquely determined by requiring $P_F = \alpha$. If $\alpha = 0$, we take $\eta = \infty$, $\rho = 0$. This test is unique up to sets of probability zero under \mathcal{H}_0 and \mathcal{H}_1 .

When $Pr[\Lambda(x) = \eta] > 0$ for certain η , we choose η and ρ as follows: Write

$$P_F = Pr[\Lambda(x) > \eta] + \rho Pr[\Lambda(x) = \eta]$$

Choose η such that

$$Pr[\Lambda(x) > \eta] \leq \alpha \leq Pr[\Lambda(x) \geq \eta]$$

Then choose ρ such that

$$\rho Pr[\Lambda(x) = \eta] = \alpha - Pr[\Lambda(x) < \eta]$$

Example 3.14

3.2.4.1.3.1 Repetition Code

Suppose we have a friend who is trying to transmit a bit (0 or 1) to us over a noisy channel. The channel causes an error in the transmission (that is, the bit is flipped) with probability p , where $0 \leq p < \frac{1}{2}$, and p is known. In order to increase the chance of a successful transmission, our friend sends the same bit N times. Assume the N transmissions are statistically independent. Under these assumptions, the bits you receive are Bernoulli random variables: $x_n \sim \text{Bernoulli}(\theta)$. We are faced with the following hypothesis test:

$$\mathcal{H}_0 : \theta = p \text{ (0 sent)}$$

$$\mathcal{H}_1 : \theta = 1 - p \text{ (1 sent)}$$

We decide to decode the received sequence $x = (x_1, \dots, x_N)^T$ by designing a Neyman-Pearson rule. The likelihood ratio is

$$\begin{aligned} \Lambda(x) &= \frac{\prod_{n=1}^N (1-p)^{x_n} p^{1-x_n}}{\prod_{n=1}^N p^{x_n} (1-p)^{1-x_n}} \\ &= \frac{(1-p)^k p^{N-k}}{p^k (1-p)^{N-k}} \\ &= \left(\frac{1-p}{p}\right)^{2k-N} \end{aligned} \tag{3.18}$$

where $k = \sum_{n=1}^N x_n$ is the number of 1s received.

NOTE: k is a sufficient statistic (Section 1.3.1) for θ .

The LRT is

$$\left(\frac{1-p}{p}\right)^{2k-N} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta$$

Taking the natural logarithm of both sides and rearranging, we have

$$k \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \frac{N}{2} + \frac{1}{2} \frac{\ln(\eta)}{\ln\left(\frac{1-p}{p}\right)} = \gamma$$

The false alarm probability is

$$\begin{aligned} P_F &= Pr[k > \gamma] + \rho Pr[k = \gamma] \\ &= \sum_{k=\gamma+1}^N \binom{N}{k} p^k (1-p)^{N-k} + \rho \binom{N}{\gamma} p^\gamma (1-p)^{N-\gamma} \end{aligned} \tag{3.19}$$

γ and ρ are chosen so that $P_F = \alpha$, as described above.

The corresponding detection probability is

$$\begin{aligned} P_D &= Pr[k > \gamma] + \rho Pr[k = \gamma] \\ &= \sum_{k=\gamma+1}^N \binom{N}{k} (1-p)^k p^{N-k} + \rho \binom{N}{\gamma} (1-p)^\gamma p^{N-\gamma} \end{aligned} \quad (3.20)$$

3.2.4.1.4 Problems

Exercise 3.8

Design a hypothesis testing problem involving continuous random variables such that $Pr[\Lambda(x) = \eta] > 0$ for certain values of η . Write down the false-alarm probability as a function of the threshold. Make as general a statement as possible about when the technical condition (p. 70) is satisfied.

Exercise 3.9

Consider the scalar hypothesis testing problem

$$\mathcal{H}_0 : x \sim f_0(x)$$

$$\mathcal{H}_1 : x \sim f_1(x)$$

where

$$f_i(x) = \frac{1}{\pi(1+(x-i)^2)}, i = \{0, 1\}$$

3.2.4.1.4.1

Write down the likelihood ratio test.

3.2.4.1.4.2

Determine the decision regions as a function of η_1 for all $\eta > 0$. Draw a representative of each. What are the "critical" values of η ?

NOTE: There are five distinct cases.

3.2.4.1.4.3

Compute the size and power (P_F and P_D) in terms of the threshold η_1 and plot the ROC.

NOTE:

$$\int \frac{1}{1+x^2} dx = \arctan(x)$$

3.2.4.1.4.4

Suppose we decide to use a simple threshold test $x \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \eta$ instead of the Neyman-Pearson rule. Does our performance \mathcal{H}_0 suffer much? Plot the ROC for this decision rule on the same graph as for the previous (Section 3.2.4.1.4.3) ROC.

Exercise 3.10

Suppose we observe N independent realizations of a Poisson random variable k with intensity parameter λ :

$$f(k) = \frac{e^{-\lambda} \lambda^k}{k!}$$

We must decide which of two intensities is in effect:

$$\mathcal{H}_0 : \lambda = \lambda_0$$

$$\mathcal{H}_1 : \lambda = \lambda_1$$

where $\lambda_0 < \lambda_1$.

3.2.4.1.4.1

Write down the likelihood ratio test.

3.2.4.1.4.2

Simplify the LRT to a test statistic involving only a sufficient statistic. Apply a monotonically increasing transformation to simplify further.

3.2.4.1.4.3

Determine the distribution of the sufficient statistic under both hypotheses.

NOTE: Use the characteristic function to show that a sum of IID Poisson variates is again Poisson distributed.

3.2.4.1.4.4

Derive an expression for the probability of error.

3.2.4.1.4.5

Assuming the two hypotheses are equally likely, and $\lambda_0 = 5$ and $\lambda_1 = 6$, what is the minimum number N of observations needed to attain a false-alarm probability no greater than 0.01?

NOTE: If you have numerical trouble, try rewriting the log-factorial so as to avoid evaluating the factorial of large integers.

Exercise 3.11

In Example 3.13, suppose $p = 0.1$. What is the smallest value of N needed to ensure $P_F \leq 0.01$? What is P_D in this case?

3.3 Detection with Unknowns

Chapter 4

Estimation Theory

4.1 Minimum Variance Unbiased Estimation

4.1.1 The Minimum Variance Unbiased Estimator¹

4.1.1.1 In Search of a Useful Criterion

In parameter estimation, we observe an N -dimensional vector X of measurements. The distribution of X is governed by a density or probability mass function $f_{\theta}(x)$, which is parameterized by an unknown parameter θ . We would like to establish a useful criterion for guiding the design and assessing the quality of an estimator $\hat{\theta}(x)$. We will adopt a classical (frequentist) view of the unknown parameter: it is not itself random, it is simply unknown.

One possibility is to try to design an estimator that minimizes the **mean-squared error**, that is, the expected squared deviation of the estimated parameter value from the true parameter value. For a scalar parameter, the MSE is defined by

$$\text{MSE}(\hat{\theta}, \theta) = E \left[\left(\hat{\theta}(x) - \theta \right)^2 \right] \quad (4.1)$$

For a vector parameter θ , this definition is generalized by

$$\text{MSE}(\hat{\theta}, \theta) = E \left[\left(\hat{\theta}(x) - \theta \right)^T \left(\hat{\theta}(x) - \theta \right) \right] \quad (4.2)$$

The expectation is with respect to the distribution of X . Note that for a given estimator, the MSE is a function of θ .

While the MSE is a perfectly reasonable way to assess the quality of an estimator, it does not lead to a useful design criterion. Indeed, the estimator that minimizes the MSE is simply the estimator

$$\hat{\theta}(x) = \theta \quad (4.3)$$

Unfortunately, this depends on the value of the unknown parameter, and is therefore not realizable! We need a criterion that leads to a realizable estimator.

NOTE: In the Bayesian Approach to Parameter Estimation, the MSE is a useful design rule.

¹This content is available online at <http://cnx.org/content/m11426/1.6/>.

4.1.1.2 The Bias-Variance Decomposition of the MSE

It is possible to rewrite the MSE in such a way that a useful optimality criterion for estimation emerges. For a scalar parameter θ , [Insert 1] This expression is called the **bias-variance decomposition** of the mean-squared error. The first term on the right-hand side is called the **variance** of the estimator, and the second term on the right-hand side is the square of the **bias** of the estimator. The formal definition of these concepts for vector parameters is now given:

Let $\hat{\theta}$ be an estimator of the parameter θ .

Definition 4.1: variance

The **variance** of $\hat{\theta}$ is [Insert 2]

Definition 4.2: bias

The **bias** of $\hat{\theta}$ is [Insert 3]

The bias-variance decomposition also holds for vector parameters: [Insert 4] The proof is a straightforward generalization of the argument for the scalar parameter case.

Exercise 4.1

Prove the bias-variance decomposition of the MSE for the vector parameter case.

4.1.1.3 The Bias-Variance Tradeoff

The MSE decomposes into the sum of two non-negative terms, the squared bias and the variance. In general, for an arbitrary estimator, both of these terms will be nonzero. Furthermore, as an estimator is modified so that one term increases, typically the other term will decrease. This is the so-called **bias-variance tradeoff**. The following example illustrates this effect.

Example 4.1

Let $\tilde{A} = \alpha \frac{1}{N} \sum_{n=1}^N x_n$, where $x_n = A + w_n$, $w_n \sim \mathcal{N}(0, \sigma^2)$, and α is an arbitrary constant.

Let's find the value of α that minimizes the MSE.

$$\text{MSE}(\tilde{A}) = E \left[(\tilde{A} - A)^2 \right] \quad (4.4)$$

NOTE: $\tilde{A} = \alpha S_N$, $S_N \sim \mathcal{N}\left(A, \frac{\sigma^2}{N}\right)$

$$\begin{aligned} \text{MSE}(\tilde{A}) &= E \left[\tilde{A}^2 \right] - 2E \left[\tilde{A} \right] A + A^2 \\ &= \alpha^2 E \left[\frac{1}{N^2} \sum_{i,j=1}^N x_i x_j \right] - 2\alpha E \left[\frac{1}{N} \sum_{n=1}^N x_n \right] A + A^2 \\ &= \alpha^2 \frac{1}{N^2} \sum_{i,j=1}^N E \left[x_i x_j \right] - 2\alpha \times \frac{1}{N} \sum_{n=1}^N E \left[x_n \right] + A^2 \end{aligned} \quad (4.5)$$

$$E \left[x_i x_j \right] = \begin{cases} A^2 + \sigma^2 & \text{if } i = j \\ A^2 & \text{if } i \neq j \end{cases}$$

$$\begin{aligned} \text{MSE}(\tilde{A}) &= \alpha^2 \left(A^2 + \frac{\sigma^2}{N} \right) - 2\alpha A^2 + A^2 \\ &= \frac{\alpha^2 \sigma^2}{N} + (\alpha - 1)^2 A^2 \end{aligned} \quad (4.6)$$

$$\sigma(\tilde{A})^2 = \frac{\alpha^2 \sigma^2}{N}$$

$$\text{Bias}^2(\tilde{A}) = (\alpha - 1)^2 A^2$$

$$\frac{\partial \text{MSE}(\tilde{A})}{\partial \alpha} = \frac{2\alpha\sigma^2}{N} + 2(\alpha - 1)A^2 = 0$$

$$\alpha^* = \frac{A^2}{A^2 + \frac{\sigma^2}{N}} \quad (4.7)$$

The optimal value α^* depends on the unknown parameter A ! Therefore the estimator is not realizable.

Note that the problematic dependence on the parameter enters through the Bias component of the MSE. Therefore, a reasonable alternative is to constrain the estimator to be unbiased, and then find the estimator that produces the minimum variance (and hence provides the minimum MSE among all unbiased estimators).

NOTE: Sometimes no unbiased estimator exists, and we cannot proceed at all in this direction.

In this example, note that as the value of α varies, one of the squared bias or variance terms increases, while the other one decreases. Furthermore, note that **the dependence of the MSE on the unknown parameter is manifested in the bias.**

4.1.1.4 Unbiased Estimators

Since the bias depends on the value of the unknown parameter, it seems that any estimation criterion that depends on the bias would lead to an unrealizable estimator, as the previous example (Example 4.1) suggests (although in certain cases realizable minimum MSE estimators can be found). As an alternative to minimizing the MSE, we could focus on estimators that have a bias of zero. In this case, the bias contributes zero to the MSE, and in particular, it does not involve the unknown parameter. By focusing on estimators with zero bias, we may hope to arrive at a design criterion that yields **realizable** estimators.

Definition 4.3: unbiased

An estimator θ is called **unbiased** if its bias is zero for all values of the unknown parameter. Equivalently, [Insert 5]

For an estimator to be unbiased we require that **on average** the estimator will yield the true value of the unknown parameter. We now give some examples.

The sample mean of a random sample is always an unbiased estimator for the mean.

Example 4.2

Estimate the DC level in the **Gaussian white noise**.

Suppose we have data x_1, \dots, x_N and model the data by

$$x_n = A + w_n \quad , \quad n \in \{1, \dots, N\}$$

where A is the unknown DC level, and $w_n \sim \mathcal{N}(\sigma, \sigma^2)$.

The parameter is $-\infty < A < \infty$.

Consider the **sample-mean estimator**:

$$A = \frac{1}{N} \sum_{n=1}^N x_n$$

Is A unbiased? **Yes.**

Since $E[\cdot]$ is a linear operator,

$$E[A] = \frac{1}{N} \sum_{n=1}^N E[x_n] = \frac{1}{N} \sum_{n=1}^N A = A$$

Therefore, A is unbiased!

What does the unbiased restriction really imply? Recall that $\theta = g(x)$, a function of the data. Therefore,

$$E[\theta] = \theta$$

and

$$E[\theta] = \int g(x)p(x|\theta)dx = \theta$$

Hence, to be unbiased, the estimator ($g(\cdot)$) must satisfy an integral equation involving the densities $p(x|\theta)$.

It is possible that an estimator can be unbiased for some parameter values, but be biased for others.

The bias of an estimator may be zero for **some** values of the unknown parameter, but not others. In this case, the estimator is **not** an unbiased estimator.

Example 4.3

$$\tilde{A} = \frac{1}{2N} \sum_{n=1}^N x_n$$

$$E[\tilde{A}] = \frac{1}{2}A = \begin{cases} 0 & \text{if } (A=0) \Rightarrow \text{unbiased} \\ \frac{1}{2}A & \text{if } (A \neq 0) \Rightarrow \text{biased} \end{cases}$$

An unbiased estimator is not necessarily a good estimator.

Some unbiased estimators are more useful than others.

Example 4.4

$$x_n = A + w_n, \quad w_n \sim \mathcal{N}(\sigma, \sigma^2)$$

$$A_1 = x_1$$

$$E[A_1] = A$$

$$A_2 = \frac{1}{N} \sum_{n=1}^N x_n$$

$$E[A_2] = A$$

$$\sigma(A_1)^2 = \sigma^2$$

$$\sigma(A_2)^2 = \frac{\sigma^2}{N}$$

Both estimators are unbiased, but A_2 has a much lower variance and therefore is a better estimator.

NOTE: $A_1(N)$ is inconsistent. $A_2(N)$ is consistent.

4.1.1.5 Minimum Variance Unbiased Estimators

Direct minimization of the MSE generally leads to non-realizable estimators. Since the dependence of an estimator on the unknown parameter appears to come from the bias term, we hope that constraining the bias to be zero will lead to a useful design criterion. But if the bias is zero, then the mean-squared error is just the variance. This gives rise to the **minimum variance unbiased estimator (MVUE)** for θ .

Definition 4.4: MVUE

An estimator θ is the **minimum variance unbiased estimator** if it is unbiased and has the smallest variance of any unbiased estimator for all values of the unknown parameter. In other words, the MVUE satisfies the following two properties: [Insert 6]

The minimum variance unbiased criterion is the primary estimation criterion in the classical (non-Bayesian) approach to parameter estimation. Before delving into ways of finding the MVUE, let's first consider whether the MVUE always exists.

4.1.1.6 Existence of the MVUE

The MVUE does not always exist. In fact, it may be that no unbiased estimators exist, as the following example demonstrates.

Place [Insert 7] here and make it an example (5).

Even if unbiased estimators exist, it may be that no single unbiased estimator has the minimum variance for **all** values of the unknown parameter.

Place [Insert 8] here and make it an example (6).

Exercise 4.2

Compute the variances of the estimators in the previous examples. Using the Cramer-Rao Lower bound, show that one of these two estimators has minimum variance among all unbiased estimators. Deduce that no single realizable estimator can have minimum variance among all unbiased estimators for all parameter values (*i.e.*, the MVUE does not exist). When using the Cramer-Rao bound, note that the likelihood is not differentiable at $\theta = 0$.

4.1.1.7 Methods for Finding the MVUE

Despite the fact that the MVUE doesn't always exist, in many cases of interest it does exist, and we need methods for finding it. Unfortunately, there is no 'turn the crank' algorithm for finding MVUE's. There are, instead, a variety of techniques that can **sometimes** be applied to find the MVUE. These methods include:

1. Compute the Cramer-Rao Lower Bound, and check the condition for equality.
2. Find a **complete** sufficient statistic and apply the Rao-Blackwell Theorem.
3. If the data obeys a general linear model, restrict to the class of linear unbiased estimators, and find the minimum variance estimator within that class. This method is in general suboptimal, although when the noise is Gaussian, it produces the MVUE.

4.1.2 The Cramer-Rao Lower Bound²

The **Cramer-Rao Lower Bound (CRLB)** sets a lower bound on the variance of **any** unbiased estimator. This can be extremely useful in several ways:

1. If we find an estimator that achieves the CRLB, then we know that we have found an MVUB estimator!

²This content is available online at <<http://cnx.org/content/m11429/1.4/>>.

2. The CRLB can provide a benchmark against which we can compare the performance of any unbiased estimator. (We know we're doing very well if our estimator is "close" to the CRLB.)
3. The CRLB enables us to rule-out impossible estimators. That is, we know that it is physically impossible to find an unbiased estimator that beats the CRLB. This is useful in feasibility studies.
4. The theory behind the CRLB can tell us if an estimator exists that achieves the bound.

4.1.2.1 Estimator Accuracy

Consider the likelihood function $p(x|\theta)$, where θ is a scalar unknown (parameter). We can plot the likelihood as a function of the unknown, as shown in Figure 4.1.

Image not finished

Figure 4.1

The more "peaky" or "spiky" the likelihood function, the easier it is to determine the unknown parameter.

Example 4.5

Suppose we observe

$$x = A + w$$

where $w \sim \mathcal{N}(\sigma, \sigma^2)$ and A is an unknown parameter. The "smaller" the noise w is, the easier it will be to estimate A from the observation x .

Suppose $A = 3$ and $\sigma = 1/3$.

Image not finished

Figure 4.2

Given this density function, we can easily rule-out estimates of A greater than 4 or less than 2, since it is very unlikely that such A could give rise to our observation.

On the other hand, suppose $\sigma = 1$.

Image not finished

Figure 4.3

In this case, it is very difficult to estimate A . Since the noise power is larger, it is very difficult to distinguish A from the noise.

The key thing to notice is that the estimation accuracy of A depends on σ , which in effect determines the peakiness of the likelihood. The more peaky, the better localized the data is about the true parameter.

To quantify the notion, note that the peakiness is effectively measured by the negative of the second derivative of the log-likelihood at its peak, as seen in Figure 4.4.

Image not finished

Figure 4.4

Example 4.6

$$x = A + w$$

$$\log p(x | A) = \left(-\log \sqrt{2\pi\sigma^2} \right) - \frac{1}{2\sigma^2} (x - A)^2 \quad (4.8)$$

$$\begin{aligned} \frac{\partial \log p(x | A)}{\partial A} &= \frac{1}{\sigma^2} (x - A) \\ -\frac{\partial^2 \log p(x | A)}{\partial A^2} &= \frac{1}{\sigma^2} \end{aligned} \quad (4.9)$$

The curvature increases as σ^2 decreases (curvature=peakiness).

In general, the curvature will depend on the observation data; $-\frac{\partial^2 \log p(x | A)}{\partial \theta^2}$ is a function of x . Therefore, an average measure of curvature is more appropriate.

$$-E \left[\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} \right] \quad (4.10)$$

This average-out randomness due to the data and is a function of θ alone.

We are now ready to state the CRLB theorem.

Theorem 4.1: Cramer-Rao Lower Bound Theorem

Assume that the pdf $p(x | \theta)$ satisfies the "regularity" condition

$$E \left[\frac{\partial \log p(x | \theta)}{\partial \theta} \right] = 0$$

where the expectation is taken with respect to $p(x | \theta)$. Then, the variance of any unbiased estimator θ must satisfy

$$\sigma(\theta)^2 \geq \frac{1}{-E \left[\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} \right]} \quad (4.11)$$

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

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = I(\theta) (g(\theta) - \theta) \quad (4.12)$$

for some functions g and I .

The corresponding estimator is MVUB and is given by $\theta = g(x)$, and the minimum variance is $\frac{1}{I(\theta)}$.

Example

$$x = A + w$$

where

$$w \sim \mathcal{N}(0, \sigma^2)$$

$$\theta = A$$

$$E \left[\frac{\partial \log p}{\partial \theta} \right] = E \left[\frac{1}{\sigma^2} (x - A) \right] = 0$$

$$\text{CRLB} = \frac{1}{-E \left[\frac{\partial^2 \log p}{\partial \theta^2} \right]} = \frac{1}{\sigma^2} = \sigma^2$$

Therefore, any unbiased estimator A has $\sigma(A)^2 \geq \sigma^2$. But we know that $A = x$ has $\sigma(A)^2 = \sigma^2$.

Therefore, $A = x$ is the MVUB estimator.

NOTE:

$$\theta = A$$

$$I(\theta) = \frac{1}{\sigma^2}$$

$$g(x) = x$$

Proof:

First consider the reguarity condition:

$$E \left[\frac{\partial \log p(x | \theta)}{\partial \theta} \right] = 0$$

NOTE:

$$E \left[\frac{\partial \log p(x | \theta)}{\partial \theta} \right] = \int \frac{\partial \log p(x | \theta)}{\partial \theta} p(x | \theta) d\theta = \int \frac{\partial p(x | \theta)}{\partial \theta} d\theta$$

Now assuming that we can interchange order of differentiation and integration

$$E \left[\frac{\partial \log p(x | \theta)}{\partial \theta} \right] = \frac{\partial \int p(x | \theta) d\theta}{\partial \theta} = \frac{\partial 1}{\partial \theta} = 0$$

So the regularity condition is satisfied whenever this interchange is possible³; i.e., when derivative is well-defined, fails for uniform density.

Now lets derive the CRLB for a scalar parameter $\alpha = g(\theta)$, where the pdf is $p(x | \theta)$. Consider any unbiased estimator of α :

$$\alpha \in \left(E[\alpha] = \alpha = g(\theta) \right)$$

Note that this is equivalent to

$$\int \alpha p(x | \theta) dx = g(\theta)$$

³This is simply the **Fundamental Theorem of Calculus** applied to $p(x | \theta)$. So long as $p(x | \theta)$ is absolutely continuous with respect to the **Lebesgue measure**, this is possible.

where α is unbiased. Now differentiate both side

$$\int \alpha \frac{\partial p(x|\theta)}{\partial \theta} dx = \frac{\partial g(\theta)}{\partial \theta}$$

or

$$\int \alpha \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx = \frac{\partial g(\theta)}{\partial \theta}$$

Now, exploiting the regularity condition,

$$\int (\alpha - \alpha) \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx = \frac{\partial g(\theta)}{\partial \theta} \quad (4.13)$$

since

$$\int \alpha \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx = \alpha E[\log p(x|\theta)] = 0$$

Now apply the **Cauchy-Schwarz inequality** to the integral above (4.13):

$$\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2 = \left(\int (\alpha - \alpha) \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx \right)^2$$

$$\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2 \leq \int (\alpha - \alpha)^2 p(x|\theta) dx \int \left(\frac{\partial \log p(x|\theta)}{\partial \theta} \right)^2 p(x|\theta) dx$$

$\sigma(\alpha)^2$ is $\int (\alpha - \alpha)^2 p(x|\theta) dx$, so

$$\sigma(\alpha)^2 \geq \frac{\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2}{E \left[\left(\frac{\partial \log p(x|\theta)}{\partial \theta} \right)^2 \right]} \quad (4.14)$$

Now we note that

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

Why? Regularity condition.

$$E \left[\frac{\partial \log p(x|\theta)}{\partial \theta} \right] = \int \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx = 0$$

Thus,

$$\frac{\partial \int \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx}{\partial \theta} = 0$$

or

$$\int \frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} p(x|\theta) + \frac{\partial \log p(x|\theta)}{\partial \theta} \frac{\partial p(x|\theta)}{\partial \theta} dx = 0$$

Therefore,

$$-E \left[\frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} \right] = \int \frac{\partial \log p(x|\theta)}{\partial \theta} \frac{\partial \log p(x|\theta)}{\partial \theta} p(x|\theta) dx = E \left[\left(\frac{\partial \log p(x|\theta)}{\partial \theta} \right)^2 \right]$$

Thus, (4.14) becomes

$$\sigma(\alpha)^2 \geq \frac{\left(\frac{\partial g(\theta)}{\partial \theta} \right)^2}{-E \left[\frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} \right]}$$

NOTE: If $g(\theta) = \theta$, then numerator is 1.

Example: DC Level in White Gaussian Noise

$$x_n = A + w_n \quad , \quad n \in \{1, \dots, N\}$$

where

$$w_n \sim \mathcal{N}(0, \sigma^2)$$

$$p(x | A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2}$$

$$\frac{\partial \log p(x | A)}{\partial A} = \frac{\partial \left(\left(-\log(2\pi\sigma^2)^{\frac{N}{2}} \right) - \frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2 \right)}{\partial A} = \frac{1}{\sigma^2} \sum_{n=1}^N x_n - A$$

$$E \left[\frac{\partial \log p(x | A)}{\partial A} \right] = 0$$

$$\frac{\partial^2 \log p(x | A)}{\partial A^2} = -\frac{N}{\sigma^2}$$

Therefore, the variance of any unbiased estimator satisfies:

$$\sigma(A)^2 \geq \frac{\sigma^2}{N}$$

The sample-mean estimator $A = \frac{1}{N} \sum_{n=1}^N x_n$ attains this bound and therefore is MVUB.

Corollary 4.1:

When the CRLB is attained

$$\sigma(\theta)^2 = \frac{1}{I(\theta)}$$

where

$$I(\theta) = -E \left[\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} \right]$$

The quantity $I(\theta)$ is called **Fisher Information** that x contains about θ .

Proof:

By CRLB Theorem,

$$\sigma(\theta)^2 = \frac{1}{-E \left[\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} \right]}$$

and

$$\frac{\partial \log p(x | \theta)}{\partial \theta} = I(\theta) (\theta - \theta)$$

This yields

$$\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} = \frac{\partial I(\theta)}{\partial \theta} (\theta - \theta) - I(\theta)$$

which in turn yields

$$-E \left[\frac{\partial^2 \log p(x | \theta)}{\partial \theta^2} \right] = I(\theta)$$

So,

$$\sigma(\theta)^2 = \frac{1}{I(\theta)}$$

The CRLB is not always attained.

Example 4.7: Phase Estimation

$$x_n = A \cos(2\pi f_0 n + \phi) + w_n, \quad n \in \{1, \dots, N\}$$

The amplitude and frequency are assumed known

$$w_n \sim \mathcal{N}(0, \sigma^2)$$

idd.

$$p(x|\phi) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N x_n - A \cos(2\pi f_0 n + \phi)}$$

$$\frac{\partial \log p(x|\phi)}{\partial \phi} = \left(-\frac{A}{\sigma^2}\right) \sum_{n=1}^N x_n \sin(2\pi f_0 n + \phi) - \frac{A}{2} \sin(4\pi f_0 n + \phi)$$

$$\frac{\partial^2 \log p(x|\phi)}{\partial \phi^2} = \left(-\frac{A}{\sigma^2}\right) \sum_{n=1}^N x_n \cos(2\pi f_0 n + \phi) - A \cos(2\pi f_0 n + 2\phi)$$

$$-E \left[\frac{\partial^2 \log p(x|\phi)}{\partial \phi^2} \right] = \frac{A^2}{\sigma^2} \sum_{n=1}^N 1/2 + 1/2 \cos(4\pi f_0 n + 2\phi) - \cos(4\pi f_0 n + 2\phi)$$

Since $I(\phi) = -E \left[\frac{\partial^2 \log p(x|\phi)}{\partial \phi^2} \right]$,

$$I(\phi) \simeq \frac{NA^2}{2\sigma^2}$$

because $\frac{1}{N} \sum \cos(4\pi f_0 n) \simeq 0$, $0 < f_0 < k$ Therefore,

$$\sigma(\phi)^2 \geq \frac{2\sigma^2}{NA^2}$$

In this case, it can be shown that there does not exist a g such that

$$\frac{\partial \log p(x|\phi)}{\partial \phi} \neq I(\phi)(g(x) - \phi)$$

Therefore, an unbiased phase estimator that attains the CRLB does not exist.

However, a MVUB estimator may still exist—only its variance will be larger than the CRLB.

4.1.2.2 Efficiency

An estimator which is unbiased and attains the CRLB is said to be **efficient**.

Example 4.8

Sample-mean estimator is efficient.

Example 4.9

Supposed three unbiased estimators exist for a param θ .

Image not finished

Figure 4.5

Image not finished

Figure 4.6

Example 4.10: Sinusoidal Frequency Estimation

$$s_n(f_0) = A \cos(2\pi f_0 n + \phi) \quad , \quad 0 < f_0 < 1/2$$

$$x_n = s_n(f_0) + w_n \quad , \quad n \in \{1, \dots, N\}$$

A and ϕ are known, while f_0 is unknown.

$$\sigma^2 \left(f_0 \right)^2 \geq \frac{\sigma^2}{A^2 \sum_{n=1}^N (2\pi n \sin(2\pi f_0 n + \phi))^2}$$

Suppose $\frac{A^2}{\sigma^2} = 1$ (SNR), where $N = 10$ and $\phi = 0$.

Image not finished

Figure 4.7: Some frequencies are easier to estimator (lower CRLB, but not necessarily just lower bound) than others.

4.1.2.3 CRLB for Vector Parameter

$$\theta = \begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_p \end{pmatrix}$$

θ is unbiased, i.e.,

$$E[\theta_i] = \theta_i \quad , \quad i \in \{1, \dots, p\}$$

4.1.2.4 CRLB

$$\sigma(\theta_i)^2 \geq (I(\theta))^{-1}_{i,i}$$

where

$$I(\theta)_{i,j} = -E \left[\frac{\partial^2 \log p(x|\theta)}{\partial \theta_i \partial \theta_j} \right]$$

$I(\theta)$ is the **Fisher Information Matrix**.

Theorem 4.2: Cramer-Rao Lower Bound - Vector Parameter
Assume the pdf $p(x|\theta)$ satisfies the "regularity" condition

$$E \left[\frac{\partial \log p(x|\theta)}{\partial \theta} \right] = 0$$

Then the covariance matrix of any unbiased estimator θ satisfies

$$C_{\hat{\theta}} - (I(\theta))^{-1} \geq 0$$

(meaning $C_{\hat{\theta}} - (I(\theta))^{-1}$ is p.s.d.) The Fisher Information matrix is

$$I(\theta)_{i,j} = -E \left[\frac{\partial^2 \log p(x|\theta)}{\partial \theta^2} \right]$$

Furthermore, θ attains the CRLB ($C_{\hat{\theta}} = (I(\theta))^{-1}$) iff

$$\frac{\partial \log p(x|\theta)}{\partial \theta} = I(\theta) (g(x) - \theta)$$

and

$$\theta = g(x)$$

Example: DC Level in White Gaussian Noise

$$x_n = A + w_n, \quad n \in \{1, \dots, N\}$$

A is unknown and $w_n \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is unknown.

$$\theta = \begin{pmatrix} A \\ \sigma^2 \end{pmatrix}$$

$$\log p(x|\theta) = \left(- \left(\frac{N}{2} \log(2\pi) \right) \right) - \frac{N}{2} \log \sigma^2 - \frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2$$

$$\frac{\partial \log p(x|\theta)}{\partial A} = \frac{1}{\sigma^2} \sum_{n=1}^N x_n - A$$

$$\frac{\partial \log p(x|\theta)}{\partial \sigma^2} = -\frac{N}{\sigma^2} + \frac{1}{2\sigma^4} \sum_{n=1}^N (x_n - A)^2$$

$$\left(\frac{\partial^2 \log p(x|\theta)}{\partial A^2} = -\frac{N}{\sigma^2} \right) \rightarrow -\frac{N}{\sigma^2}$$

$$\left(\frac{\partial^2 \log p(x|\theta)}{\partial A \partial \sigma^2} = - \left(\frac{1}{\sigma^4} \sum_{n=1}^N x_n - A \right) \right) \rightarrow 0$$

$$\left(\frac{\partial^2 \log p(x|\theta)}{\partial \sigma^2^2} = \frac{N}{2\sigma^4} - \frac{1}{\sigma^6} \sum_{n=1}^N (x_n - A)^2 \right) \rightarrow -\frac{N}{2\sigma^4}$$

Which leads to

$$I(\theta) = \begin{pmatrix} \frac{N}{\sigma^2} & 0 \\ 0 & \frac{N}{2\sigma^4} \end{pmatrix}$$

$$\sigma(A)^2 \geq \frac{\sigma^2}{N}$$

$$\sigma(\sigma^2)^2 \geq \frac{2\sigma^4}{N}$$

Note that the CRLB for A is the same whether or not σ^2 is known. This happens in this case due to the diagonal nature of the Fisher Information Matrix.

In general the Fisher Information Matrix is not diagonal and consequently the CRLBs will depend on other unknown parameters.

4.2 Maximum Likelihood Estimation⁴

The **maximum likelihood estimator** (MLE) is an alternative to the minimum variance unbiased estimator (MVUE). For many estimation problems, the MVUE does not exist. Moreover, when it does exist, there is no systematic procedure for finding it. In contrast, the MLE does not necessarily satisfy any optimality criterion, but it can almost always be computed, either through exact formulas or numerical techniques. For this reason, the MLE is one of the most common estimation procedures used in practice.

The MLE is an important type of estimator for the following reasons:

1. The MLE implements the likelihood principle.
2. MLEs are often simple and easy to compute.
3. MLEs have asymptotic optimality properties (consistency and efficiency).
4. MLEs are invariant under reparameterization.
5. If an efficient estimator exists, it is the MLE.
6. In signal detection with unknown parameters (composite hypothesis testing), MLEs are used in implementing the generalized likelihood ratio test (GLRT).

This module will discuss these properties in detail, with examples.

4.2.1 The Likelihood Principle

Supposed the data X is distributed according to the density or mass function $p(x|\theta)$. The **likelihood function** for θ is defined by

$$l(\theta|x) \equiv p(x|\theta)$$

At first glance, the likelihood function is nothing new - it is simply a way of rewriting the pdf/pmf of X . The difference between the likelihood and the pdf or pmf is what is held fixed and what is allowed to vary. When we talk about the likelihood, we view the observation x as being fixed, and the parameter θ as freely varying.

⁴This content is available online at <<http://cnx.org/content/m11446/1.5/>>.

NOTE: It is tempting to view the likelihood function as a probability density for θ , and to think of $l(\theta | x)$ as the conditional density of θ given x . This approach to parameter estimation is called **fiducial inference**, and is not accepted by most statisticians. One potential problem, for example, is that in many cases $l(\theta | x)$ is not integrable ($\int l(\theta | x) d\theta \rightarrow \infty$) and thus cannot be normalized. A more fundamental problem is that θ is viewed as a fixed quantity, as opposed to random. Thus, it doesn't make sense to talk about its density. For the likelihood to be properly thought of as a density, a Bayesian approach is required.

The likelihood principle effectively states that all information we have about the unknown parameter θ is contained in the likelihood function.

Rule 4.1: Likelihood Principle

The information brought by an observation x about θ is entirely contained in the likelihood function $p(x | \theta)$. Moreover, if x_1 and x_2 are two observations depending on the same parameter θ , such that there exists a constant c satisfying $p(x_1 | \theta) = cp(x_2 | \theta)$ for every θ , then they bring the same information about θ and must lead to identical estimators.

In the statement of the likelihood principle, it is **not** assumed that the two observations x_1 and x_2 are generated according to the same model, as long as the model is parameterized by θ .

Example 4.11

Suppose a public health official conducts a survey to estimate $0 \leq \theta \leq 1$, the percentage of the population eating pizza at least once per week. As a result, the official found nine people who had eaten pizza in the last week, and three who had not. If no additional information is available regarding how the survey was implemented, then there are at least two probability models we can adopt.

1. The official surveyed 12 people, and 9 of them had eaten pizza in the last week. In this case, we observe $x_1 = 9$, where

$$x_1 \sim \text{Binomial}(12, \theta)$$

The density for x_1 is

$$f(x_1 | \theta) = \binom{12}{x_1} \theta^{x_1} (1 - \theta)^{12 - x_1}$$

2. Another reasonable model is to assume that the official surveyed people **until** he found 3 non-pizza eaters. In this case, we observe $x_2 = 12$, where

$$x_2 \sim \text{NegativeBinomial}(3, 1 - \theta)$$

The density for x_2 is

$$g(x_2 | \theta) = \binom{x_2 - 1}{3 - 1} \theta^{x_2 - 3} (1 - \theta)^3$$

The likelihoods for these two models are proportional:

$$\ell(\theta | x_1) \propto \ell(\theta | x_2) \propto \theta^9 (1 - \theta)^3$$

Therefore, any estimator that adheres to the likelihood principle will produce the same estimate for θ , regardless of which of the two data-generation models is assumed.

The likelihood principle is widely accepted among statisticians. In the context of parameter estimation, any reasonable estimator should conform to the likelihood principle. As we will see, the maximum likelihood estimator does.

NOTE: While the likelihood principle itself is a fairly reasonable assumption, it can also be derived from two somewhat more intuitive assumptions known as the **sufficiency principle** and the **conditionality principle**. See *Casella and Berger, Chapter 6*[1].

4.2.2 The Maximum Likelihood Estimator

The **maximum likelihood estimator** $\theta(x)$ is defined by

$$\theta = \underset{\theta}{\operatorname{argmax}} l(\theta | x)$$

Intuitively, we are choosing θ to maximize the probability of occurrence of the observation x .

NOTE: It is possible that multiple parameter values maximize the likelihood for a given x . In that case, any of these maximizers can be selected as the MLE. It is also possible that the likelihood may be **unbounded**, in which case the MLE does not exist.

The MLE rule is an implementation of the likelihood principle. If we have two observations whose likelihoods are proportional (they differ by a constant that does not depend on θ), then the value of θ that maximizes one likelihood will also maximize the other. In other words, both likelihood functions lead to the same inference about θ , as required by the likelihood principle.

Understand that maximum likelihood is a **procedure**, not an optimality criterion. From the definition of the MLE, we have no idea how close it comes to the true parameter value relative to other estimators. In contrast, the MVUE is defined as the estimator that satisfies a certain optimality criterion. However, unlike the MLE, we have no clear procedure to follow to compute the MVUE.

4.2.3 Computing the MLE

If the likelihood function is differentiable, then θ is found by differentiating the likelihood (or log-likelihood), equating with zero, and solving:

$$\frac{\partial \log l(\theta | x)}{\partial \theta} = 0$$

If multiple solutions exist, then the MLE is the solution that maximizes $\log l(\theta | x)$, that is, the **global** maximizer.

In certain cases, such as pdfs or pmfs with an exponential form, the MLE can be easily solved for. That is,

$$\frac{\partial \log l(\theta | x)}{\partial \theta} = 0$$

can be solved using calculus and standard linear algebra.

Example 4.12: DC level in white Gaussian noise

Suppose we observe an unknown amplitude in white Gaussian noise with unknown variance:

$$x_n = A + w_n$$

$n \in \{0, 1, \dots, N-1\}$, where $w_n \sim \mathcal{N}(0, \sigma^2)$ are independent and identically distributed. We would like to estimate

$$\theta = \begin{pmatrix} A \\ \sigma^2 \end{pmatrix}$$

by computing the MLE. Differentiating the log-likelihood gives

$$\frac{\partial \log p(x | \theta)}{\partial A} = \frac{1}{\sigma^2} \sum_{n=1}^N x_n - A$$

$$\frac{\partial \log p(x | \theta)}{\partial \sigma^2} = -\frac{N}{\sigma^2} + \frac{1}{2\sigma^4} \sum_{n=1}^N (x_n - A)^2$$

Equating with zero and solving gives us our MLEs:

$$A = \frac{1}{N} \sum_{n=1}^N x_n$$

and

$$\sigma^2 = \frac{1}{N} \sum_{n=1}^N (x_n - A)^2$$

NOTE: σ^2 is biased!

As an exercise, try the following problem:

Exercise 4.3

Suppose we observe a random sample $x = (x_1, \dots, x_N)^T$ of Poisson measurements with intensity λ : $Pr[x_i = n] = e^{-\lambda} \frac{\lambda^n}{n!}$, $n \in \{0, 1, 2, \dots\}$. Find the MLE for λ .

Unfortunately, this approach is only feasible for the most elementary pdfs and pmfs. In general, we may have to resort to more advanced numerical maximization techniques:

1. **Newton-Raphson** iteration
2. Iteration by the **Scoring Method**
3. **Expectation-Maximization Algorithm**

All of these are iterative techniques which posit some initial guess at the MLE, and then incrementally update that guess. The iteration proceeds until a local maximum of the likelihood is attained, although in the case of the first two methods, such convergence is not guaranteed. The EM algorithm has the advantage that the likelihood is always increased at each iteration, and so convergence to at least a local maximum is guaranteed (assuming a bounded likelihood). For each algorithm, the final estimate is highly dependent on the initial guess, and so it is customary to try several different starting values. For details on these algorithms, see *Kay, Vol. I*[4].

4.2.4 Asymptotic Properties of the MLE

Let $x = (x_1, \dots, x_N)^T$ denote an IID sample of size N , and each sample is distributed according to $p(x | \theta)$. Let θ_N denote the MLE based on a sample x .

Theorem 4.3: Asymptotic Properties of MLE

If the likelihood $\ell(\theta | x) = p(x | \theta)$ satisfies certain "regularity" conditions⁵, then the MLE θ_N is **consistent**, and moreover, θ_N converges in probability to θ , where

$$\theta \sim \mathcal{N}\left(\theta, (I(\theta))^{-1}\right)$$

where $I(\theta)$ is the **Fisher Information matrix** evaluated at the true value of θ .

Since the mean of the MLE tends to the true parameter value, we say the MLE is **asymptotically unbiased**. Since the covariance tends to the inverse Fisher information matrix, we say the MLE is **asymptotically efficient**.

In general, the rate at which the mean-squared error converges to zero is not known. It is possible that for small sample sizes, some other estimator may have a smaller MSE. The proof of consistency is an application of the weak law of large numbers. Derivation of the asymptotic distribution relies on the central limit theorem. The theorem is also true in more general settings (e.g., dependent samples). See, *Kay, Vol. I, Ch. 7*[4] for further discussion.

⁵The regularity conditions are essentially the same as those assumed for the Cramer-Rao lower bound (Section 4.1.2): the log-likelihood must be twice differentiable, and the expected value of the first derivative of the log-likelihood must be zero.

4.2.5 The MLE and Efficiency

In some cases, the MLE is efficient, not just asymptotically efficient. In fact, when an efficient estimator exists, it must be the MLE, as described by the following result:

Theorem 4.4:

If θ is an efficient estimator, and the Fisher information matrix $I(\theta)$ is positive definite for all θ , then θ maximizes the likelihood.

Proof:

Recall the θ is efficient (meaning it is unbiased and achieves the Cramer-Rao lower bound) if and only if

$$\frac{\partial \ln(p(x|\theta))}{\partial \theta} = I(\theta)(\theta - \theta)$$

for all θ and x . Since θ is assumed to be efficient, this equation holds, and in particular it holds when $\theta = \theta(x)$. But then the derivative of the log-likelihood is zero at $\theta = \theta(x)$. Thus, θ is a critical point of the likelihood. Since the Fisher information matrix, which is the negative of the matrix of second order derivatives of the log-likelihood, is positive definite, θ must be a maximum of the likelihood.

An important case where this happens is described in the following subsection.

4.2.5.1 Optimality of MLE for Linear Statistical Model

If the observed data x are described by

$$x = H\theta + w$$

where H is $N \times p$ with full rank, θ is $p \times 1$, and $w \sim \mathcal{N}(0, C)$, then the MLE of θ is

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

This can be established in two ways. The first is to compute the CRLB for θ . It turns out that the condition for equality in the bound is satisfied, and θ can be read off from that condition.

The second way is to maximize the likelihood directly. Equivalently, we must minimize

$$(x - H\theta)^T C^{-1} (x - H\theta)$$

with respect to θ . Since C^{-1} is positive definite, we can write $C^{-1} = U^T \Lambda U = D^T D$, where $D = \Lambda^{\frac{1}{2}} U$, where U is an orthogonal matrix whose columns are eigenvectors of C^{-1} , and Λ is a diagonal matrix with positive diagonal entries. Thus, we must minimize

$$(Dx - DH\theta)^T (Dx - DH\theta)$$

But this is a linear least squares problem, so the solution is given by the pseudoinverse of DH :

$$\begin{aligned} \theta &= \left((DH)^T (DH) \right)^{-1} (DH)^T (Dx) \\ &= (H^T C^{-1} H)^{-1} H^T C^{-1} x \end{aligned} \tag{4.15}$$

Exercise 4.4

Consider $X_1, \dots, X_N \sim \mathcal{N}(s, \sigma^2 I)$, where s is a $p \times 1$ unknown signal, and σ^2 is known. Express the data in the linear model and find the MLE s for the signal.

4.2.6 Invariance of MLE

Suppose we wish to estimate the function $w = W(\theta)$ and not θ itself. To use the maximum likelihood approach for estimating w , we need an expression for the likelihood $\ell(w|x) = p(x|w)$. In other words, we would need to be able to parameterize the distribution of the data by w . If W is not a one-to-one function, however, this may not be possible. Therefore, we define the **induced** likelihood

$$\ell(w|x) = \max_{\theta} \{\theta, W(\theta) = w\} \ell(\theta|x)$$

The MLE w is defined to be the value of w that maximizes the induced likelihood. With this definition, the following invariance principle is immediate.

Theorem 4.5:

Let $\hat{\theta}$ denote the MLE of θ . Then $w = W(\hat{\theta})$ is the MLE of $w = W(\theta)$.

Proof:

The proof follows directly from the definitions of $\hat{\theta}$ and w . As an exercise, work through the logical steps of the proof on your own.

Example

Let $x = (x_1, \dots, x_N)^T$ where

$$x_i \sim \text{Poisson}(\lambda)$$

Given x , find the MLE of the probability that $x \sim \text{Poisson}(\lambda)$ exceeds the mean λ .

$$W(\lambda) = Pr[x > \lambda] = \sum_{n=\lfloor \lambda+1 \rfloor}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!}$$

where $\lfloor z \rfloor =$ largest integer $\leq z$. The MLE of w is

$$w = \sum_{n=\lfloor \lambda+1 \rfloor}^{\infty} e^{-\lambda} \frac{(\lambda)^n}{n!}$$

where $\hat{\lambda}$ is the MLE of λ :

$$\hat{\lambda} = \frac{1}{N} \sum_{n=1}^N x_n$$

Be aware that the MLE of a **transformed** parameter does not necessarily satisfy the asymptotic properties discussed earlier.

Exercise 4.5

Consider observations x_1, \dots, x_N , where x_i is a p -dimensional vector of the form $x_i = s + w_i$ where s is an unknown signal and w_i are independent realizations of white Gaussian noise:

$$w_i \sim \mathcal{N}(0, \sigma^2 I_{p \times p})$$

Find the maximum likelihood estimate of the energy $E = s^T s$ of the unknown signal.

4.2.7 Summary of MLE

The likelihood principle states that information brought by an observation x about θ is entirely contained in the likelihood function $p(x|\theta)$. The maximum likelihood estimator is **one** effective implementation of the likelihood principle. In some cases, the MLE can be computed exactly, using calculus and linear algebra, but at other times iterative numerical algorithms are needed. The MLE has several desirable properties:

- It is consistent and asymptotically efficient (as $N \rightarrow \infty$ we are doing as well as MVUE).
- When an efficient estimator exists, it is the MLE.
- The MLE is invariant to reparameterization.

4.3 Bayesian Estimation

4.3.1 Bayesian Estimation⁶

We are interested in estimating θ given the observation x . Naturally then, any estimation strategy will be based on the posterior distribution $p(\theta|x)$. Furthermore, we need a criterion for assessing the quality of potential estimators.

4.3.1.1 Loss

The quality of an estimate θ is measured by a real-valued **loss function**: $L(\theta, \theta)$. For example, squared error or quadratic loss is simply $L(\theta, \theta) = (\theta - \theta)^T (\theta - \theta)$

4.3.1.2 Expected Loss

Posterior Expected Loss:

$$E[L(\theta, \theta) | x] = \int L(\theta, \theta) p(\theta | x) d\theta$$

Bayes Risk:

$$\begin{aligned} E[L(\theta, \theta)] &= \int \int L(\theta, \theta) p(\theta | x) p(x) d\theta dx \\ &= \int \int L(\theta, \theta) p(x | \theta) p(\theta) dx d\theta \\ &= E[E[L(\theta, \theta) | x]] \end{aligned} \tag{4.16}$$

The "best" or optimal estimator given the data x and under a specified loss is given by

$$\theta = \underset{\theta}{\operatorname{argmin}} E[L(\theta, \theta) | x]$$

Example 4.13: Bayes MSE

$$\text{BMSE}(\theta) \equiv \int \int (\theta - \theta)^2 p(\theta | x) d\theta p(x) dx$$

Since $p(x) \geq 0$ for every x , minimizing the inner integral $\int (\theta - E[\theta])^2 p(\theta | x) d\theta = E[L(\theta, \theta) | x]$ (where $E[L(\theta, \theta) | x]$ is the posterior expected loss) for each x , minimizes the

⁶This content is available online at <<http://cnx.org/content/m11660/1.2/>>.

overall BMSE.

$$\begin{aligned} \frac{\partial \int (\theta - \hat{\theta})^2 p(\theta | x) d\theta}{\partial \hat{\theta}} &= \int \frac{\partial \left((\theta - \hat{\theta})^2 p(\theta | x) \right)}{\partial \hat{\theta} d\theta} \\ &= -2 \int (\theta - \hat{\theta}) p(\theta | x) d\theta \end{aligned} \quad (4.17)$$

Equating this to zero produces

$$\hat{\theta} = \int \theta p(\theta | x) d\theta \equiv E[\theta | x]$$

The conditional mean (also called **posterior mean**) of θ given x !

Example 4.14

$$x_n = A + W_n, \quad n \in \{1, \dots, N\}$$

$$W_n \sim \mathcal{N}(0, \sigma^2)$$

prior for unknown parameter A :

$$p(a) = U(-A_0, A_0)$$

$$p(x | A) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2}$$

$$p(A | x) = \begin{cases} \frac{\frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2}}{\int_{-A_0}^{A_0} \frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2} dA} & \text{if } |A| \leq A_0 \\ 0 & \text{if } |A| > A_0 \end{cases}$$

Minimum Bayes MSE Estimator:

$$\begin{aligned} A &= E[A | x] \\ &= \int_{-\infty}^{\infty} ap(A | x) dA \\ &= \frac{\int_{-A_0}^{A_0} A \frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2} dA}{\int_{-A_0}^{A_0} \frac{1}{2A_0(2\pi\sigma^2)^{\frac{N}{2}}} e^{-\frac{1}{2\sigma^2} \sum_{n=1}^N (x_n - A)^2} dA} \end{aligned} \quad (4.18)$$

Notes

1. No closed-form estimator
2. As $A_0 \rightarrow \infty$, $A \rightarrow \frac{1}{N} \sum_{n=1}^N x_n$
3. For smaller A_0 , truncated integral produces an A that is a function of x , σ^2 , and A_0
4. As N increases, $\frac{\sigma^2}{N}$ decreases and posterior $p(A | x)$ becomes tightly clustered about $\frac{1}{N} \sum x_n$.
This implies $A \rightarrow \frac{1}{N} \sum_n x_n$ as $n \rightarrow \infty$ (the data "swamps out" the prior)

4.3.1.3 Other Common Loss Functions

4.3.1.3.1 Absolute Error Loss

(Laplace, 1773)

$$\begin{aligned}
 L(\theta, \theta) &= |\theta - \theta| \\
 E[L(\theta, \theta) | x] &= \int_{-\infty}^{\infty} |\theta - \theta| p(\theta | x) d\theta \\
 &= \int_{-\infty}^{\theta} (\theta - \theta) p(\theta | x) d\theta + \int_{\theta}^{\infty} (\theta - \theta) p(\theta | x) d\theta
 \end{aligned} \tag{4.19}$$

Using integration-by-parts it can be shown that

$$\begin{aligned}
 \int_{-\infty}^{\theta} (\theta - \theta) p(\theta | x) d\theta &= \int_{-\infty}^{\theta} P(\theta < y | x) dy \\
 \int_{\theta}^{\infty} (\theta - \theta) p(\theta | x) d\theta &= \int_{\theta}^{\infty} P(\theta > y | x) dy
 \end{aligned}$$

where $P(\theta < y | x)$ and $P(\theta > y | x)$ are a cumulative distributions. So,

$$E[L(\theta, \theta) | x] = \int_{-\infty}^{\theta} P(\theta < y | x) dy + \int_{\theta}^{\infty} P(\theta > y | x) dy$$

Take the derivative with respect to θ implies $P(\theta < \theta | x) = P(\theta > \theta | x)$ which implies that the optimal θ under absolute error loss is **posterior median**.

4.3.1.3.2 '0-1' Loss

$$L(\theta, \theta) = \begin{cases} 0 & \text{if } \theta = \theta \\ 1 & \text{if } \theta \neq \theta \end{cases} = I_{\{\hat{\theta} \neq \theta\}}$$

$$E[L(\theta, \theta) | x] = E[I_{\{\hat{\theta} \neq \theta\}} | x] = Pr[\hat{\theta} \neq \theta | x]$$

which is the probability that $\hat{\theta} \neq \theta$ given x . To minimize '0-1' loss we must choose θ to be the value of θ with the highest posterior probability, which implies $\hat{\theta} \neq \theta$ with the smallest probability.




Figure 4.8

The optimal estimator θ under '0-1' loss is the **maximum a posteriori** (MAP) estimator—the value of θ where $p(\theta | x)$ is maximized.

4.3.2 LMMSE

4.3.2.1 Wiener Filtering and the DFT⁷

4.3.2.1.1 Connecting the Vector Space and Classical Wiener Filters

Suppose we observe

$$x = y + w$$

which are all $N \times 1$ vectors and where $w \sim \mathcal{N}(0, \sigma^2 I)$. Given x we wish to estimate y . Think of y as a signal in additive white noise w . x is a noisy observation of the signal.

Taking a Bayesian approach, put a prior on the signal y :

$$y \sim \mathcal{N}(0, R_{yy})$$

which is independent of noise w . The minimum MSE (MMSE) estimator is

$$y = R_{yx} R_{xx}^{-1} x$$

Under the modeling assumptions above

$$\begin{aligned}
 R_{yx} &= E[y(y+w)^T] \\
 &= E[yy^T] + E[yw^T] \\
 &= E[yy^T] \\
 &= R_{yy}
 \end{aligned} \tag{4.20}$$

⁷This content is available online at <<http://cnx.org/content/m11462/1.5/>>.

since $E[yw^T] = 0$ and since y and w are zero-mean and independent.

$$\begin{aligned}
 R_{xx} &= E[xx^T] \\
 &= E[(y+w)(y+w)^T] \\
 &= E[yy^T] + E[yw^T] + E[wy^T] + E[ww^T] \\
 &= R_{yy} + R_{ww}
 \end{aligned} \tag{4.21}$$

since $E[ww^T] = R_{ww}$. Hence

$$y = R_{yy}(R_{yy} + R_{ww})^{-1}x = H_{\text{opt}}x$$

Where H_{opt} is the Wiener filter. Recall the frequency domain case

$$H_{\text{opt}}(f) = \frac{S_{yy}(f)}{S_{yy}(f) + S_{ww}(f)}$$

Now let's look at an actual problem scenario. Suppose that we know *a priori* that the signal y is **smooth** or **lowpass**. We can incorporate this prior knowledge by carefully choosing the prior covariance R_{yy} .

Recall the DFT

$$\mathcal{Y}_k = \frac{1}{\sqrt{N}} \sum_{n=0}^N Y_k e^{-i2\pi \frac{kn}{N}}, \quad k = 0, \dots, N-1$$

or in vector notation

$$\mathcal{Y}_k = \langle y, u_k \rangle, \quad k = 0, \dots, N-1$$

where $u_k = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & e^{i2\pi \frac{k}{N}} & e^{i2\pi \frac{2k}{N}} & \dots & e^{i2\pi \frac{(N-1)k}{N}} \end{pmatrix}^H$ (H denotes Hermitian transpose)

NOTE: $\langle u_k, u_k \rangle = u_k^H u_k = 1$, $\langle u_k, u_l \rangle = u_k^H u_l = 0$, $k \neq l$, i.e., $\{u_k\}$, $k = 0, \dots, N-1$ is an orthonormal basis.

The vector u_k spans the subspace corresponding to a frequency band centered at frequency $f_k = \frac{2\pi k}{N}$ ("digital" frequency on $[0, 1]$). If we know that y is lowpass, then

$$E[(\langle y, u_k \rangle)^2] = E[(\|\mathcal{Y}_k\|)^2]$$

should be relatively small (compared to $E[(\langle y, u_0 \rangle)^2]$) for high frequencies.

Let

$$\sigma_k^2 = E[(\langle y, u_k \rangle)^2]$$

A lowpass model implies $\sigma_0^2 > \sigma_1^2 > \dots > \sigma_{\frac{N}{2}}^2$, assuming N even, and conjugate symmetry implies $\sigma_{N-j}^2 = \sigma_j^2$, $j = 1, \dots, \frac{N}{2}$. Furthermore, let's model the DFT coefficients as zero-mean and independent

$$E[\mathcal{Y}_k] = 0$$

$$E[\mathcal{Y}_k \overline{\mathcal{Y}_l}] = \begin{cases} \sigma_k^2 & \text{if } l = k \\ 0 & \text{if } l \neq k \end{cases}$$

This completely specifies our prior

$$y \sim \mathcal{N}(0, R_{yy})$$

$$R_{yy} = U D \overline{U}^T$$

where

$$D = \begin{pmatrix} \sigma_0^2 & 0 & \dots & 0 \\ 0 & \sigma_1^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{N-1}^2 \end{pmatrix}$$

and

$$U = \begin{pmatrix} u_0 & u_1 & \dots & u_{N-1} \end{pmatrix}$$

NOTE:

$$\mathcal{Y} = U^H y$$

is the DFT and

$$y = U \mathcal{Y}$$

is the inverse DFT.

With this prior on y the Wiener filter is

$$y = U D U^H (U D U^H + \sigma^2 I)^{-1} x$$

Since U is a unitary matrix $U U^H = I$ and therefore

$$\begin{aligned} y &= U D U^H (U (D + \sigma^2 I) U^H)^{-1} x \\ &= U D U^H U (D + \sigma^2 I)^{-1} U^H x \\ &= U D (D + \sigma^2 I)^{-1} U^H x \end{aligned} \tag{4.22}$$

⁸ Now take the DFT of both sides

$$\mathcal{Y} = U^H y = D (D + \sigma^2 I)^{-1} \mathcal{X}$$

where $\mathcal{X} = U^H x$ and is the DFT of x . Both D and $D + \sigma^2 I$ are diagonal so

$$\mathcal{Y}_k = \frac{d_{k,k}}{d_{k,k} + \sigma^2} \mathcal{X}_k = \frac{\sigma_k^2}{\sigma_k^2 + \sigma^2} \mathcal{X}_k$$

Hence the Wiener filter is a frequency (DFT) domain filter

$$\mathcal{Y}_k = H_k \mathcal{X}_k$$

where \mathcal{X}_k is the k^{th} DFT coefficient of x and the filter response at digital frequency $\frac{2\pi k}{N}$ is

$$H_k = \frac{\sigma_k^2}{\sigma_k^2 + \sigma^2}$$

Assuming $\sigma_0^2 > \sigma_1^2 > \dots > \sigma_{\frac{N}{2}}^2$ and $\sigma_{N-j}^2 = \sigma_j^2$, $j = 1, \dots, \frac{N}{2}$. The filter's response is a **digital lowpass filter!**

⁸If A, B, C are all invertible, compatible matrices, then $(ABC)^{-1} = C^{-1}B^{-1}A^{-1}$. $U^{-1} = U^H$, $(U^H)^{-1} = U$.

A Digital Lowpass Filter!




Figure 4.9

4.3.2.1.2 Summary of Wiener Filter

Problem: Observe $x = y + w$

Recover/estimate signal y .

Classical Wiener Filter (continuous-time):

$$H(\omega) = \frac{S_{yy}(\omega)}{S_{yy}(\omega) + S_{ww}(\omega)}$$

where $y(t)$ and $w(t)$ are stationary processes.

Vector Space Wiener Filter:

$$H = R_{yy}(R_{yy} + R_{ww})^{-1}$$

Wiener Filter and DFT: ($R_{ww} = \sigma^2 I$). If $R_{yy} = UDU^H$, where U is DFT, then H is a discrete-time filter whose DFT is given by

$$\begin{aligned} H_k &= \sum_{n=0}^{N-1} h_n e^{i2\pi \frac{kn}{N}} \\ &= \frac{d_{k,k}}{d_{k,k} + \sigma^2} \end{aligned} \quad (4.23)$$

Here, $d_{k,k}$ plays the same role as $S_{yy}(\omega)$.

4.3.2.2 Kalman Filters⁹

The Kalman filter is an important generalization of the Wiener filter. Unlike Wiener filters, which are designed under the assumption that the signal and noise are **stationary**, the Kalman filter has the ability to **adapt** itself to non-stationary environments.

⁹This content is available online at <<http://cnx.org/content/m11438/1.7/>>.

The Kalman filter can be viewed as a **sequential** minimum MSE estimator of a signal in additive noise. If the signal and noise are jointly Gaussian, the then Kalman filter is optimal in a minimum MSE sense (minimizes expected quadratic loss).

If the signal and/or noise are non-Gaussian, then the Kalman filter is the best linear estimator (linear estimator that minimizes MSE among all possible linear estimators).

4.3.2.2.1 Dynamical Signal Models

Recall the simple DC signal estimation problem.

$$x_n = A + w_n \quad , \quad n = \{0, \dots, N - 1\} \quad (4.24)$$

Where A is the unknown DC level and w_n is the white Gaussian noise. A could represent the voltage of a DC power supply. We know how to find several good estimators of A given the measurements $\{x_0, \dots, x_{N-1}\}$.

In practical situations this model may be too simplistic. the load on the power supply may charge over time and there will be other variations due to temperature and component aging.

To account for these variations we can employ a more accurate measurement model:

$$x_n = A_n + w_n \quad , \quad n = \{0, \dots, N - 1\} \quad (4.25)$$

where the voltage A_n is the **true** voltage at time n .

Now the estimation problem is significantly more complicated since we must estimate $\{A_0, \dots, A_{N-1}\}$. Suppose that the true voltage A_n does not vary too rapidly over time. Then successive samples of A_n will not be too different, suggesting that the voltage signal displays a high degree of **correlation**.

This reasoning suggests that it may be reasonable to regard the sequence $\{A_0, \dots, A_{N-1}\}$, as a realization of a correlated (not white) random process. Adopting a random process model for A_n allows us to pursue a Bayesian approach to the estimation problem (Figure 4.10 (Voltage Varying Over Time)).

Voltage Varying Over Time




Figure 4.10

Using the model in (4.25), it is easy to verify that the maximum likelihood and MVUB estimators are given by

$$A_n = x_n \quad (4.26)$$

Our estimate is simply the noisy measurements! No averaging takes place, so there is no noise reduction. Let's look at the example again, Figure 4.11 (True Voltage Varying Over Time).

True Voltage Varying Over Time




Figure 4.11

The voltage A_n is varying about an average value of 10V. Assume this average value is known and write

$$A_n = 10 + y_n \quad (4.27)$$

Where y_n is a zero-mean random process. Now a simple model for y_n which allows us to specify the correlation between samples is the **first-order Gauss-Markov** process model:

$$y_n = ay_{n-1} + u_n \quad , \quad n = \{1, 2, \dots\} \quad (4.28)$$

Where $u_n \sim \mathcal{N}(0, \sigma_u^2)$ iid (white Gaussian noise process). To initialize the process we take y_0 to be the realization of a Gaussian random variable: $y_0 \sim \mathcal{N}(0, \sigma_y^2)$. u_n is called the **driving** or **excitation** noise. The model in (4.28) is called the **dynamical** or **state** model. The current output y_n depends only on the **state** of the system at the previous time, or y_{n-1} , and the current input u_n (Figure 4.12).




Figure 4.12

$$y_1 = ay_0 + u_0$$

$$\begin{aligned} y_2 &= ay_1 + u_1 \\ &= a(ay_0 + u_0) + u_1 \\ &= a^2y_0 + au_1 + u_2 \end{aligned} \tag{4.29}$$

⋮

$$y_n = a^{n+1}y_0 + \sum_{k=1}^n a^k u_{n-k}$$

$$\begin{aligned} E[y_n] &= a^{n+1}E[y_0] + \sum_{k=1}^n a^k E[u_{n-k}] \\ &= 0 \end{aligned} \tag{4.30}$$

Correlation:

$$\begin{aligned} E[y_m y_n] &= E\left[(a^{m+1}y_0 + \sum_{k=1}^m a^k u_{m-k})(a^{n+1}y_0 + \sum_{l=1}^n a^l u_{n-l})\right] \\ &= E[a^{m+n+2}y_0^2] + E\left[\sum_{k=1}^m \sum_{l=1}^n a^{k+l} u_{m-k} u_{n-l}\right] \end{aligned} \tag{4.31}$$

$$E[u_{m-k} u_{n-l}] = \begin{cases} \sigma_n^2 & \text{if } m - k = n - l \\ 0 & \text{otherwise} \end{cases} \tag{4.32}$$

If $m > n$, then

$$E[y_m y_n] = a^{m+n+2}\sigma_y^2 + a^{m-n}\sigma_u^2 \sum_{k=1}^n a^{2k} \tag{4.33}$$

If $|a| > 1$, then it's obvious that the process diverges (variance $\rightarrow \infty$). This is equivalent to having a pole outside the unit circle shown in Figure 4.13.




Figure 4.13

So, let's assume $|a| < 1$ and hence a stable system. Thus as m and n get large

$$a^{m+n+2}\sigma_y^2 \rightarrow 0$$

Now let $m - n = \tau$. Then for m and n large we have

$$\begin{aligned} E[y_m y_n] &= a^\tau \sigma_u^2 \sum_{k=1}^n a^{2k} \\ &= \frac{a^{\tau+2} \sigma_u^2}{1-a^2} \end{aligned} \quad (4.34)$$

This shows us how correlated the process is:

$$|a| \rightarrow 1 \Rightarrow \text{heavily correlated (or anticorrelated)}$$

$$|a| \rightarrow 0 \Rightarrow \text{weakly correlated}$$

How can we use this model to help us in our estimation problem?

4.3.2.2.2 The Kalman Filter

Let's look at a more general formulation of the problem at hand. Suppose that we have a vector-valued dynamical equation

$$y_{n+1} = Ay_n + bu_n \quad (4.35)$$

Where y_n is $p \times 1$ dimensional, A is $p \times p$, and b is $p \times 1$. The initial **state vector** is $Y_0 \sim \mathcal{N}(0, R_0)$, where R_0 is the covariance matrix and $u_n \sim \mathcal{N}(0, \sigma_u^2)$ iid (white Gaussian **excitation** noise). This reduces to the case we just looked at when $p = 1$. This model could represent a p^{th} order Gauss-Markov process:

$$y_{n-1} = a_1 y_n + a_2 y_{n-1} + \dots + a_p y_{n-p+1} + u_n \quad (4.36)$$

Define

$$y_n = \begin{pmatrix} y_{n-p+1} \\ y_{n-p+2} \\ \vdots \\ y_{n-1} \\ y_n \end{pmatrix} \quad (4.37)$$

Then,

$$\begin{aligned} y_{n+1} &= Ay_n + bu_n \\ &= \begin{pmatrix} 0 & 1 & 0 & \dots & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & 0 & \dots & \dots & 0 & 1 \\ a_1 & a_2 & \dots & \dots & a_{p-1} & a_p \end{pmatrix} \begin{pmatrix} y_{n-p+1} \\ y_{n-p+2} \\ \vdots \\ \vdots \\ y_{n-1} \\ y_n \end{pmatrix} + \begin{pmatrix} 0 \\ \vdots \\ \vdots \\ \vdots \\ 0 \\ 1 \end{pmatrix} + u_n \end{aligned} \quad (4.38)$$

Here A is the **state transition matrix**. Since y_n is a linear combination of Gaussian vectors:

$$y_n = A^2 y_0 + \sum_{k=1}^n A^{k-1} b u_{n-k} \quad (4.39)$$

We know that y_n is also Gaussian distributed with mean and covariance $R_n = E[y_n y_n^T]$, $Y_n \sim \mathcal{N}(, R_n)$. The covariance can be recursively computed from the basic state equation:

$$R_{n+1} = AR_n A^T + \sigma_u^2 b b^T \quad (4.40)$$

Assume that measurements of the state are available:

$$x_n = C^T y_n + w_n \quad (4.41)$$

Where $w_n \sim \mathcal{N}(0, \sigma_w^2)$ iid independent of $\{u_n\}$ (white Gaussian **observation** noise).

For example, if $C = (0, \dots, 0, 1)^T$, then

$$x_n = y_n + w_n \quad (4.42)$$

Where x_n is the observation, y_n is the signal, and w_n is the noise. Since our model for the signal is Gaussian as well as the observation noise, it follows that $x_n \sim \mathcal{N}(0, \sigma_n^2)$, where $\sigma_n^2 = C^T R_n C + \sigma_w^2$ (Figure 4.14 (Block Diagram)).

Block Diagram




Figure 4.14

Kalman first posed the problem of estimating the state of y_n from the sequence of measurements

$$x_n = \begin{pmatrix} x_0 \\ \vdots \\ x_n \end{pmatrix}$$

To derive the Kalman filter we will call upon the Gauss-Markov Theorem.

First note that the conditional distribution of y_n given x_n is Gaussian:

$$y_n | x_n \sim \mathcal{N}(y_{n|n}, P_{n|n})$$

Where $y_{n|n}$ is the conditional mean and $P_{n|n}$ is the covariance.

We **know** that this is the **form** of the conditional distribution because y_n and x_n are **jointly** Gaussian distributed.

NOTE:

$$y_n | x_n \sim \mathcal{N}(y_{n|n}, P_{n|n})$$

where y_n is the signal samples y_n, \dots, y_{n-p+1} , x_n is the observations/measurements x_n, \dots, x_{n-p+1} , and $y_{n|n}$ is the best (minimum MSE) estimator of y_n given x_n .

This is all well and good, but we need to know what the conditional mean and covariance are explicitly. So the problem is now to find/compute $y_{n|n}$ and $P_{n|n}$. We can take advantage of the recursive state equation to obtain a recursive procedure for this calculation. To begin, consider the "predictor" $y_{n|n-1}$:

$$y_n|x_{n-1} \sim \mathcal{N}(y_{n|n-1}, P_{n|n-1})$$

Where y_n is the signal samples, $\{y_n, \dots, y_{n-p+1}\}$, x_{n-1} is the observations $\{x_{n-1}, \dots, x_{n-p}\}$, and $y_{n|n-1}$ is the best min MSE estimator of y_n given x_{n-1} . Although we don't know what forms $y_{n|n-1}$ and $P_{n|n-1}$ have, we do know two important facts:

1. The predictor $y_{n|n-1}$ acts as a sufficient statistic for y_n . That is, we can replace x_{n-1} (the data) with $y_{n|n-1}$ (the predictor). In other words, all the relevant information in x_{n-1} pertaining to y_n is summarized by the predictor $y_{n|n-1}$, which is, of course, a function of x_{n-1} .
2. The predictor $y_{n|n-1}$ and the prediction error $e_{n|n-1} = y_n - y_{n|n-1}$ are orthogonal (the **orthogonality principle** of minimum MSE estimators $\Rightarrow E[y_{n|n-1} e_{n|n-1}^T] = 0$) \Rightarrow error is orthogonal to estimator).

Moreover,

$$y_n = y_{n|n-1} + e_{n|n-1} \quad (4.43)$$

Since all quantities are zero-mean,

$$e_{n|n-1} \sim \mathcal{N}(0, P_{n|n-1})$$

where $P_{n|n-1}$ is the covariance of $y_n|x_{n-1}$ and "variability" of y_n about the predictor $y_{n|n-1}$. Therefore,

$$y_{n|n-1} \sim \mathcal{N}(0, R_n - P_{n|n-1})$$

Where R_n is the covariance of Y_n . Now suppose that we have the predictor $y_{n|n-1}$ computed and a new measurement is made:

$$\begin{aligned} x_n &= C^T y_n + w_n \\ &= C^T (y_{n|n-1} + e_{n|n-1}) + w_n \end{aligned} \quad (4.44)$$

NOTE: $y_{n|n-1}$, $e_{n|n-1}$, and w_n are all **orthogonal**.

We can express all relevant quantities in the matrix equation

$$\begin{pmatrix} y_n \\ y_{n|n-1} \\ x_n \end{pmatrix} = \begin{pmatrix} I & I & 0 \\ 0 & I & 0 \\ C^T & C^T & 1 \end{pmatrix} \begin{pmatrix} e_{n|n-1} \\ y_{n|n-1} \\ w_n \end{pmatrix} \quad (4.45)$$

Now because of the orthogonality, the covariance

$$E \left[\begin{pmatrix} e_{n|n-1} \\ y_{n|n-1} \\ w_n \end{pmatrix} \begin{pmatrix} e_{n|n-1} \\ y_{n|n-1} \\ w_n \end{pmatrix}^T \right] = \begin{pmatrix} P_{n|n-1} & 0 & 0 \\ 0 & R_n - P_{n|n-1} & 0 \\ 0 & 0 & \sigma_w^2 \end{pmatrix} \quad (4.46)$$

Combining this with the matrix (4.45) shows that

$$E \left[\begin{pmatrix} y_n \\ y_{n|n-1} \\ x_n \end{pmatrix} \begin{pmatrix} y_n \\ y_{n|n-1} \\ x_n \end{pmatrix}^T \right] = \begin{pmatrix} R_n & P_{n|n-1} & R_n C \\ P_{n|n-1} & \vdots & \dots \\ C^T R_n & \vdots & S_n \end{pmatrix} \quad (4.47)$$

NOTE: $(y_n, y_{n|n-1}, x_n)^T$ are jointly Gaussian with the covariance in (4.47) and means zero.

Where

$$P_{n|n-1} = R_n - P_{n|n-1} \quad (4.48)$$

$$S_n = \begin{pmatrix} P_{n|n-1} & P_{n|n-1} C \\ C^T P_{n|n-1} & \sigma_w^2 \end{pmatrix} \quad (4.49)$$

We now have all the quantities necessary to compute our recursive estimator using the Gauss-Markov Theorem.

We will now derive a recursion for conditional distribution of y_n given $y_{n|n-1}$ (best estimate based on past observations) and x_n (current observation). We know that $y_n | (y_{n|n-1}, x_n)$ is Gaussian (since all quantities are jointly Gaussian). Let's denote this conditional distribution by

$$y_n | (y_{n|n-1}, x_n) \sim \mathcal{N}(y_{n|n}, P_{n|n})$$

Applying the Gauss-Markov Theorem we find

$$y_{n|n} = \begin{pmatrix} P_{n|n-1} \\ R_n C \end{pmatrix}^T S_n^{-1} \begin{pmatrix} y_{n|n-1} \\ x_n \end{pmatrix} \quad (4.50)$$

which is the best estimator of y_n given $y_{n|n-1}$ and x_n . The inverse of S_n is given by

$$S_n^{-1} = \begin{pmatrix} (P_{n|n-1})^{-1} & 0 \\ 0 & 0 \end{pmatrix} + \gamma_n^{-1} \begin{pmatrix} -C \\ 1 \end{pmatrix} \begin{pmatrix} -C^T \\ 1 \end{pmatrix}^T \quad (4.51)$$

where

$$\begin{aligned} \gamma_n^{-1} &= \sigma_n^2 - C^T (R_n - P_{n|n-1}) C \\ &= C^T P_{n|n-1} C + \sigma_w^2 \end{aligned} \quad (4.52)$$

NOTE: $\sigma_n^2 = C^T R_n C + \sigma_w^2$

Substituting this inverse formula into (4.50) yields

$$y_{n|n} = y_{n|n-1} + P_{n|n-1} C \gamma_n^{-1} (x_n - C^T y_{n|n-1}) \quad (4.53)$$

The Gauss-Markov Theorem also gives us an expression for $P_{n|n}$:

$$P_{n|n} = R_n - \begin{pmatrix} P_{n|n-1} \\ R_n C \end{pmatrix}^T S_n^{-1} \begin{pmatrix} P_{n|n-1} \\ C^T R_n \end{pmatrix} \quad (4.54)$$

and upon substituting (4.50) for S_n^{-1} we get

$$P_{n|n} = P_{n|n-1} - \gamma_n^{-1} P_{n|n-1} C C^T P_{n|n-1} \quad (4.55)$$

Note that both expressions contain the quantity

$$K_n = P_{n|n-1} C \gamma_n^{-1} \quad (4.56)$$

which is the so-called **Kalman gain**.

Using the Kalman gain, the **Kalman recursions** are given by

$$y_{n|n} = y_{n|n-1} + K_n \left(x_n - C^T y_{n|n-1} \right) \quad (4.57)$$

$$P_{n|n} = P_{n|n-1} - \gamma_n K_n K_n^T \quad (4.58)$$

The recursions are complete except for definitions of $y_{n|n-1}$ and $P_{n|n-1}$.

$$\begin{aligned} y_{n|n-1} &= E[y_n | x_{n-1}] \\ &= E[Ay_{n-1} + bu_{n-1} | x_{n-1}] \\ &= Ay_{n-1|n-1} \end{aligned} \quad (4.59)$$

$$\begin{aligned} P_{n|n-1} &= E \left[\left(y_n - y_{n|n-1} \right) \left(y_n - y_{n|n-1} \right)^T \right] \\ &= AP_{n-1|n-1} A^T + \sigma_n^2 bb^T \end{aligned} \quad (4.60)$$

Now we can summarize the **Kalman filter**:

1. (4.57), where $y_{n|n}$ is the best estimate if y_n given observations up to time n .
2. (4.59)
3. (4.56)
4. (4.52)
5. (4.58)
6. (4.60)

Measurements/observation model:

$$x_n = y_n + w_n$$

$$w_n \sim \mathcal{N}(0, \sigma_w^2)$$

($C = 1$).

Example 4.15: First-order Gauss-Markov Process

$$y_{n+1} = ay_n + u_n$$

Where y_{n+1} is a time-varying voltage, $u_n \sim \mathcal{N}(0, \sigma_u^2)$, and $\sigma_u = 0.1$. ($a = 0.99$) \Rightarrow highly correlated process. ($A = a$, $b = 1$).





Figure 4.15

Kalman Filtering Equations

1. $P_{n|n-1} = a^2 P_{n-1|n-1} + \sigma_n^2$ ($q(n)$ in MATLAB code)
2. $\gamma_n^{-1} = P_{n|n-1} + \sigma_w^2$ ($g(n)$ in MATLAB code)
3. $P_{n|n} = P_{n|n-1} - \gamma_n^{-1} P_{n|n-1}^2$ ($p(n)$ in MATLAB code)
4. $K_n = P_{n|n-1} \gamma_n^{-1}$ ($k(n)$ in MATLAB code)
5. $y_{n|n-1} = a y_{n-1|n-1}$ ($py(n)$ in MATLAB code)
6. $y_{n|n} = y_{n|n-1} + K_n (x_n - y_{n|n-1})$ ($ey(n)$ in MATLAB code)

Initialization: $ey(1) = 0$, $q(1) = \sigma_u^2$

**Figure 4.16**

Chapter 5

Adaptive Filtering

5.1 Introduction to Adaptive Filtering¹

The Kalman filter is just one of many **adaptive** filtering (or estimation) algorithms. Despite its elegant derivation and often excellent performance, the Kalman filter has two drawbacks:

1. The derivation and hence performance of the Kalman filter depends on the accuracy of the *a priori* assumptions. The performance can be less than impressive if the assumptions are erroneous.
2. The Kalman filter is fairly computationally demanding, requiring $O(P^2)$ operations per sample. This can limit the utility of Kalman filters in high rate real time applications.

As a popular alternative to the Kalman filter, we will investigate the so-called **least-mean-square** (LMS) adaptive filtering algorithm.

The principle advantages of LMS are

1. No prior assumptions are made regarding the signal to be estimated.
2. Computationally, LMS is very efficient, requiring $O(P)$ per sample.

The price we pay with LMS instead of a Kalman filter is that the rate of convergence and adaptation to sudden changes is slower for LMS than for the Kalman filter (with correct prior assumptions).

¹This content is available online at <<http://cnx.org/content/m11433/1.5/>>.

5.1.1 Adaptive Filtering Applications

5.1.1.1 Channel/System Identification

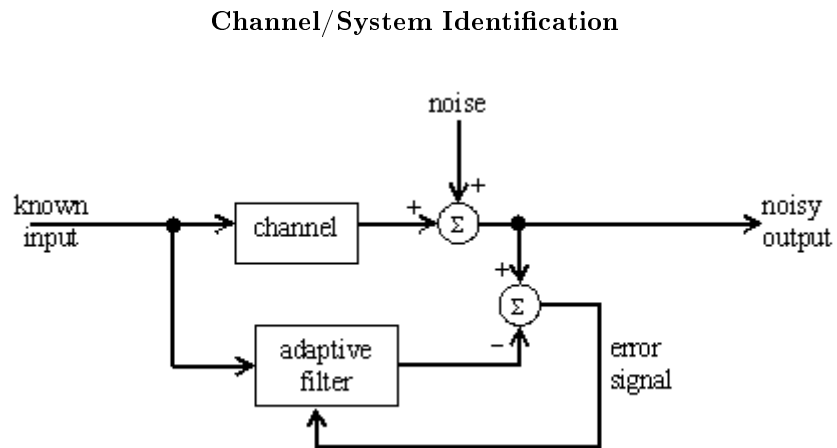


Figure 5.1

5.1.1.2 Noise Cancellation

Suppression of maternal ECG component in fetal ECG (Figure 5.2).

Image not finished

Figure 5.2: Cancelling maternal heartbeat in fetal electrocardiography (ECG): position of leads.

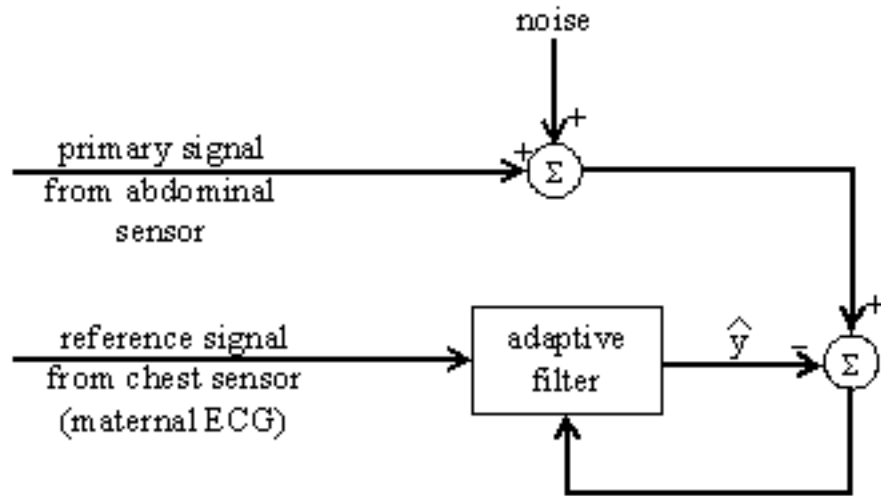


Figure 5.3

y is an estimate of the maternal ECG signal present in the abdominal signal (Figure 5.4).

Image not finished

Figure 5.4: Results of fetal ECG experiment (bandwidth, 3-35Hz; sampling rate, 256Hz): (a)reference input (chest lead); (b)primary input (abdominal lead); (c)noise-canceller output.

5.1.1.3 Channel Equalization

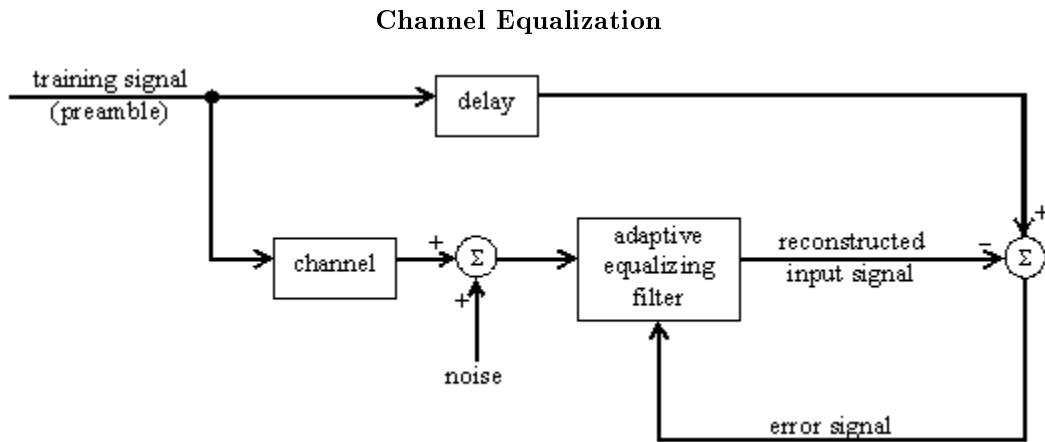


Figure 5.5

5.1.1.4 Adaptive Controller

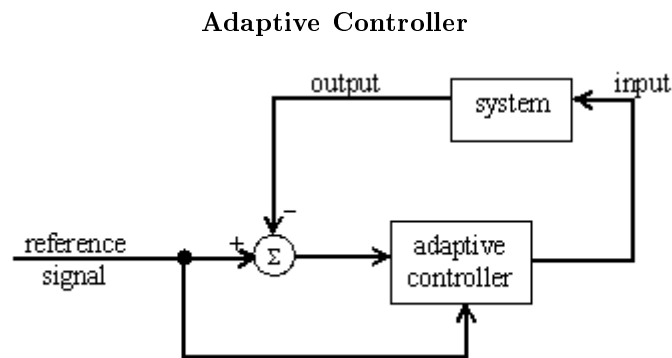


Figure 5.6: Here, the reference signal is the desired output. The adaptive controller adjusts the controller gains (filter weights) to keep them appropriate to the system as it changes over time.

5.1.2 Iterative Minimization

Most adaptive filtering algorithms (LMS included) are modifications of standard iterative procedures for solving minimization problems in a **real-time** or **on-line** fashion. Therefore, before deriving the LMS algorithm we will look at iterative methods of minimizing error criteria such as MSE.

Consider the following set-up:

x_k : observation

y_k : signal to be estimated

5.1.2.1 Linear estimator

$$y_k = w_1 x_k + w_2 x_{k-1} + \dots + w_p x_{k-p+1} \quad (5.1)$$



Figure 5.7

Impulse response of the filter:

$$\dots, 0, 0, w_1, w_2, \dots, w_p, 0, 0, \dots$$

5.1.2.2 Vector notation

$$y_k = x_k^T w \quad (5.2)$$

Where

$$x_k = \begin{pmatrix} x_k \\ x_{k-1} \\ \vdots \\ x_{k-p+1} \end{pmatrix}$$

and

$$w = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_p \end{pmatrix}$$

5.1.2.3 Error signal

$$\begin{aligned} e_k &= y_k - \hat{y}_k \\ &= y_k - x_k^T w \end{aligned} \quad (5.3)$$

5.1.2.4 Assumptions

(x_k, y_k) are jointly stationary with zero-mean.

5.1.2.5 MSE

$$\begin{aligned}
 E[e_k^2] &= E[(y_k - x_k^T w)^2] \\
 &= E[y_k^2] - 2w^T E[x_k y_k] + w^T E[x_k x_k^T] w \\
 &= R_{yy} - 2w^T R_{xy} + w^T R_{xx} w
 \end{aligned} \tag{5.4}$$

Where R_{yy} is the variance of y_k^2 , R_{xx} is the covariance matrix of x_k , and $R_{xy} = E[x_k y_k]$ is the cross-covariance between x_k and y_k

NOTE: The MSE is quadratic in W which implies the MSE surface is "bowl" shaped with a unique minimum point (Figure 5.8).




Figure 5.8

5.1.2.6 Optimum Filter

Minimize MSE:

$$\left(\frac{\partial E[e_k^2]}{\partial w} = 2R_{xy} + 2R_{xx}w = 0 \right) \Rightarrow (w_{\text{opt}} = R_{xx}^{-1}R_{xy}) \tag{5.5}$$

Notice that we can re-write (5.5) as

$$E[x_k x_k^T w] = E[x_k y_k] \tag{5.6}$$

or

$$\begin{aligned}
 E[x_k (y_k - x_k^T w)] &= E[x_k e_k] \\
 &= 0
 \end{aligned} \tag{5.7}$$

Which shows that the error signal is orthogonal to the input x_k (by the **orthogonality principle** of minimum MSE estimator).

5.1.2.7 Steepest Descent

Although we can easily determine w_{opt} by solving the system of equations

$$R_{xx}w = R_{xy} \quad (5.8)$$

Let's look at an iterative procedure for solving this problem. This will set the stage for our adaptive filtering algorithm.

We want to minimize the MSE. The idea is simple. Starting at some initial weight vector w_0 , iteratively adjust the values to decrease the MSE (Figure 5.9 (In One-Dimension)).

In One-Dimension




Figure 5.9

We want to **move** w_0 towards the optimal vector w_{opt} . In order to move in the correct direction, we must move **downhill** or in the direction opposite to the gradient of the MSE surface at the point w_0 . Thus, a natural and simple adjustment takes the form

$$w_1 = w_0 - \frac{1}{2}\mu \left. \frac{\partial E[e_k^2]}{\partial w} \right|_{w=w_0} \quad (5.9)$$

Where μ is the step size and tells us how far to move in the negative gradient direction (Figure 5.10).




Figure 5.10

Generalizing this idea to an iterative strategy, we get

$$w_k = w_{k-1} - \frac{1}{2} \mu \frac{\partial E[e_k^2]}{\partial w} \Big|_{w=w_{k-1}} \quad (5.10)$$

and we can repeatedly update w : w_0, w_1, \dots, w_k . Hopefully each subsequent w_k is closer to w_{opt} . Does the procedure converge? Can we adapt it to an on-line, real-time, dynamic situation in which the signals may not be stationary?

5.2 LMS Algorithm Analysis²

5.2.1 Objective

Minimize instantaneous squared error

$$e_k^2(w) = (y_k - x_k^T w)^2 \quad (5.11)$$

5.2.2 LMS Algorithm

$$w_k = w_{k-1} + \mu x_k e_k \quad (5.12)$$

Where w_k is the new weight vector, w_{k-1} is the old weight vector, and $\mu x_k e_k$ is a small step in the instantaneous error gradient direction.

²This content is available online at <<http://cnx.org/content/m11439/1.4/>>.

5.2.3 Interpretation in Terms of Weight Error Vector

Define

$$v_k = w_k - w_{\text{opt}} \quad (5.13)$$

Where w_{opt} is the optimal weight vector and

$$\epsilon_k = y_k - x_k^T w_{\text{opt}} \quad (5.14)$$

where ϵ_k is the minimum error. The stochastic difference equation is:

$$v_k = I v_{k-1} + \mu x_k \epsilon_k \quad (5.15)$$

5.2.4 Convergence/Stability Analysis

Show that (tightness)

$$\lim_{B \rightarrow \infty} \max \{Pr [\|v_k\| \geq B]\} = 0 \quad (5.16)$$

With probability 1, the weight error vector is bounded for all k .

Chebyshev's inequality is

$$Pr [\|v_k\| \geq B] \leq \frac{E [\|v_k\|^2]}{B^2} \quad (5.17)$$

and

$$Pr [\|v_k\| \geq B] = \frac{1}{B^2} \left((E[v_k])^2 + \sigma(v_k)^2 \right) \quad (5.18)$$

where $(E[v_k])^2$ is the squared bias. If $(E[v_k])^2 + \sigma(v_k)^2$ is finite for all k , then $\lim_{B \rightarrow \infty} Pr [\|v_k\| \geq B] = 0$ for all k .

Also,

$$\sigma(v_k)^2 = \text{tr} (E[v_k v_k^T]) \quad (5.19)$$

Therefore $\sigma(v_k)^2$ is finite if the diagonal elements of $\Gamma_k \equiv E[v_k v_k^T]$ are bounded.

5.2.5 Convergence in Mean

$\|E[v_k]\| \rightarrow 0$ as $k \rightarrow \infty$. Take expectation of (5.15) using smoothing property to simplify the calculation. We have convergence in mean if

1. R_{xx} is positive definite (invertible).
2. $\mu < \frac{2}{\lambda_{\max}(R_{xx})}$.

5.2.6 Bounded Variance

Show that $\Gamma_k = E[v_k v_k^T]$, the weight vector error covariance is bounded for all k .

NOTE: We could have $E[v_k] \rightarrow 0$, but $\sigma(v_k)^2 \rightarrow \infty$; in which case the algorithm would not be stable.

Recall that it is fairly straightforward to show that the diagonal elements of the transformed covariance $C_k = U \Gamma_k U^T$ tend to zero if $\mu < \frac{1}{\lambda_{\max}(R_{xx})}$ (U is the eigenvector matrix of R_{xx} ; $R_{xx} = U D U^T$). The diagonal elements of C_k were denoted by $\gamma_{k,i}$, $i = \{1, \dots, p\}$.

NOTE: $\sigma(v_k)^2 = \text{tr}(\Gamma_k) = \text{tr}(U^T C_k U) = \text{tr}(C_k) = \sum_{i=1}^p \gamma_{k,i}$

Thus, to guarantee boundedness of $\sigma(v_k)^2$ we need to show that the "steady-state" values $\gamma_{k,i} \rightarrow (\gamma_i < \infty)$.

We showed that

$$\gamma_i = \frac{\mu(\alpha + \sigma_\epsilon^2)}{2 \times (1 - \mu\lambda_i)} \quad (5.20)$$

where $\sigma_\epsilon^2 = E[\epsilon_k^2]$, λ_i is the i^{th} eigenvalue of R_{xx} ($R_{\text{xx}} = U \begin{pmatrix} \lambda_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \lambda_p \end{pmatrix} U^T$), and $\alpha = \frac{c\sigma_\epsilon^2}{1-c}$.

$$0 < c = \frac{1}{2} \sum_{i=1}^p \frac{\mu\lambda_i}{1 - \mu\lambda_i} < 1 \quad (5.21)$$

We found a sufficient condition for μ that guaranteed that the steady-state γ_i 's (and hence $\sigma(v_k)^2$) are bounded:

$$\mu < \frac{\frac{2}{3}}{\sum_{i=1}^p \lambda_i}$$

Where $\sum_{i=1}^p \lambda_i = \text{tr}(R_{\text{xx}})$ is the input vector energy.

With this choice of μ we have:

1. convergence in mean
2. bounded steady-state variance

This implies

$$\lim_{B \rightarrow \infty} \max \{Pr\{\|v_k\| \geq B\}\} = 0 \quad (5.22)$$

In other words, the LMS algorithm is stable about the optimum weight vector w_{opt} .

5.2.7 Learning Curve

Recall that

$$e_k = y_k - x_k^T w_{k-1} \quad (5.23)$$

and (5.14). These imply

$$e_k = \epsilon_k - x_k^T v_{k-1} \quad (5.24)$$

where $v_k = w_k - w_{\text{opt}}$. So the MSE

$$\begin{aligned} E[e_k^2] &= \sigma_\epsilon^2 + E[v_{k-1}^T x_k x_k^T v_{k-1}] \\ &= \sigma_\epsilon^2 + E[E[v_{k-1}^T x_k x_k^T v_{k-1} | x_n \epsilon_n, n < k]] \\ &= \sigma_\epsilon^2 + E[v_{k-1}^T R_{\text{xx}} v_{k-1}] \\ &= \sigma_\epsilon^2 + E[\text{tr}(R_{\text{xx}} v_{k-1} v_{k-1}^T)] \\ &= \sigma_\epsilon^2 + \text{tr}(R_{\text{xx}} \Gamma_{k-1}) \end{aligned} \quad (5.25)$$

Where $(\text{tr}(R_{\text{xx}} \Gamma_{k-1}) \equiv \alpha_{k-1}) \rightarrow (\alpha = \frac{c\sigma_\epsilon^2}{1-c})$. So the limiting MSE is

$$\begin{aligned} \epsilon_\infty &= \lim_{k \rightarrow \infty} E[e_k^2] \\ &= \sigma_\epsilon^2 + \frac{c\sigma_\epsilon^2}{1-c} \\ &= \frac{\sigma_\epsilon^2}{1-c} \end{aligned} \quad (5.26)$$

Since $0 < c < 1$ was required for convergence, $\epsilon_\infty > \sigma_\epsilon^2$ so that we see noisy adaptation leads to an MSE larger than the optimal

$$\begin{aligned} E[\epsilon_k^2] &= E\left[(y_k - x_k^T w_{\text{opt}})^2\right] \\ &= \sigma_\epsilon^2 \end{aligned} \quad (5.27)$$

To quantify the increase in the MSE, define the so-called **misadjustment**:

$$\begin{aligned} M &= \frac{\epsilon_\infty - \sigma_\epsilon^2}{\sigma_\epsilon^2} \\ &= \frac{\epsilon_\infty}{\sigma_\epsilon^2} - 1 \\ &= \frac{c}{\sigma_\epsilon^2} \\ &= \frac{c}{1-c} \end{aligned} \quad (5.28)$$

We would of course like to keep M as small as possible.

5.2.8 Learning Speed and Misadjustment Trade-off

Fast adaptation and quick convergence require that we take steps as large as possible. In other words, learning speed is proportional to μ ; larger μ means faster convergence. How does μ affect the misadjustment?

To guarantee convergence/stability we require

$$\mu < \frac{\frac{2}{3}}{\sum_{i=1}^p \lambda_i (R_{xx})}$$

Let's assume that in fact $\mu \ll \frac{1}{\sum_{i=1}^p \lambda_i}$ so that there is no problem with convergence. This condition implies $\mu \ll \frac{1}{\lambda_i}$ or $\mu \lambda_i \ll 1$, $i = \{1, \dots, p\}$. From here we see that

$$c = \frac{1}{2} \sum_{i=1}^p \frac{\mu \lambda_i}{1 - \mu \lambda_i} \simeq \frac{1}{2} \mu \sum_{i=1}^p \lambda_i \ll 1 \quad (5.29)$$

This misadjustment

$$M = \frac{c}{1-c} \simeq c = \frac{1}{2} \mu \sum_{i=1}^p \lambda_i \quad (5.30)$$

This shows that larger step size μ leads to larger misadjustment.

Since we still have convergence in mean, this essentially means that with a larger step size we "converge" faster but have a larger variance (rattling) about w_{opt} .

5.2.9 Summary

small μ implies

- small misadjustment in steady-state
- slow adaptation/tracking

large μ implies

- large misadjustment in steady-state
- fast adaptation/tracking

Example 5.1

$$w_{\text{opt}} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

$$x_k \sim \mathcal{N} \left(0, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

$$y_k = x_k^T w_{\text{opt}} + \epsilon_k$$

$$\epsilon_k \sim \mathcal{N}(0, 0.01)$$

5.2.1 LMS Algorithm

initialization $w_0 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and $w_k = w_{k-1} + \mu x_k e_k$, $k \geq 1$, where $e_k = y_k - x_k^T w_{k-1}$

Learning Curve

Image not finished

Figure 5.11: $\mu = 0.05$

LMS Learning Curve

Image not finished

Figure 5.12: $\mu = 0.3$

Comparison of Learning Curves

Image not finished

Figure 5.13

Glossary

B Bayesian statistical model

A statistical model compose of a data generation model, $p(x|\theta)$, and a **prior** distribution on the parameters, $p(\theta)$.

bias

The **bias** of θ is [Insert 3]

C conjugate prior

$p(\theta)$ is a **conjugate prior** for $p(x|\theta)$ if $p(\theta) \in \mathcal{P} \Rightarrow p(\theta|x) \in \mathcal{P}$

I iid

independent and identically distributed

L

Let X_1, \dots, X_M be a random sample, governed by the density or probability mass function $f(x|\theta)$. The statistic $T(x)$ is **sufficient** for θ if the conditional distribution of x , given $T(x) = t$, is independent of θ . Equivalently, the functional form of $f_{\theta|t}(x)$ does not involve θ .

M MVUE

An estimator θ is the **minimum variance unbiased estimator** if it is unbiased and has the smallest variance of any unbiased estimator for all values of the unknown parameter. In other words, the MVUE satisfies the following two properties: [Insert 6]

U unbiased

An estimator θ is called **unbiased** if its bias is zero for all values of the unknown parameter. Equivalently, [Insert 5]

V variance

The **variance** of θ is [Insert 2]

Bibliography

- [1] Casella and Berger. *Statistical Inference*. Duxbury Press, Belmont, CA, 1990.
- [2] Kay. *Estimation Theory*. Prentice Hall, 1993.
- [3] S. Kay. *Detection Theory*. Prentice Hall, 1998.
- [4] Steven Kay. *Fundamentals of Statistical Signal Processing Volume I: Estimation Theory*. Prentice Hall, 1993.
- [5] J. Neyman and E.S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. Roy. Soc. Ser. A*, 231:289–337, 1933.
- [6] L. Scharf. *Statistical Signal Processing*. Addison-Wesley, 1991.
- [7] L. Scharf. *Statistical Signal Processing*. Addison-Wesley, 1991.
- [8] P.M. Woodward. *Probability and Information Theory, with Applications to Radar*. Pergamon Press, Oxford, second edition edition, 1964.

Index of Keywords and Terms

Keywords are listed by the section with that keyword (page numbers are in parentheses). Keywords do not necessarily appear in the text of the page. They are merely associated with that section. *Ex.* apples, § 1.1 (1) **Terms** are referenced by the page they appear on. *Ex.* apples, 1

- ' '0-1' loss, § 4.3.1(96)
- A**
 - a priori, § 3.2.1(52)
 - absolute error loss, § 4.3.1(96)
 - adaptive, § 5.1(113), 113
 - adaptive controller, § 5.1(113)
 - alternative, § 3.1(47), 49
 - analog, § 2.1(33), 34
 - anticausal, § 2.1(33), 36
 - aperiodic, § 2.1(33)
 - asymptotically efficient, 93
 - asymptotically unbiased, 93
- B**
 - basis, § 1.1(7), 10
 - Bayes, § 3.2.1(52)
 - Bayes Rule, § 1.2(17), 18
 - Bayes', § 3.2.1(52)
 - Bayesian, § 4.3.2.2(102)
 - Bayesian Analysis, § 1.2(17)
 - Bayesian statistical model, § 1.2(17), 18
 - Bernoulli trials, § 1.3.2(30)
 - bias, § 4.1.1(77), 78, 78, 125
 - bias-variance decomposition, 78
 - bias-variance tradeoff, 78
 - binary, § 3.1(47), 49
- C**
 - Cauchy-Schwarz inequality, 85
 - causal, § 2.1(33), 36
 - channel equalization, § 5.1(113)
 - channel/system identification, § 5.1(113)
 - complement, § 1.1(7), 12
 - completeness, § 1.3.1(27), 30
 - composite, § 3.1(47), 50
 - conditionality principle., 91
 - conjugate, § 1.2(17)
 - conjugate prior, 24, 125
 - continuous, 33
 - continuous time, § 2.1(33)
 - correlation, § 4.3.2.2(102), 103
 - Cramer-Rao, § 4.1.2(81)
 - Cramer-Rao Lower Bound, § 2.2.1(40), 81
 - criteria, § 3.2.1(52)
 - CRLB, § 2.2.1(40), § 4.1.2(81)
- D**
 - decision boundary, § 3.1(47), 50
 - decision region, § 3.2.1(52)
 - decision regions, § 3.1(47), 50
 - detection, § (1), 3, § 3.2.1(52)
 - detection probability, § 3.1(47), 51, § 3.2.1(52)
 - deterministic signal, 39
 - digital, § 2.1(33), 34
 - direct sum, § 1.1(7), 12
 - discrete, 33
 - discrete time, § 2.1(33)
 - downhill, § 5.1(113), 119
 - driving, 104
 - driving noise, § 4.3.2.2(102)
 - dynamical, 104
 - dynamical model, § 4.3.2.2(102)
- E**
 - efficiency, § 4.1.2(81)
 - efficient, 87
 - error probability, § 3.2.1(52)
 - estimation, § (1), 3, § 4.1.1(77)
 - estimator accuracy, § 4.1.2(81)
 - even signal, § 2.1(33), 37
 - excitation, 104
 - excitation noise, § 4.3.2.2(102)
 - Expectation-Maximization Algorithm, 93
 - expected loss, § 4.3.1(96)
- F**
 - false alarm, § 3.1(47), 51
 - false-alarm probability, § 3.1(47), 51, § 3.2.1(52)
 - filter, § 4.3.2.2(102)
 - finite, 10
 - finite dimensional, § 1.1(7), 11
 - finite-length signal, 35
 - first-order Gauss-Markov, § 4.3.2.2(102), 104
 - Fisher Information, 86, § 4.2(90)
 - Fisher Information Matrix, 89, 93
 - fundamental period, 34
 - Fundamental Theorem of Calculus, 84

- G** Guassian white noise, 79
- H** Hilbert Space, § 1.1(7), 15
 hypotheses, § 3.1(47), 47
 hypothesis test, § 3.1(47), 47
 hypothesis testing, § 3.2.1(52)
- I** idd, 87
 infinite-length signal, 35
 informative, § 1.2(17)
 invariance, § 4.2(90)
 inversion, § 1.2(17), 17
 invertible, § 1.1(7), 17
 iterative minimization, § 5.1(113)
- J** joint, § 1.2(17)
- K** Kalman, § 4.3.2.2(102)
 Kalman filter, 110
 Kalman gain, § 4.3.2.2(102), 110
 Kalman recursions, § 4.3.2.2(102), 110
- L** Lagrange multiplier, § 3.2.1(52)
 learning curve, § 5.2(120)
 learning speed, § 5.2(120)
 least-mean-square, § 5.1(113), 113
 Lebesgue measure, 84
 likelihood function, § 4.2(90), 90
 likelihood principle, § 4.2(90)
 likelihood ratio, § 3.2.1(52), § 3.2.2(57), 58,
 § 3.2.3(59), 60, § 3.2.4.1(68), 71
 likelihood ratio test, § 3.2.2(57), § 3.2.4.1(68),
 71
 linear model, § 2.2.1(40), 40
 linear transformation, § 1.1(7), 15
 linearly dependent, § 1.1(7), 9
 linearly independent, § 1.1(7), 9
 loss function, § 4.3.1(96), 96
 lower bound, § 4.1.2(81)
 lowpass, § 4.3.2.1(99), 100
 LRT, § 3.2.2(57)
- M** M-ary, § 3.1(47), 49
 MAP, 65, § 4.3.1(96)
 MAP rule, § 3.2.3(59)
 marginal, § 1.2(17)
 matrix representation, § 1.1(7), 16
 maximum a posteriori, § 3.2.3(59), 65, 99
 Maximum Likelihood Estimation, § 4.2(90)
 maximum likelihood estimator, 90, 92
 Maximum Likelihood Estimators, § 4.2(90)
 maximum probability correct, § 3.2.1(52)
 mean-squared error, 77
 minimal, § 1.3.1(27), 30
 minimum error probability, § 3.2.1(52)
 minimum variance, § 4.1.1(77)
 Minimum Variance Unbiased Estimator,
 § 2.2.1(40), 125
 minimum variance unbiased estimator
 (MVUE), 81
 misadjustment, § 5.2(120), 123
 miss, § 3.1(47), 51, § 3.2.4.1(68), 68
 miss probability, § 3.1(47), 51, § 3.2.1(52)
 MLE, § 4.2(90)
 monotonicity, § 3.2.1(52)
 MSE, § 4.1.1(77)
 MVUB, § 2.2.1(40), § 4.1.2(81)
 MVUE, 81
- N** Newton-Raphson, 93
 Neyman-Pearson, § 3.2.1(52)
 Neyman-Pearson criterion, § 3.2.1(52)
 Neyman-Pearson lemma, § 3.2.4.1(68), 70
 noise cancellation, § 5.1(113)
 non-informative, § 1.2(17)
 noncausal, § 2.1(33), 36
 null, § 3.1(47), 49
- O** observation, 40
 observation vector, 40
 odd signal, § 2.1(33), 37
 on-line, § 5.1(113), 116
 one-sided, § 3.1(47), 51
 optimization, § 3.2.1(52)
 orthogonal, § 1.1(7), 14, 17
 orthogonal compliment, § 1.1(7), 15
 orthogonality principle, 108, 119
 orthonormal, § 1.1(7), 14
- P** parameter vector, 40
 period, 34
 periodic, § 2.1(33)
 phase-shift keying, § 3.1(47), 49
 posterior, § 1.2(17)
 posterior mean, 97
 posterior median, 98
 power, § 3.1(47), 51
 prior, § 1.2(17), § 1.2(17), 125, 18, 22
 prior distribution, § 1.2(17)
 prior distribution, 18
 probability, § (1), 2
 probability of error, § 3.2.1(52), 53
 Probability Theory, § (1), 2
 pseudorandom noise sequence, 45
- R** random signal, 39

- real-time, § 5.1(113), 116
- receiver operating characteristic, § 3.2.4.1(68), 70
- ROC, § 3.2.4.1(68)
- S** sample-mean estimator, 79
- scalar case, § 4.1.1(77)
- Scoring Method, 93
- signal-to-noise ratio, § 3.2.4.1(68), 73
- signals and systems, § 2.1(33)
- simple, § 3.1(47), 50
- size, § 3.1(47), 51
- smooth, § 4.3.2.1(99), 100
- span, § 1.1(7), 9
- standard basis, § 1.1(7), 16
- state, 104
- state model, § 4.3.2.2(102)
- state of nature, 2
- state transition matrix, 106
- statistic, § (1), 2
- statistical signal processing, § (1), 2
- statistics, § (1), 2
- steepest descent, § 5.1(113)
- subspace, § 1.1(7), 9
- sufficiency principle, 91
- sufficient, § 1.3.1(27), 125
- sufficient statistic, 28, § 1.3.2(30), § 3.2.1(52)
- system matrix, 40
- T** threshold, § 3.2.2(57), 58, § 3.2.3(59), 60, § 3.2.4.1(68), 71
- tradeoff, § 3.2.1(52)
- truncated, § 1.2(17), 19
- two-sided, § 3.1(47), 50
- type I, 51
- type I error, § 3.1(47), 53
- type II error, § 3.1(47), 53
- type II errors, 51
- U** unbiased, 79, 125
- unbiased estimators, § 4.1.1(77)
- uncertainties, § (1), 1
- unitary, § 1.1(7), 17
- unknown parameters, 2
- unknowns, § (1), 1
- V** Vandermonde matrix, 42
- variance, 78, 78, 125
- vector parameter, § 4.1.2(81)
- vectors, § 1.1(7), 8
- W** weight error vector, § 5.2(120)
- well-defined, § 1.1(7), 11

Attributions

Collection: *Statistical Signal Processing*

Edited by: Clayton Scott

URL: <http://cnx.org/content/col10232/1.1/>

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Introduction to Statistical Signal Processing"

By: Clayton Scott

URL: <http://cnx.org/content/m11947/1.3/>

Pages: 1-5

Copyright: Clayton Scott

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Review of Linear Algebra"

By: Clayton Scott

URL: <http://cnx.org/content/m11948/1.2/>

Pages: 7-17

Copyright: Clayton Scott

License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Bayesian Paradigm"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11444/1.7/>

Pages: 17-27

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Sufficient Statistics"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11481/1.7/>

Pages: 27-30

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Fisher-Neyman Factorization Theorem"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11480/1.6/>

Pages: 30-32

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Signal Classifications and Properties"

By: Melissa Selik, Richard Baraniuk, Michael Haag

URL: <http://cnx.org/content/m10057/2.21/>

Pages: 33-39

Copyright: Melissa Selik, Richard Baraniuk, Michael Haag

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Linear Models"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11432/1.2/>
Pages: 40-46
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "Hypothesis Testing"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11531/1.6/>
Pages: 47-52
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "Criteria in Hypothesis Testing"
By: Don Johnson
URL: <http://cnx.org/content/m11228/1.4/>
Pages: 52-57
Copyright: Don Johnson
License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Bayes Risk Criterion in Hypothesis Testing"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11533/1.6/>
Pages: 57-59
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "Minimum Probability of Error Decision Rule"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11534/1.11/>
Pages: 59-67
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Neyman-Pearson Criterion"
By: Clayton Scott
URL: <http://cnx.org/content/m11548/1.2/>
Pages: 68-76
Copyright: Clayton Scott
License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Minimum Variance Unbiased Estimator"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11426/1.6/>
Pages: 77-81
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "The Cramer-Rao Lower Bound"
By: Clayton Scott, Robert Nowak
URL: <http://cnx.org/content/m11429/1.4/>
Pages: 81-90
Copyright: Clayton Scott, Robert Nowak
License: <http://creativecommons.org/licenses/by/1.0>

Module: "Maximum Likelihood Estimation"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11446/1.5/>

Pages: 90-96

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Bayesian Estimation"

By: Clayton Scott

URL: <http://cnx.org/content/m11660/1.2/>

Pages: 96-99

Copyright: Clayton Scott

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Wiener Filtering and the DFT"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11462/1.5/>

Pages: 99-102

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Kalman Filters"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11438/1.7/>

Pages: 102-112

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "Introduction to Adaptive Filtering"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11433/1.5/>

Pages: 113-120

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Module: "LMS Algorithm Analysis"

By: Clayton Scott, Robert Nowak

URL: <http://cnx.org/content/m11439/1.4/>

Pages: 120-124

Copyright: Clayton Scott, Robert Nowak

License: <http://creativecommons.org/licenses/by/1.0>

Statistical Signal Processing

To learn Statistical Signal Processing.

About Connexions

Since 1999, Connexions has been pioneering a global system where anyone can create course materials and make them fully accessible and easily reusable free of charge. We are a Web-based authoring, teaching and learning environment open to anyone interested in education, including students, teachers, professors and lifelong learners. We connect ideas and facilitate educational communities.

Connexions's modular, interactive courses are in use worldwide by universities, community colleges, K-12 schools, distance learners, and lifelong learners. Connexions materials are in many languages, including English, Spanish, Chinese, Japanese, Italian, Vietnamese, French, Portuguese, and Thai. Connexions is part of an exciting new information distribution system that allows for **Print on Demand Books**. Connexions has partnered with innovative on-demand publisher QOOP to accelerate the delivery of printed course materials and textbooks into classrooms worldwide at lower prices than traditional academic publishers.