Detection of Abrupt Changes: Theory and Application[1]

Michèle Basseville

IRISA/CNRS
Rennes, France

Igor V. Nikiforov

Institute of Control Sciences
Moscow, Russia

[1]This book was previously published by Prentice-Hall, Inc.

Contents

Preface

Over the last twenty years, there has been a significant increase in the number of real problems concerned with questions such as

- fault detection and diagnosis (monitoring);
- condition-based maintenance of industrial processes;
- safety of complex systems (aircrafts, boats, rockets, nuclear power plants, chemical technological processes, etc.);
- quality control;
- prediction of natural catastrophic events (earthquakes, tsunami, etc.);
- monitoring in biomedicine.

These problems result from the increasing complexity of most technological processes, the availability of sophisticated sensors in both technological and natural worlds, and the existence of sophisticated information processing systems, which are widely used. Solutions to such problems are of crucial interest for safety, ecological, and economical reasons. And because of the availability of the above-mentioned information processing systems, complex monitoring algorithms can be considered and implemented.

The common feature of the above problems is the fact that the problem of interest is the detection of one or several *abrupt changes* in some characteristic properties of the considered object. The key difficulty is to detect intrinsic changes that are not necessarily directly observed and that are measured together with other types of perturbations. For example, it is of interest to know how and when the modal characteristics of a vibrating structure change, whereas the available measurements (e.g., accelerometers) contain a mix of information related to both the changes in the structure and the perturbations due to the environment.

Many monitoring problems can be stated as the problem of detecting a change in the *parameters* of a static or dynamic stochastic system. The main goal of this book is to describe a unified framework for the design and the performance analysis of the algorithms for solving these change detection problems. We call abrupt change any change in the parameters of the system that occurs either instantaneously or at least very fast with respect to the sampling period of the measurements. Abrupt changes by no means refer to changes with large magnitude; on the contrary, in most applications the main problem is to detect small changes. Moreover, in some applications, the *early warning* of small - and not necessarily fast - changes is of crucial interest in order to avoid the economic or even catastrophic consequences that can result from an accumulation of such small changes. For example, small faults arising in the sensors of a navigation system can result, through the underlying integration, in serious errors in the estimated position of the plane. Another example is the early warning of small deviations from the normal operating conditions of an industrial process. The early detection of slight changes in the state of the process allows to plan in a more adequate manner the periods during which the process should be inspected and possibly repaired, and thus to reduce the exploitation costs.

Our intended readers include engineers and researchers in the following fields :

- signal processing and pattern recognition;
- automatic control and supervision;
- time series analysis;
- applied statistics;
- quality control;
- condition-based maintenance and monitoring of plants.

We first introduce the reader to the basic ideas using a nonformal presentation in the simplest case. Then we have tried to include the key mathematical background necessary for the design and performance evaluation of change detection algorithms. This material is usually spread out in different types of books and journals. The main goal of chapters 3 and 4 is to collect this information in a single place. These two chapters should be considered not as a small textbook, but rather as short notes that can be useful for reading the subsequent developments.

At the end of each chapter, we have added a Notes and References section and a summary of the main results. We apologize for possible missing references.

We would like to acknowledge the readers of earlier versions of the book, for their patient reading and their numerous and useful comments. Thanks are due to Albert Benveniste, who read several successive versions, for numerous criticisms, helpful discussions and suggestions; to Mark Bodson, who reviewed the manuscript; to Fredrik Gustafsson, Eric Moulines, Alexander Novikov, David Siegmund, Shogo Tanaka and Qinghua Zhang, for their numerous comments; and to Alan Willsky for his comments regarding an early version of chapter 7.

Philippe Louarn provided us with extensive and valuable help in using LATEX; his endless patience and kindness in answering our questions undoubtedly helped us in making the manuscript look as it is. Bertrand Decouty helped us in using software systems for drawing pictures.

Our thanks also to Thomas Kailath who accepted the publication of this book in the Information and System Sciences Series which he is editing.

During the research and writing of this book, the authors have been supported by the Centre National de la Recherche Scientifique (CNRS) in France, the Institute of Control Sciences in Moscow, Russia, and the Institut National de la Recherche en Informatique et Automatique (INRIA) in Rennes, France.

The book was typeset by the authors using LATEX. The figures were drawn using MATLAB and XFIG under the UNIX operating system. Part of the simulations were developed using the package AURORA designed at the Institute of Control Sciences in Moscow, Russia.

Michèle Basseville
Igor Nikiforov
Rennes, France

List of Figures

List of Tables

Notation and Symbols

Some symbols may have locally other meanings.

Basic Notation

Notation	**Meaning**
$\mathbf{1}_{\{x\}}$	Indicator function of event x.
$\nVdash_r$	Vector of size r made of 1.
$\{x : \ldots\}$	Set of x such that.
$\mathbf{R}$	Set of real numbers.
(a, b)	Open interval.
$[a, b]$	Closed interval.
$\max_k, \min_k$	Extrema with respect to an integer value.
$\sup_x, \inf_x$	Extrema with respect to a real value.
$X \sim Y$	Same order of magnitude.
$X \approx Y$	Approximately equal.
$\Vert v \Vert_A^2 = v^T A v$	Quadratic form with respect to matrix A.
y	Observation (random variable) of dimension $r = 1$.
Y	Observation (random variable) of dimension $r > 1$.
$\mathcal{Y}_1^k = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_k \end{pmatrix}$	or $\mathcal{Y}_1^k = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{pmatrix}$
$\check{\mathcal{Y}}_1^k = \begin{pmatrix} Y_k \\ Y_{k-1} \\ \vdots \\ Y_1 \end{pmatrix}$	or $\check{\mathcal{Y}}_1^k = \begin{pmatrix} y_k \\ y_{k-1} \\ \vdots \\ y_1 \end{pmatrix}$
k	Current time instant - discrete time.
t	Current time instant - continuous time.
N	Sample size.
γ_i	Weighting coefficients.
$\dot{g}(x)$	First derivative of the function $g(x)$.
$\ddot{g}(x)$	Second derivative of the function $g(x)$.

$\mathrm{tr}\, A$	Trace of matrix A.
$\det A$	Determinant of matrix A.
$\ker A$	Kernel of matrix A.

Notation for Probability

Notation	**Meaning**
Ω	Sample (abstract) space.
$\mathcal{B}$	Event space, sigma algebra of subsets of Ω.
μ	Probability measure.
$\mathbf{P}(B)$	Probability of the event $B \in \mathcal{B}$.
$\mathbf{E}$	Expectation.
Y	Random variable, scalar, or vector.
y	Argument of the distribution functions.
$F(y)$	Cumulative distribution function, cdf.
θ	Vector of parameters with dimension ℓ.
$(\vartheta_i)_{1\leq i\leq \ell}$	Coordinates of the vector parameter θ.
θ	Parameter.
Υ	Change direction.
$\mathcal{P} = \{\mathbf{P}_\theta\}_{\theta\in\Theta}$	Parametric family of probability distributions.
$p(y), f(y)$	Probability density, pdf.
$p_\theta(y), f_\theta(y)$	Parameterized probability density, pdf.
$\mathbf{N}(\theta)$	Entropy of the distribution p.
$\mathcal{L}$	Probability law; $\mathcal{L}(Y) = \mathbf{P}_\theta$.
$\mathcal{N}$	Normal (or Gaussian) law.
$\varphi(y)$	Gaussian density, with $\mu = 0, \sigma = 1$.
$\phi(y)$	Corresponding Gaussian cdf.
$\Phi_{\mu,\Sigma}(B)$	$= \int_B f(y)dy$, where $\mathcal{L}(X) = \mathcal{N}(\mu, \Sigma)$.
$\mathcal{K}$	Functional used in stochastic approximation algorithms.
κ	Mean value of $\mathcal{K}$.
ψ	Laplace transform or moment generating function.
ς	Laplace variable.
σ^2	Variance of the scalar random variable Y.
R, Σ	Covariance matrix of the vector random variable Y.
Φ_Y	Power spectrum of a process Y.
$\mathbf{T}_p(Y)$	Toeplitz matrix of order p for a process Y.
$\mathrm{var}(Y)$	Variance of the random variable Y.
$\mathrm{cov}(Y)$	Covariance matrix of the vector random variable Y.
$a_1, ..., a_p$	Scalar AR coefficients.

$A_1, ..., A_p$	Matrix AR coefficients.
$b_1, ..., b_q$	Scalar MA coefficients.
$B_1, ..., B_q$	Matrix MA coefficients.
$T_{-\epsilon,h}$	Exit time from the interval $(-\epsilon, h)$.
$\mathbf{P}_\theta(-\epsilon\|z)$	Abbreviation for $\mathbf{P}_\theta(S_{T_{-\epsilon,h}} \leq -\epsilon\|S_0 = z)$.
$\mathbf{E}_\theta(-\epsilon\|z)$	Abbreviation for $\mathbf{E}_\theta(T_{-\epsilon,h}\|S_0 = z)$.
w.p.1.	Abbreviation for "with probability 1."

Notation for Statistics

Notation	**Meaning**
$\mathbf{H}$	Hypothesis.
α	Error probability.
β	Power.
K_ϵ	Class of tests with level $1-\epsilon$.
κ_ϵ	Quantile of the normal distribution corresponding to level $1-\epsilon$.
$\mathbf{I}$	Fisher information.
$\mathbf{K}$	Kullback information.
$\mathbf{J}$	Kullback divergence (symmeterized information).
$b(\theta)$	Bias of an estimate $T(y)$.
$S_i^j(\theta_0, \theta_1)$	Log-likelihood ratio for observations from y_i until y_j.
$S_i^j = S_i^j(\theta_0, \theta_1)$	
$S_j = S_1^j$	
$\tilde{S}_i^j$	Weighted log-likelihood ratio.
s_i	Increment of S_i^j.
Λ_i^j	Likelihood ratio for observations from y_i until y_j.
$\tilde{\Lambda}_i^j$	Weighted likelihood ratio.
$l_\theta(y) = \ln f_\theta(y)$	Log-likelihood function.
$z = \frac{\partial l_\theta(y)}{\partial \theta}$	Efficient score for a random variable with scalar parameter.
$Z = \frac{\partial l_\theta(y)}{\partial \theta}$	Efficient score for a random variable with vector parameter.
$\mathcal{Z}_k^N = \frac{\partial l_\theta(\mathcal{Y}_k^N)}{\partial \theta}$	Efficient score for a sample $Y_k, \ldots, Y_N$ of a random process with scalar or vector parameter.
$\Delta_N = \frac{1}{\sqrt{N}} \mathcal{Z}_N$	
$Z^* = \left. \frac{\partial l_\theta(y)}{\partial \theta} \right\|_{\theta=\theta^*}$	
g	Test statistics.
T	Stopping time.
$N(z)$	ASN function, as a function of the initial value z of the cumulative sum.
η	Shifted log-likelihood ratio.
$\mathcal{S}$	Cumulative sum of η.

Notation for Change Detection (Except in State-Space Models)

Notation	**Meaning**
θ_0	Parameter before change.
θ_1	Parameter after change.
ν	Magnitude of change.
Υ	Direction of change in a vector parameter.
ν_m	Minimum magnitude of change.
t_0	Change time.
t_a	Alarm time.
g_k	Decision function.
$-a, \lambda, -\epsilon, h$	Thresholds.
N_k	Number of observations since last vanishing of g_k.
$\bar{T}$	Mean time between false alarms.
$\bar{\tau}$	Mean delay for detection.
$L_z(\theta)$	ARL function.

Notation for State-Space Models

Notation	**Meaning**
y, Y	Observation, of dimension $r \geq 1$.
X	State, of dimension n.
U	Input (control) of dimension m.
V, W	Noises on observation and state equations, respectively.
R, Q	Covariance matrices for V, W, respectively.
F	State transition matrix.
G, J	Control matrices.
H	Observation matrix.
$\mathcal{H}$	Hankel matrix.
$\mathcal{J}(G, J)$	Toeplitz impulse response.
$\mathcal{O}$	Observability matrix.
$\mathcal{C}$	Controllability matrix.
$\mathcal{T}$	Transfer function.
K	Kalman gain.
P	Covariance matrix of state estimation error.
ε_k	Innovation.
Σ	Covariance matrix of innovation ε.
e_k	Residual.
ζ_k	Parity vector.
$\mathcal{S}$	Parity space.
Υ_x	Direction of the additive change in X.
Υ_y	Direction of the additive change in Y.
Ψ_x	Vector made of successive Υ_x.
Ψ_y	Vector made of successive Υ_y.
Γ	"Gain" matrix for Υ_x.
Ξ	"Gain" matrix for Υ_y.
ρ	Signature of additive change on the innovation.
ϱ	Signature of additive change on the parity vector.
$\mathcal{K}_x$	Transfer function for the contribution of Υ_x to ρ.
$\mathcal{K}_y$	Transfer function for the contribution of Υ_y to ρ.
$\mathcal{H}_x$	Transfer function for the contribution of Υ_x to ϱ.
$\mathcal{H}_y$	Transfer function for the contribution of Υ_y to ϱ.

1

Introduction

In this chapter, we describe the purpose and contents of the book. In section 1.1 we give the theoretical and applied motivations for change detection. The last part of this section consists of three possible statistical problem statements for change detection, together with the intuitive definition of the corresponding criteria to be used for the design and performance analysis of change detection techniques. The formal definition of these criteria is given at the end of chapter 4, after the introduction of the key mathematical tools to be used throughout the book.

In section 1.2, we introduce five typical application examples, which we will use to introduce the main techniques. In section 1.3, we describe the organization of the book, based on a classification of change detection problems according to the types of characteristics that change. We give a short description of each chapter and a general flowchart of the chapters. Finally, in section 1.4, we comment further on several critical issues concerning the design of change detection algorithms and the investigation of their properties.

1.1 Introducing Change Detection

In this section, we introduce abrupt changes for segmentation, fault detection, and monitoring. We describe the main motivations for the investigation of change detection problems. Illustrating examples are described in the next section. Then we classify the topics of change detection methodology into three main classes of problems encountered in signal processing, time series analysis, automatic control, and industrial quality control. Next, we give three statistical problem statements and the intuitive definition of the corresponding criteria. Finally, we describe the purpose of the book.

1.1.1 Motivations

An intensively investigated topic is time series analysis and identification. The main assumptions underlying these investigations are that the properties or parameters describing the data are either constant or slowly time-varying. On the other hand, many practical problems arising in quality control, recognition-oriented signal processing, and fault detection and monitoring in industrial plants, can be modeled with the aid of parametric models in which the parameters are subject to *abrupt changes at unknown time instants*. By abrupt changes, we mean changes in characteristics that occur very fast with respect to the sampling period of the measurements, if not instantaneously. Because a large part of the information contained in the measurements lies in their nonstationarities, and because most of adaptive estimation algorithms basically can follow only slow changes, the detection of abrupt changes is a problem of interest in many applications, as we show in the five examples of section 1.2. The *detection of abrupt changes* refers to tools that help us decide whether such a change occurred in the characteristics of the considered object.

The first meaning of abrupt change thus refers to a time instant at which properties suddenly change, but before and after which properties are constant in some sense, e.g., stationary. This notion serves as a basis to the corresponding formal mathematical problem statement, and to the formal derivation of algorithms for change detection.

It should now be clear that abrupt changes by no means imply changes with large magnitude. Many change detection problems are concerned with the detection of small changes, as we discuss now.

1.1.1.1 Fault Detection in Controlled Dynamic Systems

The problem of fault detection for monitoring an industrial process involves two types of questions. First, of course, the detection of failures or catastrophic events should be achieved. But second, and often of crucial practical interest, the detection of smaller faults - namely of sudden or gradual (incipient) modifications, which affect the process without causing it to stop - is also required to prevent the subsequent occurrence of more catastrophic events. As we explain in examples 1.2.2 and 1.2.5, both faults and failures can be approached in the abrupt change detection framework, with all the aspects of detection, estimation, and diagnosis usually implied in most failure detection and isolation (FDI) systems. Such a detection tool helps to increase the reliability and availability of the industrial process by reducing the number of shutdowns that are necessary for systematic maintenance. Usually, one has to distinguish between instruments and process faults. The detection of these two types of faults do not involve the same degree of difficulty. Instruments faults can often be modeled by an *additive* change in a state-space model, whereas process faults are more often *nonadditive* changes in the state of such models (see section 1.3).

In this situation, fast detection is often of crucial importance, for the reconfiguration of the control law, for example. Two uses of the change detection methodology in this framework are of interest. The first is related to the automatic processing of individual signals, as we discuss in the next paragraph. The second is more involved, from the point of view of the modeling information that is required. If the detection of the process faults is desired, and not only that of the instrument faults, or if isolation information is desired, then a partial knowledge of the physical model of the process is required to achieve the *diagnosis* of the fault in terms of its location in the process and physical cause. Both geometrical tools from system theory and statistical tools for change detection are used in these situations. We refer to example 1.2.5 for further discussion of these issues.

1.1.1.2 Segmentation of Signals

Now we discuss another important practical motivation for change detection. In recognition-oriented signal processing, the *segmentation of signals* refers to the automatic decomposition of a given signal into stationary, or weakly nonstationary, segments, the length of which is adapted to the local properties of the signal. As we show in examples 1.2.3 and 1.2.4, the change detection methodology provides preferential tools for such an automatic segmentation, which can be achieved either on-line or off-line. In this situation, the problems of interest are the detection of the changes in the local characteristics, and the estimation of the places, in time or space, where the changes occur. False alarms are relatively less crucial than in the previous case of fault detection, because they can be dealt with at the next stage of the recognition system. For example, in continuous speech processing, these algorithms can be used for detecting true abrupt changes. However, in practice these algorithms also give relevant results in less simple situations, for more slow transitions between pieces of signal where the properties of the signal are in fact slowly time-varying before and after the abrupt change [André-Obrecht, 1988]. The same is true in biomedical and seismic signal processing, where several segmentation algorithms have been derived and used for detecting onsets of spikes in EEG signals or P-waves in ECG signals or S-waves and P-waves in seismic data. We refer to examples 1.2.3 and 1.2.4

for additional comments on this point. Finally, let us emphasize that this context of signal segmentation is also valid in the framework of monitoring of industrial processes, where the analysis of individual sensors or actuators signals, without using any information about the model of the whole system, can bring key information for monitoring and fault detection.

1.1.1.3 Gain Updating in Adaptive Algorithms

Adaptive identification algorithms basically can track only slow fluctuations of characteristic parameters. For improving their tracking performances when quick fluctuations of the parameters occur, a possible solution consists of detecting abrupt changes in the characteristics of the analyzed system. The estimation of the change time and magnitude basically allows a more accurate updating of the gains of the identification algorithm. Such an approach has proved useful for tracking maneuvering targets, for example, as in [Willsky, 1976, Favier and Smolders, 1984].

1.1.1.4 Summary

The motivations leading to the change detection framework and methodology can be summarized as follows :

- From the theoretical point of view, it allows us to process abrupt changes and thus it is a natural counterpart to the *adaptive* framework and state of the art which basically can deal only with slow changes; and it is *one* way to approach the analysis of *nonstationary* phenomena.
- From the practical point of view, statistical decision tools for detecting and estimating changes are of great potential interest in three types of problems :
  1. quality control and fault detection in measurement systems and industrial processes in view of improved performances and condition-based maintenance;
  2. automatic segmentation of signals as a first step in recognition-oriented signal processing;
  3. gain updating in adaptive identification algorithms for improving their tracking ability.

The implementation of change detection techniques in these three types of situations is generally achieved with the aid of different philosophy and constraints (e.g., by choosing different models and criteria and by tuning of the parameters of the detectors), but basically the same methodology and tools apply in all situations.

1.1.2 Problem Statements and Criteria

We now discuss change detection problems from the point of view of mathematical statistics. We describe several problem statements and give the intuitive definitions of the corresponding criteria. Statistical change detection problems can be classified into three main classes for several reasons. The first lies in the theoretical definition of criteria used for deriving the algorithms; the second motivation comes from practical experience with different types of problems; and the last reason is historical, as sketched below. Recalling that the formal definition of criteria is given at the end of chapter 4, we now describe these three classes of problems.

1.1.2.1 On-line Detection of a Change

A preliminary statement for this question can be formulated as follows. Let $(y_k)_{1 \leq k \leq n}$ be a sequence of observed random variables with conditional density $p_\theta(y_k | y_{k-1}, \ldots, y_1)$. Before the unknown change time t_0, the conditional density parameter θ is constant and equal to θ_0. After the change, the parameter is equal to θ_1. The on-line problem is to detect the occurrence of the change as soon as possible, with a fixed rate of false alarms before t_0. The estimation of the change time t_0 is not required. The estimation of the parameters θ_0 and θ_1 is not required, but sometimes can be partly used in the detection algorithm. We implicitly assume that, in case of multiple change times, each change is detected quickly enough, one after the other, such that at each time instant only one change has to be considered. In the on-line framework, the detection is performed by a *stopping rule*, which usually has the form

$$t_a = \inf\{n : g_n(y_1, \ldots, y_n) \geq \lambda\} \tag{1.1.1}$$

where λ is a threshold, and $(g_n)_{n\geq 1}$ is a family of functions of n coordinates. Simple stopping rules are presented in chapter 2. The *alarm time* t_a is the time at which the change is detected. Note that if $t_a = n$, it is sufficient to observe the sample up to time n, which explains the name of "sequential" or on-line point of view.

In the on-line framework, the criteria are the *delay for detection*, which is related to the ability of the algorithm to set an alarm when a change actually occurs, and the *mean time between false alarms* (see section 4.4). Usually the overall criterion consists in *minimizing the delay for detection for a fixed mean time between false alarms*. We explain in section 4.4 that, from the mathematical point of view, two different definitions of the delay can be stated, which give rise to different results and proofs of optimality for various change detection algorithms.

1.1.2.2 Off-line Hypotheses Testing

We now consider the hypotheses "without change" and "with change." This off-line hypotheses testing problem can be formally stated as follows. Given a finite sample $y_1, \ldots, y_N$, test between

$$\begin{aligned}
\mathbf{H}_0 \quad &: \quad \text{for } 1 \leq k \leq N \ : \ p_\theta(y_k|y_{k-1}, \ldots, y_1) = p_{\theta_0}(y_k|y_{k-1}, \ldots, y_1) \\
\mathbf{H}_1 \quad &: \quad \text{there exists an unknown } 1 \leq t_0 \leq N \text{ such that:} \\
\text{for } 1 \leq k \leq t_0 - 1 \quad &: \quad p_\theta(y_k|y_{k-1}, \ldots, y_1) = p_{\theta_0'}(y_k|y_{k-1}, \ldots, y_1) \\
\text{for } t_0 \leq k \leq N \quad &: \quad p_\theta(y_k|y_{k-1}, \ldots, y_1) = p_{\theta_1}(y_k|y_{k-1}, \ldots, y_1)
\end{aligned} \tag{1.1.2}$$

In this problem statement, the estimation of the change time t_0 is not required.

As we explain in section 4.1, the usual criterion used in hypothesis testing is a tradeoff between the ability to detect actual changes when they occur, which requires a great sensitivity to high-frequency effects, and the ability not to detect anything when no change occurs, which requires a low sensitivity to noise effects. These are obviously two contradictory requirements. The standard criterion is usually to *maximize the probability of deciding $\mathbf{H}_1$ when $\mathbf{H}_1$ is actually true (i.e., the power), subject to the constraint of a fixed probability of deciding $\mathbf{H}_1$ when $\mathbf{H}_0$ is actually true (i.e., the size or probability of false alarms)*. In section 4.4, we explain why this criterion is especially difficult to use in the statistical change detection framework.

1.1.2.3 Off-line Estimation of the Change Time

In this problem statement, we consider the same hypotheses as before, and we assume that a change does take place in the sample of observations. Let $(y_k)_{1\leq k\leq N}$ be this sequence of random observations with conditional density $p_\theta(y_k|y_{k-1}, \ldots, y_1)$. Before the unknown change time t_0, which is assumed to be such that $1 \leq t_0 \leq N$, the parameter θ of the conditional density is constant and equal to θ_0. After the change, the parameter is equal to θ_1. The unknown change time has to be estimated from the observations $y_1, \ldots, y_N$ ($1 \leq N < \infty$) with maximum accuracy. The estimation of t_0 can possibly use information about θ_0 and θ_1, the availability of which depends upon situations. In such an estimation problem, we *intentionally leave out the search for multiple change times* between 1 and N.

The problem is to estimate t_0. This problem is a typical estimation problem for a discrete parameter. Obviously, this estimate has to be as accurate as possible. Usually, this accuracy is estimated by the probability that the estimate belongs to a given confidence interval, or by the first two moments of the probability distribution of the estimation error (bias and standard deviation).

Other types of criteria for deriving change estimation algorithms are discussed in [Bojdecki and Hosza, 1984, Pelkowitz, 1987].

1.1.2.4 Summary

In some practical applications all three types of problems may have to be solved together. We also emphasize here that an off-line point of view may be useful to design a decision and/or estimation function that is finally implemented on-line. We discuss this in section 1.4.

The *five intuitive performance indexes* for designing and evaluating change detection algorithms are the following :

1. *mean time between false alarms*;
2. *probability of false detection*;
3. *mean delay for detection*;
4. *probability of nondetection*;
5. *accuracy of the change* time and magnitude *estimates*.

We use these five indexes throughout the book. Another useful index consists of the Kullback information between the distributions before and after change. This distance does have a strong influence on the above-mentioned performance indexes, and we use it as a weak performance index when discussing *detectability* issues.

Another property of change detection algorithms is of great practical importance, and that is the *robustness*. Algorithms that are robust with respect to noise conditions and to modeling errors, and that are easy to tune on a new signal, are obviously preferred in practice. These robustness features cannot easily be formally stated, but should definitely be kept in mind when designing and experiencing change detection algorithms. This issue is discussed in several places in this book.

1.1.3 Purpose of the Book

This book is basically devoted to the design and investigation of *on-line change detection algorithms*. The off-line problem statement is discussed much more briefly, and mainly with a view to the discussion of some applications.

When designing change detection and estimation algorithms, it may be useful to distinguish two types of tasks :

1. *Generation of "residuals"* : These artificial measurements are designed to reflect possible changes of interest in the analyzed signal or system. They are, for example, ideally close to zero when no change occurs and significantly different from zero after the change. This is the case of the so-called *parity checks*, designed with the aid of the analytical redundancy approach. In other cases, the mean value or the spectral properties of these residuals may change when the analyzed system is changing. From the mathematical statistics point of view, a convenient way for generating these artificial measurements consists of deriving *sufficient statistics*.

2. *Design of decision rules* based upon these residuals : This task consists of designing the convenient decision rule which solves the change detection problem as reflected by the residuals.

In this book, we mainly focus on *parametric statistical tools* for detecting abrupt changes in properties of *discrete time* signals and dynamic systems. We intend to present didactically generalizations of points of view for designing algorithms together with new results, both theoretical and experimental, about their performances. The starting point is elementary well-known detectors used in industrial applications. We then generalize this approach in two tasks to more complex situations in which spectral properties of signals or dynamic properties of systems change. This book is intended to be a bridge between mathematical statistics tools and applied problems. Therefore, we do not derive all mathematical statistics theories, and readers who want complete mathematical results and proofs must consult other books or papers, indicated in references.

Even though great emphasis is placed on task 2, we also address the problem of deriving solutions for task 1. Deterministic solutions, such as in the analytical redundancy approach, are often based on geometrical properties of dynamic systems, as discussed further in section 1.4 and later in chapter 7. Mathematical statistics solutions, such as sufficient statistics or the so-called local approach, are further described in the following chapters, especially chapter 8.

1.2 Application Examples

In this section, we describe five typical application examples of change detection techniques. For each example, we give a short description of the particular problem and its context, including the main references. For some of these models, the detailed information about the possibly complex underlying physical models is given in chapter 11. This selection of examples is not exclusive; it is intended to give only sufficient initial insights into the variety of problems that can be solved within this framework, and to serve as much as possible as a common basis for all the algorithmic equipment presented in the subsequent chapters. In chapter 11, we come back to application problems, showing results of processing real signals with the aid of change detection algorithms, and discussing several potential application domains.

In the present chapter, the five examples are ranged according to the increasing complexity of the underlying change detection problems. We start with quality control and condition monitoring of inertial navigation systems (examples 1.2.1 and 1.2.2). Then we describe seismic signal processing and segmentation of signals (examples 1.2.3 and 1.2.4). Finally, we discuss failure detection in mechanical systems subject to vibrations (example 1.2.5).

1.2.1 Quality Control

One of the earliest applications of change detection is the problem of quality control, or continuous production monitoring. On-line quality control procedures are used when decisions are to be reached sequentially, as when measurements are taken. Consider a production process that can be *in control* and *out of control*. Situations where this process leaves the in control condition and enters the out of control state are called *disorders*. For many reasons, it is necessary to detect the disorder and estimate its time of occurrence. It may be a question of safety of the technological process, quality of the production, or classification of output items of production. For all these problems, the best solution is *quickest detection of the disorder with as few false alarms as possible*. This criterion is used because the delay for detection is a period of time during which the technological process is out of control without action of the monitoring system. From both safety and quality points of view, this situation is obviously highly nondesirable. On the other hand, frequent false alarms are inconvenient because of the cost of stopping the production and searching for the origin of the defect; nor is this situation desirable from psychological point of view, because the operator will very quickly stop using the monitoring system. Nevertheless, the optimal solution, according to the above-mentioned criterion, is basically a *tradeoff* between quick detection and few false alarms, using a comparison between the losses implied by the two events.

We stress here that we solve this problem using a *statistical approach*. From this point of view, the samples of measurements are a realization of a random process. Because of this random behavior, large fluctuations can occur in the measurements even when the process is in control, and these fluctuations result in false alarms. On the other hand, a given decision rule cannot detect the change instantaneously, again because of the random fluctuations in the measurements. When the technological process is in control, the measurements have a specific probability distribution. When the process is out of control, this distribution changes. If a parametric approach is used, we speak about changes in the parameters of this probability distribution. For example, let us consider a chemical plant where the quality of the output material is characterized by the concentration of some chemical component. We assume that this concentration is normally distributed. Under normal operating conditions, the mean value and standard deviation of this normal distribution are μ_0 and σ_0, respectively. We also assume that under faulty conditions, two basic types of changes can occur in these parameters :

- deviation from the reference mean value μ_0 towards μ_1, with constant standard deviation, as depicted in figure 1.1; in other words, this type of change is a systematic error. This example serves as a common basis for depicting the typical behavior of all the algorithms presented in chapter 2.
- increase in the standard deviation from σ_0 to σ_1, with constant mean, as depicted in figure 1.2. This type of change is a random error.

Composite changes can also occur. The problem is to design a statistical decision function and a decision rule that can detect these disorders. The typical behavior of such a decision function is depicted in figure 1.1.

In the simplest case, all the parameters of each of the two above-mentioned situations are assumed to be known. The tuning of a statistical decision rule is then reduced to the choice of a threshold achieving the requested tradeoff between the false alarm rate and the mean delay for detection. Several types of decision rules are used in industry as standards and are called *control charts*. They are described in detail in section 2.1.

The main references for quality control are [Aroian and Levene, 1950, Goldsmith and Whitfield, 1961, Van Dobben De Bruyn, 1968, Bissell, 1969, Phillips, 1969, Gibra, 1975, Wetherill and Brown, 1991]. Other references can be found in chapter 11.

1.2.2 Navigation System Monitoring

Navigation systems are typical equipments for planes, boats, rockets, and other moving objects. Important examples of such systems are inertial navigation systems, radionavigation systems, and global satellite navigation sets for planes. An inertial navigation system has two types of sensors : gyros and accelerometers. Using this sensor information and the motion equations, the estimation of the coordinates and the velocities of the moving object can be achieved. In view of safety and accuracy requirements, redundant fault-tolerant measurement systems are used. The first task of such a type of system is *detection and isolation of faulty sensors*. This problem can be stated as a particular change detection problem in some convenient modeling framework, as discussed in detail in chapter 11. The criterion to be used is again *quick detection and few false alarms*. Fast detection is necessary because, between the fault onset time and the detection time, we use abnormal measurements in the navigation equations, which is highly undesirable. On the other hand,

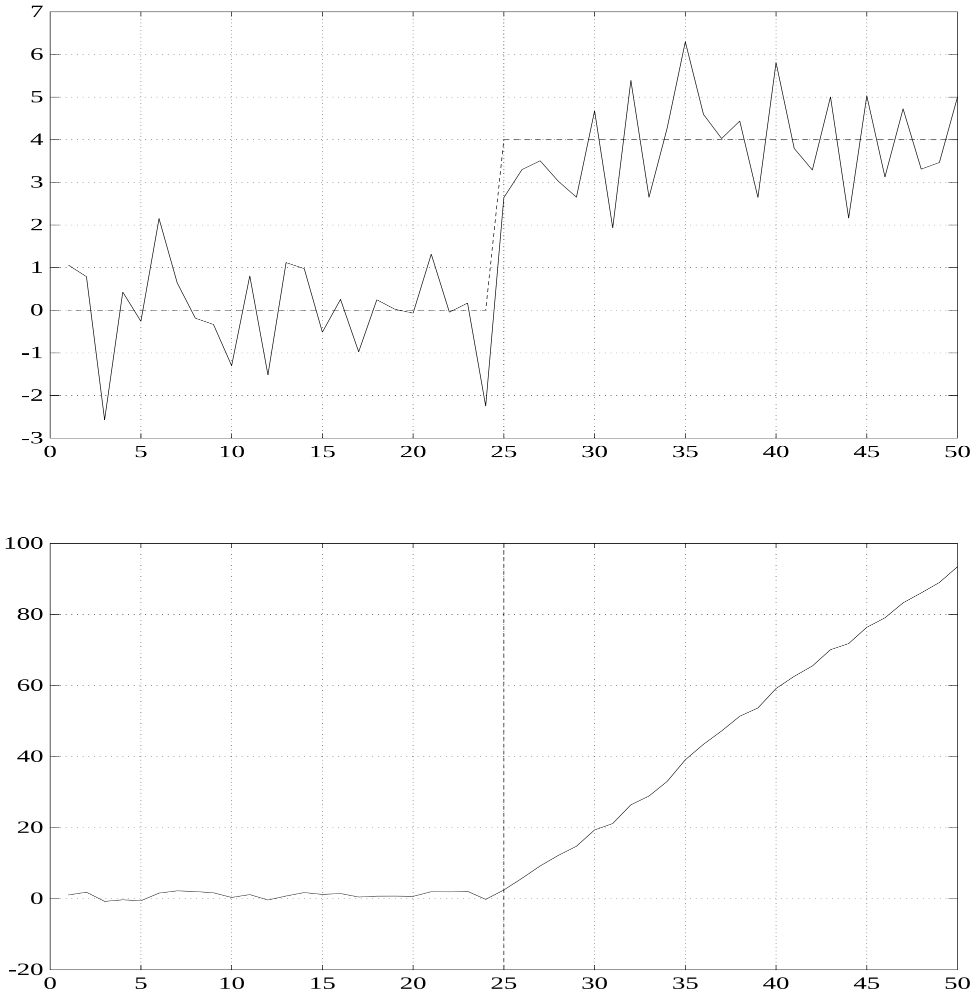

Figure 1.1 Increase in mean with constant variance and the typical behavior of the decision function in quality control.

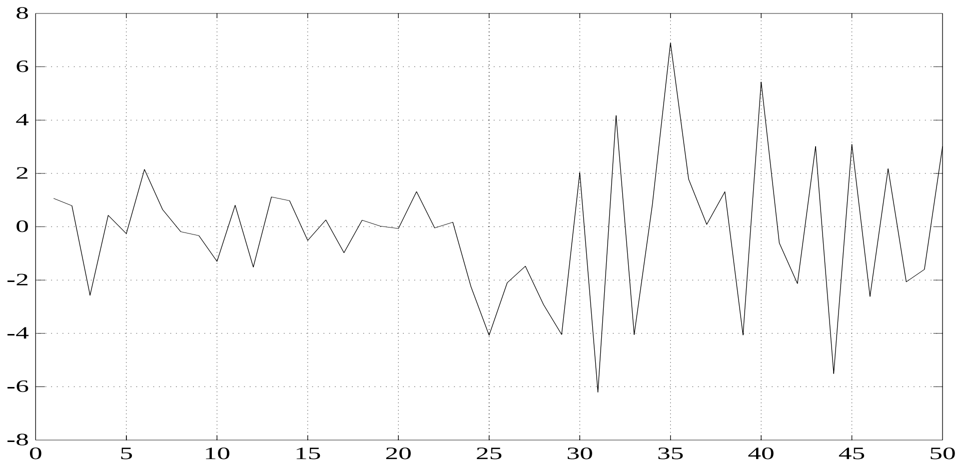

Figure 1.2 Increase in variance with constant mean.

false alarms result in lower accuracy of the estimate because some correct information is not used. The optimal solution is again a tradeoff between these two contradictory requirements. For radionavigation systems, integrity monitoring using redundant measurements is an important problem and is generally solved with the aid of the same criteria.

The main references for the monitoring of inertial navigation systems are [Newbold and Ho, 1968, Willsky *et al.*, 1975, Satin and Gates, 1978, Kerr, 1987]. Integrity monitoring of navigation systems is investigated in [Sturza, 1988]. Other references can be found in chapter 11.

1.2.3 Seismic Data Processing

Let us now discuss some typical problems of seismic data processing. In many situations, it is necessary to estimate *in situ* the geographical coordinates and other parameters of earthquakes. Typical three-dimensional signals are shown in figure 1.3, and comprise E-W, Z, and N-S measurements. The two main events to be detected are the P-wave and the S-wave; note that the P-wave can be very “small.” From the physical background in seismology, which we explain in chapter 11, it results that the processing of these three-dimensional measurements can be split into three tasks :

1. on-line detection and identification of the seismic waves;

2. off-line estimation of the onset times of these waves;

3. off-line estimation of the azimuth using correlation between components of P-wave segments.

From now on, we consider only the first two questions. Detection of the P-wave has to be achieved *very quickly with a fixed false alarms rate*. The main reason for this is to allow S-wave detection in this on-line processing. P-wave detection is a difficult problem, because the data contain many nuisance signals coming from the environment of the seismic station, and discriminating between these events and a true P-wave is not easy. The same situation holds for the S-wave, where the difficulty is greater, because of low signal-to-noise ratio and numerous nuisance signals between P-wave and S-wave.

After P-wave and S-wave detection, *off-line accurate estimation of onset times* is requested for both types of waves. As we explain in chapter 11, a possible solution consists of using some fixed size samples

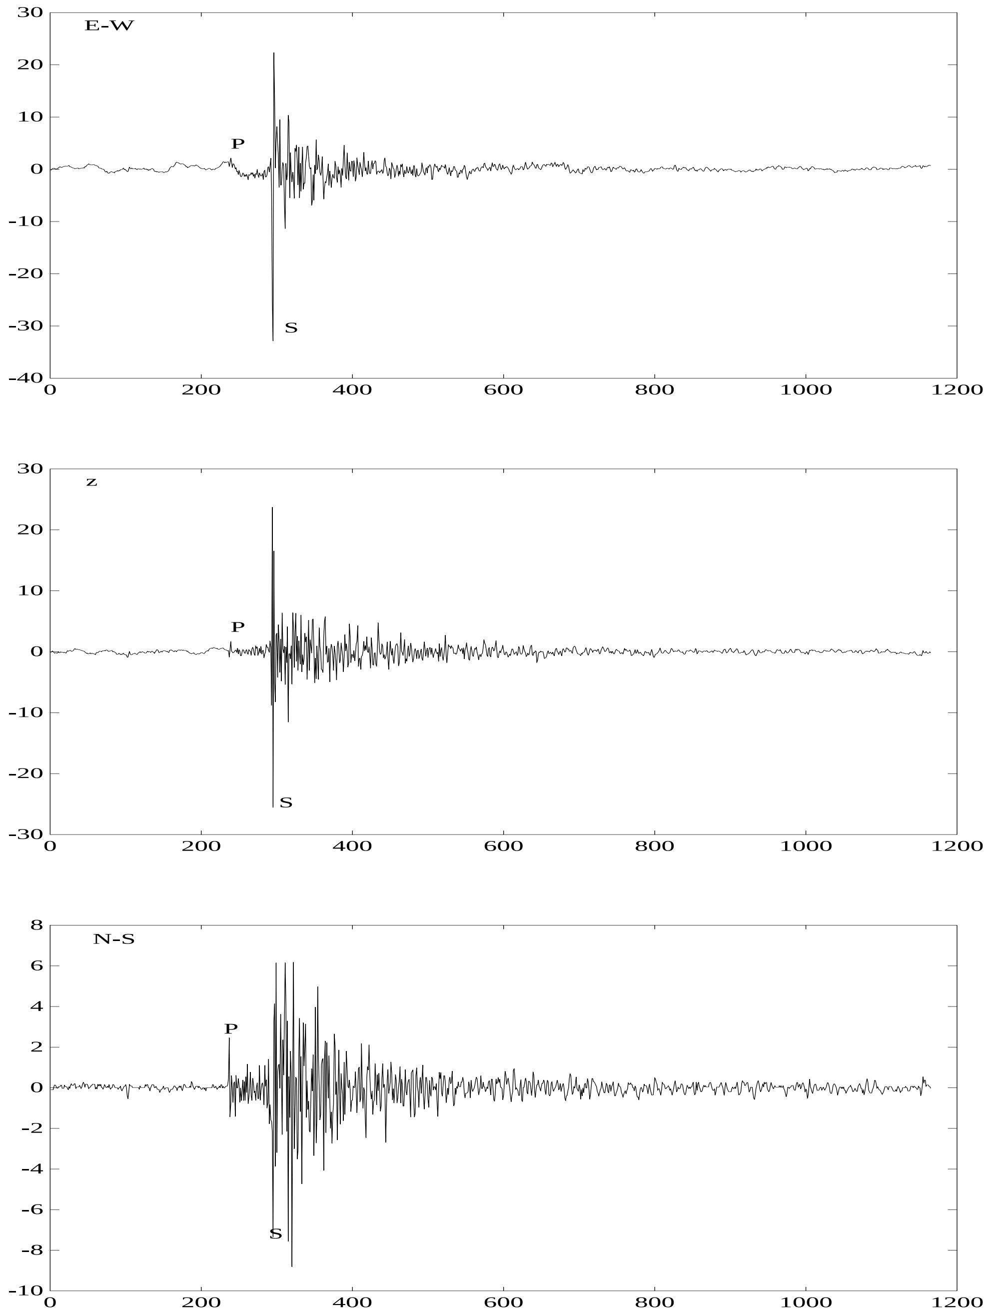

Figure 1.3 The three typical components of a seismogram : E-W, Z and N-S (Courtesy of the Academy of Sciences of USSR, Far Eastern Scientific Center, Institute of Sea Geology) .

of the three-dimensional signals, centered at a rough estimate of the onset time provided by the detection algorithm. This off-line change time estimation is described in section 8.7.

The main references for seismic data processing are [Tjostheim, 1975, Kitagawa and Gersch, 1985, Nikiforov and Tikhonov, 1986, Pisarenko *et al.*, 1987, Nikiforov *et al.*, 1989]. Other references can be found in chapter 11.

1.2.4 Segmentation of Signals

A possible approach to recognition-oriented signal processing consists of using an automatic segmentation of the signal as the first processing step. A segmentation algorithm splits the signal into homogeneous segments, the lengths of which are adapted to the local characteristics of the analyzed signal. The homogeneity of a segment can be in terms of the mean level or in terms of the spectral characteristics. This is discussed further when we introduce the additive and nonadditive change detection problems. The segmentation approach has proved useful for the automatic analysis of various biomedical signals, for example, electroencephalograms [R.Jones *et al.*, 1970, Bodenstein and Praetorius, 1977, Sanderson and Segen, 1980, Borodkin and Mottl', 1976, Ishii *et al.*, 1979, Appel and von Brandt, 1983], and electrocardiograms [Gustafson *et al.*, 1978, Corge and Puech, 1986]. Segmentation algorithms for recognition-oriented geophysical signal processing are discussed in [Basseville and Benveniste, 1983a]. More recently, a segmentation algorithm has been introduced as a powerful tool for the automatic analysis of continuous speech signals, both for recognition [André-Obrecht, 1988] and for coding [Di Francesco, 1990]. An example of automatic segmentation of a continuous (French) speech signal[1] is shown in figure 1.4. Other examples are discussed in chapter 11.

The main desired properties of a segmentation algorithm are *few false alarms and missed detections, and low detection delay*, as in the previous examples. However, keep in mind the fact that the segmentation of a signal is often nothing more than the first step of a recognition procedure. From this point of view, it is obvious that the properties of a given segmentation algorithm also depend upon the processing of the segments which is performed at the next stage. For example, it is often the case that, for segmentation algorithms, false alarms (sometimes called oversegmentation) are less critical than for onset detection algorithms. A false alarm for the detection of an imminent tsunami obviously has severe and costly practical consequences. On the other hand, in a recognition system, false alarms at the segmentation stage can often be easily recognized and corrected at the next stage. A segmentation algorithm exhibiting the above-mentioned properties is potentially a powerful tool for a recognition system.

It should be clear that a segmentation algorithm allows us to detect several types of events. Examples of events obtained through a spectral segmentation algorithm and concerning recognition-oriented speech processing are discussed in [André-Obrecht, 1988, André-Obrecht and Su, 1988, André-Obrecht, 1990].

1.2.5 Vibration Monitoring of Mechanical Systems

Let us now describe the *vibration monitoring* problem and its connection with change detection. For both complex mechanical structures, such as offshore platforms, bridges, buildings, and dams, and rotating machines, such as turbo-alternators and gearing systems, it is of crucial interest to monitor the vibrating characteristics without using artificial excitation or stop-down, but in the usual functioning mode under natural or usual excitation (swell, road traffic, wind, water pressure, earthquakes, big works in the neighborhood, steam). The vibrating characteristics of a mechanical structure or machine basically reflects its state of health, and any deviation in these characteristics brings information of importance to its functioning mode. The main difficulty in this problem is that the measured signals (accelerometers, gauges) reflect both the

[1]This result is due to Régine André-Obrecht. The help of Bernard Delyon in drawing this figure is also gratefully acknowledged.

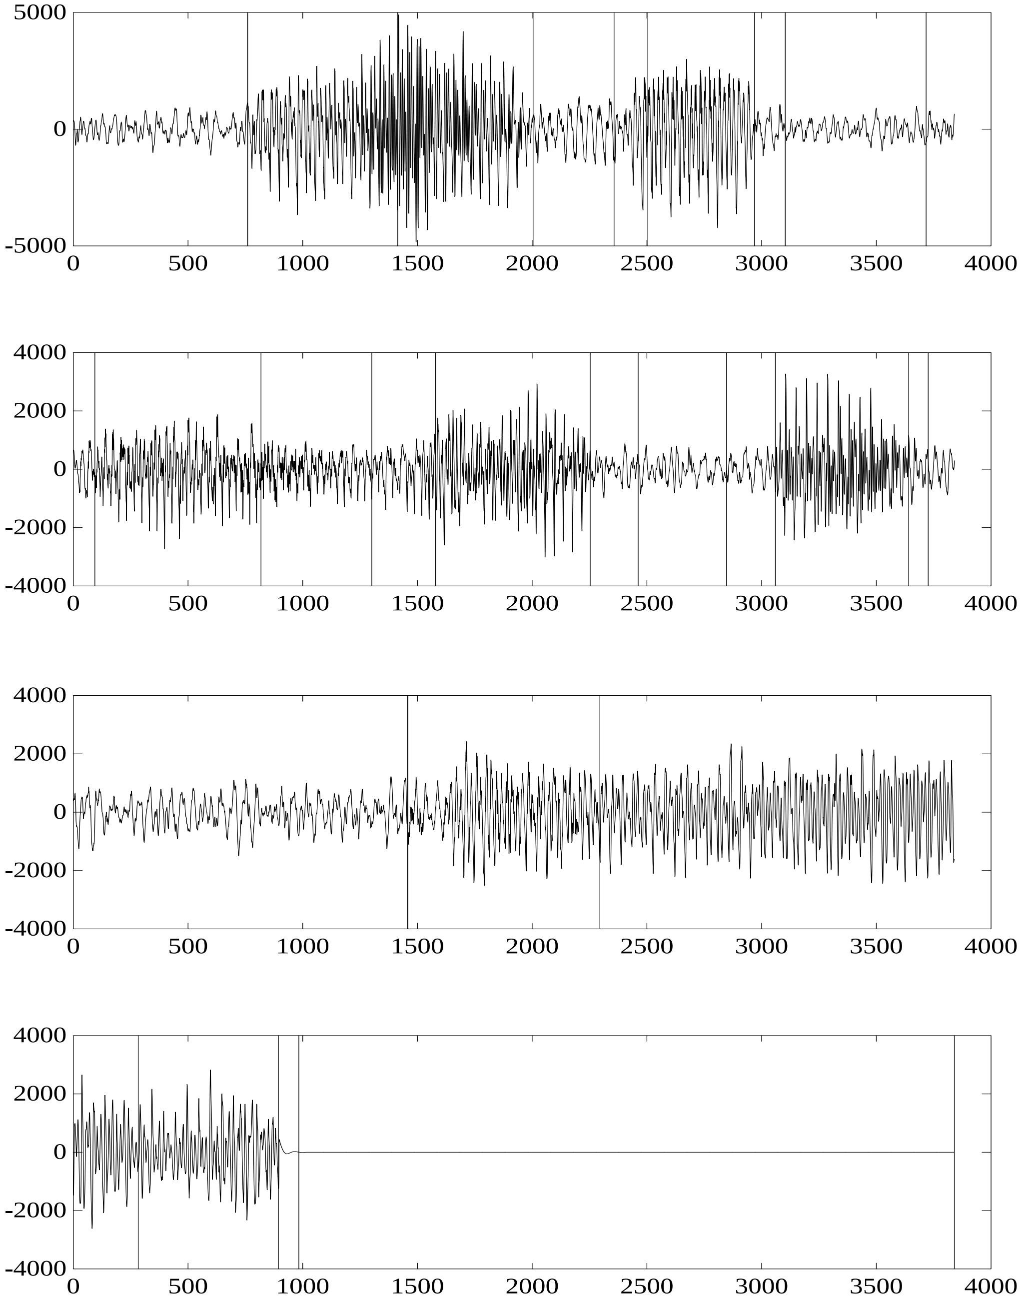

Figure 1.4 An example of speech signal segmentation. The estimated change times (vertical lines) provide us with the boundaries of the segments.

nonstationarities due to the surrounding excitation, which is always highly time-varying, and the nonstationarities due to changes in the eigen characteristics of the mechanical object itself. We show in chapter 11 that this vibration monitoring problem can be stated as the problem of detecting changes in the AR part of a multivariable ARMA model having nonstationary MA coefficients. Typical changes to be detected have a magnitude of about 1% of the eigenfrequencies. The second difficult problem is to diagnose or isolate the detected changes, either in terms of the vibrating characteristics (eigenvalues and eigenvectors), or in terms of the mechanical characteristics (masses, stiffness coefficients, together with an approximate localization in the mechanical object). These questions are investigated in detail in section 9.3. The criteria that are to be used in such a problem are *few false alarms* and ability to *detect small changes* in possibly long samples of data.

Small mechanical systems, such as a small number of masses connected by springs, often serve as laboratory experimental setups for simulating more complex vibrating structures. Thus, they can be used for testing fault detection and diagnostic tools. Examples can be found in [Kumamaru *et al.*, 1989, Basseville *et al.*, 1987a]. The models of these simulation examples are described in the appendix to chapter 11. The main references concerning signal processing methods for vibration monitoring can be found in [Braun, 1986].

1.3 Content of the Book

In this section, we describe in detail the content of the book. First, let us give some comments referring to the three problem statements described in section 1.1. Even though several chapters address the second and third problems, the main emphasis of this book is on the first problem, namely *on-line change detection using a parametric statistical approach*. We now describe the general organization of the book, then the content of each chapter; finally, we present and discuss the flowchart of the book.

1.3.1 General Organization

The organization of the chapters follows a simple distinction between changes in the *scalar parameter* of an *independent* sequence of observations, and changes in the *multidimensional parameter* of a *dependent* sequence. Thus, we divide the book into two main parts corresponding to these two sets. A third part is devoted to the tuning and application issues. The organization of the second part about multidimensional changes is based on a classification of change detection problems into two categories : *additive* changes and *nonadditive* (or spectral) changes. Basically, we mean that changes can be viewed as either additive or multiplicative on the transfer function of the considered signal or system. Equivalently, changes can be viewed as either changes in the mean value of the law of the observed signals, or changes in the correlations. A more thorough discussion about this classification can be found at the beginning of Part II.

1.3.2 Description of Each Chapter

Before proceeding, let us mention that, at the end of each chapter, the reader can find notes and bibliographical references concerning the problems discussed and a summary of the key results.

Part I is devoted to *changes in the scalar parameter* of an independent sequence. In chapter 2, we introduce the reader to the theory of on-line change detection algorithms in the framework of an independent random sequence parameterized by a scalar parameter. We first consider the case of known parameters before and after change. In section 2.1, we begin with the description of elementary algorithms of common use in industrial applications (quality control, for example) : these are Shewhart control charts, finite or infinite moving average control charts, and filtered derivative algorithms. In section 2.2, we introduce a key detection tool, the CUSUM algorithm, which we derive using both on-line and off-line points of view.

In section 2.3, we describe Bayes-type algorithms. In the case of an unknown parameter after change, we discuss two possible solutions in section 2.4 : the weighted CUSUM and the generalized likelihood ratio (GLR). In section 2.5, we discuss how algorithms used for detecting changes can improve the tracking ability of an adaptive identification scheme. Finally, in section 2.6, we discuss the two off-line problem statements introduced in subsection 1.1.2 : off-line hypotheses testing and estimation of the change time.

Chapters 3 and 4 are an excursion outside the part devoted to changes in a scalar parameter, and are aimed at the presentation of all the *theoretical backgrounds* to be used throughout the book. Chapter 3 is composed of two sections. The first is devoted to the presentation of the main results from probability theory, including conditional probability and expectation, Brownian motion and diffusion processes, martingales, and stopping times. In section 3.2, we summarize some results from the control literature, namely observers, Kalman filter, and connections between state-space and ARMA models. Chapter 4 is composed of four sections. Section 4.1 is concerned with some basic results about estimation and information from a mathematical statistics point of view. Section 4.2 is devoted to statistical hypotheses testing, including expansion of likelihood ratios, and section 4.3 to sequential analysis. Finally, in section 4.4, we formally define the criteria for designing and evaluating change detection algorithms in both the on-line and off-line frameworks.

In chapter 5, we come back to changes in the scalar parameter of an independent random sequence, and present the main analytical and numerical results concerning the algorithms presented in chapter 2. We investigate the properties of the elementary algorithms in section 5.1. Then in section 5.2, we describe in detail the properties of CUSUM-type algorithms, following the key results of Lorden. The properties of the GLR algorithm are discussed in section 5.3, together with the role of *a priori* information. Bayes-type algorithms are briefly investigated in section 5.4. Finally, in section 5.5, we present analytical and numerical comparative results. This concludes the Part I.

Part II is concerned with the *extension* of these algorithms to more complex situations of changes, namely changes in the vector parameter of an independent sequence, additive changes in a possibly dependent sequence, and nonadditive changes in a dependent sequence too. The key ideas of Part II are described in chapter 6.

Chapter 7 is devoted to the extension of the key algorithms developed in the independent case considered in chapter 2, to additive changes in more complex models, namely regression, ARMA, and state-space models. In section 7.1, we introduce general additive changes, and explain transformations from observations to innovations and redundancy relations. Section 7.2 deals with the statistical tools for detecting additive changes. We begin by discussing in subsection 7.2.1 what we call the *basic problem* of detecting a change in the mean vector parameter of an independent Gaussian sequence. Then we discuss the extension of the CUSUM-type and GLR detectors to the more general situations of regression, ARMA, and state-space models in subsections 7.2.2, 7.2.3, and 7.2.4, respectively. Still from a statistical point of view, we then discuss the diagnosis or isolation problem and the detectability issue in subsections 7.2.5 and 7.2.6. The properties of these algorithms are discussed in section 7.3. Section 7.4 is devoted to the presentation of geometrical tools for change detection and diagnosis, known as analytical redundancy techniques. We begin the discussion about redundancy by describing in subsection 7.4.1 the direct redundancy often used in the case of regression models. We extend this notion to the temporal redundancy in subsection 7.4.2. In subsection 7.4.3, we describe another technique for generating analytical redundancy relationships. We conclude this section with a discussion of the detectability issue in subsection 7.4.4, again from a geometrical point of view. This chapter about additive changes concludes with section 7.5, which contains a discussion about some basic links between statistical and geometrical tools. Actually, links exist for the design of detection algorithms as well as for the solutions to the diagnosis problem and the detectability definitions.

Chapter 8 addresses the problem of detecting changes in the spectral properties of a scalar signal by using parametric approaches. We mainly focus on on-line algorithms. In section 8.1, we first introduce spectral changes and explain their specificities and difficulties with respect to additive changes. We show why the transformation from observations to innovations used for additive changes is not sufficient here, and we introduce to the use of the local approach for change detection. In section 8.2, we investigate the general case of conditional probability distributions, and we describe the main ideas for designing on-line algorithms, namely CUSUM and GLR approaches, and possible simplifications, including the local approach, and leading to either linear or quadratic decision functions. All these algorithms are then described in the cases of AR and ARMA models in section 8.3. In section 8.4, we describe the design of non-likelihood-based algorithms, also using the local approach. This extended design allows a systematic derivation of change detection and diagnosis algorithms associated with any recursive parametric identification method. In section 8.5, we discuss the detectability issue. In section 8.6, we discuss the implementation issues related to the fact that, in practice, the model parameters before and after change are not known. In section 8.7, we consider off-line algorithms, using the likelihood approach, and discuss the connection with on-line algorithms.

Chapter 9 is concerned with spectral changes in the multidimensional case, including the diagnosis problem, and the properties of the detection algorithms in both the scalar and the multidimensional cases. In section 9.1, we introduce the key detection tools, namely the likelihood ratio, the local approach, and the non-likelihood-based algorithms, emphasizing the new multidimensional issues. Then in section 9.2, we extend the likelihood-based algorithms of chapter 8 to multidimensional AR and ARMA models. Section 9.3 is concerned with the application of the non-likelihood-based design of algorithms to the problem of the detection and diagnosis of changes in spectral characteristics of multidimensional signals, or equivalently in the eigenstructure of nonstationary multivariable systems. We describe both on-line and off-line detection algorithms. Then we investigate the diagnosis problem from several points of view. The detectability issue is discussed in section 9.4, from a statistical point of view, as in chapters 7 and 8. The theoretical properties of the various algorithms introduced in this and the previous chapters, are investigated in section 9.5. This concludes the Part II.

We begin the Part III with chapter 10, which is devoted to the problems of implementing and tuning change detection algorithms. This chapter is divided into four sections. In section 10.1, we describe a general methodology for implementing and tuning the algorithms. With respect to the design of the algorithms, this methodology is more philosophical than technical, but it relies on the available theoretical results concerning the properties of the algorithms. Section 10.2 is concerned with the tuning of all the techniques introduced in chapter 2 and investigated in chapter 5, namely the algorithms for detecting changes in the scalar parameter of an independent sequence. In section 10.3, we investigate the case of a vector parameter and a linear decision function, and in section 10.4, the case of a quadratic decision function.

In chapter 11, we come back to the applications. The main goals of this chapter are to show examples of the use of change detection algorithms and examples of potential application of the change detection methodology. Of course, the list of application domains that we investigate there is not exhaustive. The examples of the first type are fault detection in inertial navigation systems, onset detection in seismic signal processing, continuous speech signals segmentation, and vibration monitoring. The examples of the second type are statistical quality control, biomedical signal processing, and fault detection in chemical processes.

1.3.3 Flowchart of the Book

In figure 1.5, we show the general organization of the book and suggestions for using it. Two paths can be used for reading this book. The reader interested mainly in the algorithms themselves can start at beginning with the design of the algorithms, proceed through the properties, and finally reach tuning and applications. The reader interested mainly in the practical design and application of the algorithms can start with the applications at the end in order to select his path through the other chapters.

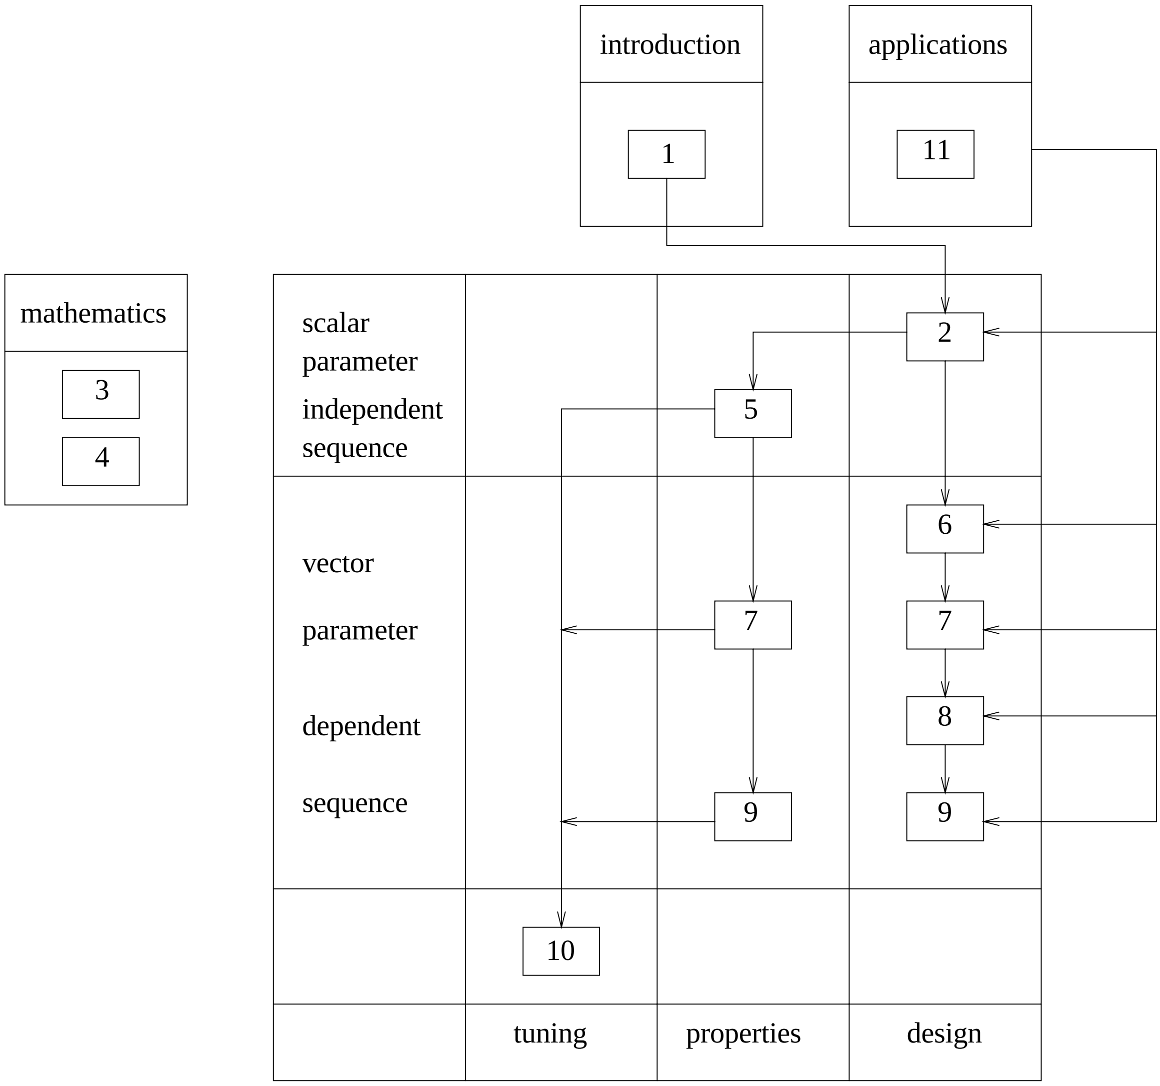

Figure 1.5 Flowchart of the book, showing two paths : one focused on the investigation of the algorithms and the other on the practical design and applications.

1.4 Some Further Critical Issues

In this section, we comment further on some of the issues involved when designing change detection algorithms and investigating their properties. Because we deal mainly with parametric techniques in this book, the key issues of choice of models, use of prior information, redundancy, and nuisance parameters have to be addressed. This is done in the first subsection. We then discuss the properties of the algorithms and detectability.

1.4.1 Designing the Algorithms

We now consider questions related to choice of models and use of prior information, generation of residuals, and nuisance parameters.

1.4.1.1 Choice of Models and Use of Prior Information

When dealing with modeling, and more specifically with parametric techniques, the choice of models is a critical issue. The reader is referred to [Ljung, 1987] for investigation, discussions, and references about linear time-invariant models, or time-varying or nonlinear models. What we would like to stress here is that, in the framework of change detection, the situation is significantly different from the identification point of view. Basically, *models useful for detection and monitoring are usually smaller than physical models and models for identification and recognition* (see the discussion of speech segmentation and recognition in section 11.1).

For example, consider the problem of detecting and diagnosing changes or faults in large structures or industrial processes. Useful results can be obtained with the aid of parametric models of relatively small size with respect to the dimension of the physical model of the process (which is based on partial differential equations, for example) . Even though it is often believed that parametric techniques are useful for diagnosis purposes only when there is a bijection between the parametric model and the physical one (see the survey [Isermann, 1984], for example), diagnosis in terms of the physical model can be inferred from a small black-box parametric model. This has been obtained in the vibration monitoring application introduced in section 1.2, and is discussed in detail in chapter 9. As another example, a relevant segmentation can be obtained with the aid of AR models of order 2, whereas the classification of the resulting segments may very well require AR models of significantly higher orders. This is discussed in section 11.1 for the case of continuous speech signals.

Another important issue when designing change detection algorithms is the use of prior information about the changes. When model structure and parameterization have been chosen, it is useful, if not necessary, to examine what is known about the possible values of the parameters before and after change, and how this prior information should be used.

Referring to the preliminary problem statement, which we formulate in section 1.1, from the on-line point of view, knowing the parameter θ_0 before change is of secondary interest. If θ_0 is unknown, it may be identified with the aid of a convenient identification algorithm. The actual problem lies then in the parameter θ_1 after change. Three cases have to be distinguished :

1. θ_1 is known : this is an easy but unrealistic case. It is often used as a starting point for the design of a detection algorithm, which is then extended to more realistic situations, for example, by replacing unknown parameters by values fixed *a priori* (such as a “minimum” magnitude of jump), or by estimated values. This case is the preferential situation for the derivation of most theoretical optimality results for change detection algorithms, and comparison between these theoretical results and numerical ones from simulation analysis. It may be also useful to compare empirical estimation of the performances of the change detection algorithm on real data to these theoretical properties. See sections 2.2, 7.2, and 8.2 for examples of such design.

2. Few prior information on values of θ_1 corresponding to interesting changes is available; for example, it is known that there exists a separating hyperplane between the set of values of θ_0 and the set of values of θ_1. How this type of information can be used in the design of a change detection algorithm is explained in sections 7.2 and 8.2.

3. Nothing is known about θ_1. This situation is obviously the most interesting from a practical point of view, but also the most difficult from the point of view of the design and the investigation of the properties of the algorithms. Two main approaches exist for solving this problem, and are described in sections 2.4, 7.2 and 8.2. Because the corresponding algorithms are complex, we also investigate possible simplifications.

1.4.1.2 Redundancy Relationships

We now discuss the use of analytical redundancy relationships for change detection. As we stated at the end of section 1.1, one possible general approach for change detection consists of splitting the task into (1) generation of residuals, which are, for example, ideally close to zero when no change occurs, and significantly different from zero after a change, and (2) design of decision rules based on these (possibly non-statistically optimal) residuals. One way of obtaining such residuals is to use analytical redundancy relations. For example, in chemical processes, static balance equations are helpful for detecting failures in pipes, sensors, and actuators for fluid flows.

For other systems, a complete model may be available and can be used in the formal statistical change detection approach. In this case, the generation of residuals is basically included in the derivation of the algorithm itself and does not have to be considered as a separate task. For example, we use this point of view for discussing fault detection in an inertial navigation system.

There exists a bridge between these two types of solutions, and in chapter 7 we show in which cases they are equivalent.

1.4.1.3 Nuisance Parameters

Assume that a parametric model is characterized by a parameter vector which is divided into two subsets : one subset is useful for detecting changes in the properties of the underlying object; the other subset contains information about the object or its environment, but the changes in this subset are not of interest. It turns out that very often these nuisance parameters are highly involved with the useful parameters, and thus have an influence on the decision function. The use of change detection algorithms that do not take into account this fact leads to additional false alarms and missed detections. A specific design of the change detection algorithm must be used in this case. The so-called min-max approach is introduced in chapter 4 for this purpose. A problem that is very close to the question of nuisance parameters is the problem of isolation or diagnosis. We show in section 7.2 how to use this specific approach to design change detection algorithms to solve the isolation problem. Another example is investigated in section 9.3, where we show that it is possible to design decision functions that decouple as much as possible these two parameter subsets, for example, AR and MA parameters in ARMA models.

1.4.2 Investigating the Algorithms

The investigation of the properties of algorithms is useful for two purposes : First, it helps us understand what can be gained in practice when using such algorithms; and second, it gives answers to the optimality issues. We now discuss these points, distinguishing between the properties that result from a formal definition of criteria and those that result from a weaker but useful performance index.

1.4.2.1 Properties of the Algorithms

The mean delay for detection and the mean time between false alarms are the two key criteria for on-line change detection algorithms. As we discuss in chapter 5, in some cases there exist optimal algorithms that minimize the mean delay for a given mean time between false alarms. From a practical point of view, knowledge of the values of these performance indexes for sets of parameters is useful. A key tool for investigating the properties of on-line change detection algorithms is the so-called *average run length function*, which concentrates the information about both these performances indexes. The computation of this function is difficult for most of the practically relevant change detection problems. For this reason, we introduce numerical algorithms for the evaluation of this function. We also introduce a weaker performance index, which we call *detectability*, that is strongly connected with the two previous criteria and can be computed in more complex cases.

1.4.2.2 Detectability

For defining the detectability of a given change, two levels can be considered. The first investigates which changes are detectable and which are not. In the same way that observability and controllability depend on the observation and control matrices of the system, the detectability depends on the statistical tool that is used for detection. Therefore, the detectability of a change should be defined in terms of the effect or signature that it produces on the “sufficient” statistic that is used in the decision rule. For example, if the statistic reflects possible changes in the system by changes in its own mean value, any change that does not modify the mean value of the statistic is not detectable.

A second level defines the detectability as a performance index of the decision rule. We discuss this detectability issue using both statistical and geometrical points of view, and unify the different definitions into the framework of *information*. More precisely, in the statistical point of view, we define the detectability of a change with the aid of an intrinsic feature of the system, namely the mutual information between the two models before and after change. We show that, surprisingly, the two points of view – “the detectability depends upon the detection tool which is used” and “the detectability is an intrinsic feature of the analyzed system” – basically lead to only one definition of detectability. We discuss these detectability issues in subsections 7.2.6, 7.4.4, and 7.5.4, and again in sections 8.5 and 9.4.

1.5 Notes and References

In this section, we give some historical notes and then references for seminars, survey papers, and books related to change detection. We believe it of interest to put the *on-line* change detection framework, motivations, and methodology in a historical perspective. Because this subject basically grew up at the confluence of several disciplines, a complete historical picture is difficult to draw. Our partial knowledge can be summarized as follows.

1.5.1 Historical Notes

We distinguish two parallel directions of investigations, in the areas of mathematical statistics and automatic control theory, and then summarize investigations concerning the possible merging of these two directions.

1.5.1.1 Mathematical Statistics

Interest in on-line change detection probably arose first in the area of quality control, where *control charts* were introduced in [Shewhart, 1931] and then *cumulative sums charts* in [Page, 1954a]. The two main classes of statistical problem statements are the Bayesian and the non-Bayesian approaches.

Bayesian approach The first Bayesian change detection problem was stated in [Girshick and Rubin, 1952] to solve a typical on-line quality control problem for continuous technological processes. The first optimality results concerning Bayesian change detection algorithms were obtained in [Shiryaev, 1961, Shiryaev, 1963, Shiryaev, 1965]. Since then, the literature in this area has become quite wide. More recent investigations can be found in [Pollak and Siegmund, 1985, Pollak, 1985, Pollak, 1987].

Non-Bayesian approach The first investigation of non-Bayesian change detection algorithms was made in [Page, 1954a]. The asymptotic optimality of cumulative sum algorithms was proved in [Lorden, 1971]. Nonasymptotic optimality results can be found in [Moustakides, 1986, Ritov, 1990]. The extension of such techniques to composite hypotheses testing problems is discussed in [Lorden, 1971, Lorden, 1973, Pollak and Siegmund, 1975]. The generalization of Lorden's results to dependent processes is discussed in [Bansal and Papantoni-Kazakos, 1986].

1.5.1.2 Automatic Control

In the area of automatic control, change detection problems are referred to as model-based fault detection and isolation (FDI). The concept of analytical redundancy for fault detection was investigated approximately independently at the same time in the United States [Beard, 1971] and in the Soviet Union [Britov and Mironovski, 1972]. Further key developments concerning the geometrical aspects can be found in [E.Chow and Willsky, 1984, Lou *et al.*, 1986, Massoumnia, 1986, White and Speyer, 1987, Viswanadham *et al.*, 1987a, Wünnenberg, 1990] and are discussed in the survey papers [Willsky, 1976, Frank, 1990, Patton and Chen, 1991, Gertler, 1991]. Typically, the models used in these investigations are more complex than the models classically used in the mathematical statistics literature.

From a formal point of view, this research direction does not belong to the theory of change detection, because of the lack of statistical problem statements and criteria. Nevertheless, the main ideas underlying fault detection tools, namely the use of innovations or residuals for monitoring purposes, are very close to the concept of sufficient statistics for detection. For this reason, we think it useful to discuss these two types of concepts together. A first attempt to bring together both geometric concepts of analytical redundancy and statistical decision tools is the survey paper [Willsky, 1976].

1.5.1.3 Joint Approach

In the early 1970s, a new research direction arose, involving complex statistical models (much more complex than in classical statistical investigations). The main motivation for these new developments were unsuccessful attempts at using pure mathematical tools for solving concrete problems in the automatization of industrial processes. The starting point of these new investigations was the use of change detection decision rules for the more complex models, and the extension of the available theoretical results existing about them

[Lumel'sky, 1972, Nikiforov, 1975, Bagshaw and R.Johnson, 1977, Nikiforov, 1978, Nikiforov, 1980, Segen and Sanderson, 1980, Basseville and Benveniste, 1983a, Basseville and Benveniste, 1983b, Vorobeichikov and Konev, 1988]. Survey papers reporting these investigations are [Basseville, 1982, Kligiene and Telksnys, 1983, Basseville, 1988].

1.5.1.4 Investigations in Application Domains

The problem of detecting abrupt changes in properties of signals and dynamic systems has received increasing attention in the last twenty years. One key reason for that is its connection to the problem of fault detection, and strong industrial needs in the area of condition-based maintenance and monitoring of plants. Another reason is its usefulness in time-series analysis and signal processing for recognition purposes.

Several books related to quality control exist, such as [Shewhart, 1931, Woodward and Goldsmith, 1964, Van Dobben De Bruyn, 1968, Duncan, 1986]. The analysis of biomedical signals, especially electroencephalograms, is another field where many contributions to the problem of automatic segmentation of signals have been made [R.Jones *et al.*, 1970, Borodkin and Mottl', 1976, Mathieu, 1976, Segen and Sanderson, 1980]. The interest in the change detection methodology in this area is reflected in [Cohen, 1987], where segmentation algorithms are presented as basic signal processing tools. Geophysical signal processing can also be achieved with the aid of segmentation algorithms; for example, diagraphy [Basseville and Benveniste, 1983a] and seismology [Nikiforov and Tikhonov, 1986, Nikiforov *et al.*, 1989]. Automatic segmentation was introduced as a first step toward continuous speech recognition in [André-Obrecht, 1988], and as a first step toward speech coding in [Di Francesco, 1990], both using the algorithm presented in [Basseville and Benveniste, 1983b].

Interest in the change detection methodology also arose in chemical engineering [Himmelblau, 1978]. In the field of econometry, two books are devoted to the problem of structural change detection, i.e., the problem of detection of changes in the parameters of an econometric model. These are [Poirier, 1976, Broemeling and Tsurumi, 1987]. An annotated bibliography can also be found in [Shaban, 1980].

Many other application domains have been investigated, as can be seen from the long list of application studies of innovation-based fault detection/diagnosis methods in [Patton *et al.*, 1989] and [Tzafestas *et al.*, 1987].

1.5.1.5 Related Investigations

The most closely related investigations concern the *off-line* change detection and estimation problems. The historical starting point of these studies is [Page, 1957]. Subsequent investigations are in [Hinkley, 1970, Hinkley, 1971, Kligiene and Telksnys, 1983]. More generally, complete theoretical optimality results about the likelihood approach to change detection are obtained in [Deshayes and Picard, 1979, Deshayes and Picard, 1983].

1.5.2 Seminars, Surveys, and Books

Two national seminars on change detection were organized in 1984 independently in Paris, France, and in Palanga, USSR, emphasizing great interest and activity in this field in both countries. The contents of these seminars are presented in [Basseville and Benveniste, 1986, Telksnys, 1987]. Two subsequent seminars took place in Moscow, USSR, and in Voronej, USSR, in 1988 and 1990, respectively. Many international conferences and workshops in the area of automatic control have had sessions on fault detection and isolation over the last fifteen years.

Many survey papers about this problem have been published over the past twenty years, for example, four survey papers in *Automatica* [Willsky, 1976, Isermann, 1984, Basseville, 1988, Frank, 1990] and two in

Automation and Remote Control [Mironovski, 1980, Kligiene and Telksnys, 1983]; one survey about sensor failure detection in jet engines [Merril, 1985]; and three survey papers in the econometry literature [Zacks, 1983, Krishnaiah and Miao, 1988, Csörgö and Horváth, 1988].

We now list some of the books in this area. The topic of statistical tools for change detection is investigated in [Woodward and Goldsmith, 1964, Van Dobben De Bruyn, 1968, Shiryaev, 1978, Nikiforov, 1983, Basseville and Benveniste, 1986, Siegmund, 1985b, Telksnys, 1987, Zhigljavsky and Kraskovsky, 1988, Brodskiy and Darkhovskiy, 1992]. Books oriented more toward geometrical tools are [Tzafestas *et al.*, 1987, Singh *et al.*, 1987, Patton *et al.*, 1989]. The book [Viswanadham *et al.*, 1987b] basically put together, but without integration, reliability theory and fault-tolerant computer systems on one hand and fault detection and diagnosis on the other hand. More specific books are [Himmelblau, 1978, Pau, 1981].