Multi-Sensor Data Fusion

Multi-Sensor Data Fusion

An Introduction

With 81 Figures and 59 Tables



Dr. H.B. Mitchell homepage: www.hbmitchell.com

Library of Congress Control Number: 2007926176

ISBN 978-3-540-71463-7 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media springer.com © Springer-Verlag Berlin Heidelberg 2007

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: by the author and Integra, India using a Springer LATEX macro package Cover design: Erich Kirchner, Heidelberg

Printed on acid-free paper SPIN: 11911104 543210

This book is dedicated to the memory of my parents

WOOLF AND BELLA CISSIE MITCHELL

They surrounded me with love and gently encouraged me in my studies. They were never discouraged nor lost faith in me or my ability. This book is a small token of my love for them.

May their memory be a blessing

Preface

The purpose of this book is to provide an introduction to the theories and techniques of multi-sensor data fusion. The book has been designed as a text for a one-semester graduate course in multi-sensor data fusion. It should also be useful to advanced undergraduates in electrical engineering or computer science who are studying data fusion for the first time and to practising engineers who wish to apply the concepts of data fusion to practical applications.

The book is intended to be largely self-contained in so far as the subject of multi-sensor data fusion is concerned, although some prior exposure to the subject may be helpful to the reader. A clear understanding of multi-sensor data fusion can only be achieved with the use of a certain minimum level of mathematics. It is therefore assumed that the reader has a reasonable working knowledge of the basic tools of linear algebra, calculus and simple probability theory. More specific results and techniques which are required are explained in the body of the book or in appendices which are appended to the end of the book.

Although conceptually simple, the study of multi-sensor data fusion presents challenges that are fairly unique within the education of the electrical engineer or computer scientist. Unlike other areas encounted by a student of these subjects, multi-sensor data fusion draws on, and brings together, theories and techniques that are typically taught separately in the traditional curricula. To become competent in the field the student must be familiar with tools taken from a wide range of diverse subjects, including: neural networks, signal processing, statistical estimation, tracking algorithms, computer vision, and control theory. All too often the student views multi-sensor data fusion as a miscellaneous assortment of processes and techniques which bear no relationship to each other. We have attempted to overcome this problem by presenting the material using a common statistical Bayesian framework. In this way the different theories and techniques are clearly integrated and the underlying pattern of relationships that exist between the different methodologies are emphasized. Furthermore, by adopting a single framework, we have kept the book at a reasonable size while treating many new and important topics in great depth. We should point out that we have not, however, ignored other frameworks when this seemed appropriate.

As with any other branch of engineering, multi-sensor data fusion is a pragmatic activity which is driven by practicalities. It is therefore important that the student is able to experiment with the different techniques presented in the book. For this purpose software code, written in Matlab, is particularly convenient and we have included details of relevant Matlab code which may be downloaded from the worldwide web. For the professional engineer we have both illustrated the theory with many real-life applications and have provided him with an extensive list of up-to-date references. Additional information, including material for course instructors, is available from the author's homepage: www.hbmitchell.com.

The book is based on seminars, lectures and courses on multi-sensor data fusion which have been taught over several years. The structure and content of the book is based on material gathered and ideas exchanged with my colleagues. Particular thanks are extended to Dr. Paul Schaefer and Mr. Michael Avraham with whom I have discussed most topics in the book and to Ms. Ruth Rotman, Prof. Brian Cowan and Prof. Stanley Rotman who have kindly read and commented on the various drafts of the book. I am also indebted to my wife and children for the support and patience they have shown me while the book was being written.

Finally, to the reader. We hope you will enjoy reading this book and that it will prove to be an useful addition to the increasingly important and expanding field of multi-sensor data fusion.

H.B. Mitchell

Contents

Part I Basics

1	Intr	roduction	3
	1.1	Definition	3
	1.2	Synergy	4
	1.3	Multi-Sensor Data Fusion Strategies	5
		1.3.1 Fusion Type	5
		1.3.2 Sensor Configuration	6
		1.3.3 Input/Output Characteristics	6
	1.4	Formal Framework	7
		1.4.1 Multi-Sensor Integration	9
	1.5	Catastrophic Fusion	10
	1.6	Organization	12
	1.7	Further Reading	13
2	\mathbf{Sen}	sors	15
2	Sen 2.1	sors Introduction	$15 \\ 15$
2	Sen 2.1 2.2	sors Introduction Smart Sensor	15 15 16
2	Sen 2.1 2.2 2.3	sors Introduction Smart Sensor Logical Sensors	$15 \\ 15 \\ 16 \\ 17$
2	Sen 2.1 2.2 2.3 2.4	sors Introduction Smart Sensor Logical Sensors Interface File System (IFS)	15 15 16 17 17
2	Sen 2.1 2.2 2.3 2.4	sors Introduction Smart Sensor Logical Sensors Interface File System (IFS) 2.4.1 Interface Types	15 15 16 17 17 18
2	Sen 2.1 2.2 2.3 2.4	sors Introduction Smart Sensor Logical Sensors Interface File System (IFS) 2.4.1 Interface Types 2.4.2 Timing	15 15 16 17 17 18 19
2	Sen 2.1 2.2 2.3 2.4 2.5	sors Introduction Smart Sensor Logical Sensors Interface File System (IFS) 2.4.1 Interface Types 2.4.2 Timing Sensor Observation	15 15 16 17 17 18 19 20
2	Sen 2.1 2.2 2.3 2.4 2.5	sorsIntroductionSmart SensorLogical SensorsInterface File System (IFS)2.4.1Interface Types2.4.2TimingSensor Observation2.5.1Sensor Uncertainty Δy	15 16 17 17 18 19 20 21
2	Sen 2.1 2.2 2.3 2.4 2.5 2.6	sorsIntroductionSmart SensorLogical SensorsInterface File System (IFS)2.4.1Interface Types2.4.2TimingSensor Observation2.5.1Sensor Uncertainty Δy Sensor Characteristics	15 15 16 17 17 18 19 20 21 23
2	Sen 2.1 2.2 2.3 2.4 2.5 2.6 2.7	sorsIntroductionSmart SensorLogical SensorsInterface File System (IFS)2.4.1Interface Types2.4.2TimingSensor Observation2.5.1Sensor CharacteristicsSensor Sensor Properties	15 15 16 17 17 18 19 20 21 23 23
2	Sen 2.1 2.2 2.3 2.4 2.5 2.6 2.7 2.8	sorsIntroductionSmart SensorLogical SensorsInterface File System (IFS)2.4.1Interface Types2.4.2TimingSensor Observation2.5.1Sensor Uncertainty Δy Sensor CharacteristicsSensor Sensor PropertiesSensor Model	15 15 16 17 17 18 19 20 21 23 23 23 24

Arc	29 hitecture
3.1	Introduction
3.2	Fusion Node
	3.2.1 Properties
3.3	Simple Fusion Networks
	3.3.1 Single Fusion Cell 33
	3.3.2 Parallel Network
	3.3.3 Serial Network 35
	3.3.4 Iterative Network
3.4	Network Topology
	3.4.1 Centralized
	3.4.2 Decentralized 39
	3.4.3 Hierarchical 42
3.5	Software
3.6	Further Reading 44
	Arc 3.1 3.2 3.3 3.4 3.4

Part II Representation

4	Cor	nmon Representational Format	47
	4.1	Introduction	47
	4.2	Spatial-Temporal Transformation	50
	4.3	Geographical Information System	51
		4.3.1 Spatial Covariance Function	54
	4.4	Common Representational Format	55
	4.5	Subspace Methods	58
		4.5.1 Principal Component Analysis	59
		4.5.2 Linear Discriminant Analysis	60
	4.6	Multiple Training Sets	64
	4.7	Software	67
	4.8	Further Reading	67
5	Spa	tial Alignment	69
	5.1	Introduction	69
	5.2	Image Registration	69
		5.2.1 Mutual Information	70
	5.3	Resample/Interpolation	74
	5.4	Pairwise Transformation T	76
	5.5	Image Fusion	77
	5.6	Mosaic Image	80
	5.7	Software	81
	5.8	Further Reading	82

6	Ten	nporal Alignment
	6.1	Introduction
	6.2	Dynamic Time Warping
	6.3	Dynamic Programming
		6.3.1 Derivative Dynamic Time Warping
		6.3.2 Continuous Dynamic Time Warping
	6.4	Video Compression
	6.5	Software
	6.6	Further Reading
7	Sen	sor Value Normalization
	7.1	Introduction
		7.1.1 Sensor Value Normalization
	7.2	Binarization
	7.3 Parametric Normalization Functions	
	7.4	Fuzzy Normalization Functions
	7.5 Ranking	
	7.6	Conversion to Probabilities
		7.6.1 Platt Calibration
		7.6.2 Binning
		7.6.3 Kernels
		7.6.4 Isotonic Regression
		7.6.5 Multi-Class Probability Estimates
	7.7	Software
	7.8	Further Reading

Part III Data Fusion

8	Вау	resian Inference
	8.1	Introduction
	8.2	Bayesian Analysis
	8.3	Probability Model
	8.4	A Posteriori Distribution
		8.4.1 Standard Probability Distribution Functions
		8.4.2 Conjugate Priors 119
		8.4.3 Non-Informative Priors
		8.4.4 Missing Data
	8.5	Model Selection
		8.5.1 Laplace Approximation
		8.5.2 Bayesian Model Averaging
	8.6	Computation
		8.6.1 Markov Chain Monte Carlo

	8.7	Software
	8.8	Further Reading
9	Par	ameter Estimation 133
0	9.1	Introduction 133
	9.2	Parameter Estimation
	9.3	Bayesian Curve Fitting 137
	9.4	Maximum Likelihood
	9.5	Least Squares
	9.6	Linear Gaussian Model
	0.0	9.6.1 Line Fitting 145
		9.6.2 Change Point Detection
		9.6.3 Probabilistic Subspace
	9.7	Generalized Millman Formula
	9.8	Software
	9.9	Further Reading
10	Rob	oust Statistics
	10.1	Introduction
	10.2	Outliers
	10.3	Robust Parameter Estimation
		10.3.1 Student- <i>t</i> Function
		10.3.2 "Good-and-Bad" Likelihood Function
		10.3.3 Gaussian Plus Constant
		10.3.4 Uncertain Error Bars
	10.4	Classical Robust Estimators
		10.4.1 Least Median of Squares
	10.5	Robust Subspace Techniques
	10.6	Robust Statistics in Computer Vision
	10.7	Software
	10.8	Further Reading
11	Seq	uential Bayesian Inference
	11.1	Introduction
	11.2	Recursive Filter
	11.3	Kalman Filter
		11.3.1 Parameter Estimation
		11.3.2 Data Association
		11.3.3 Model Inaccuracies
		11.3.4 Multi-Target Tracking
	11.4	Extensions of the Kalman Filter
		11.4.1 Robust Kalman Filter
		11.4.2 Extended Kalman Filter

		11.4.3 Unscented Kalman Filter	191
		11.4.4 Switching Kalman Filter	192
		11.4.5 Kriged Kalman Filter	194
	11.5	Particle Filter	195
	11.6	Multi-Sensor Multi-Temporal Data Fusion	195
		11.6.1 Measurement Fusion	195
		11.6.2 Track-to-Track Fusion	197
	11.7	Software	200
	11.8	Further Reading	200
	-	0	
12	Baw	esian Decision Theory	201
14	12 1	Introduction	201
	12.1 12.2	Pattern Becognition	201
	12.2	Naive Bayes' Classifier	201
	12.0	12.3.1 Representation	205
		12.3.2 Parformance	200
		12.3.2 Terrormance	200
	19.4	Modifications	207
	12.4	19.4.1 Feature Space	210
		12.4.1 Feature Space	210
		12.4.2 Model Assumptions	214
		12.4.5 Learning Methods	210
	10 E	12.4.4 Multiple Classifiers	210
	12.0	Definition Classifier	211
	12.0		218
	12.7	Software	219
	12.8	Further Reading	219
10	Б		001
13	Ense	emble Learning	221
	13.1	Introduction	221
	13.2	Bayesian Framework	221
	13.3	Empirical Framework	224
	13.4	Diversity Techniques	225
	13.5	Diversity Measures	227
		13.5.1 Ensemble Selection	230
	13.6	Classifier Types	230
	13.7	Combination Strategies	231
		13.7.1 Simple Combiners	231
		13.7.2 Meta-Learners	236
	13.8	Boosting	238
	13.9	Recommendations	240
	13.10	Software	240
	13.11	Further Reading	240

Part IV Sensor Management

14	Sensor Management
	14.1 Introduction
	14.2 Hierarchical Classification
	14.2.1 Sensor Control
	14.2.2 Sensor Scheduling
	14.2.3 Resource Planning
	14.3 Sensor Management Techniques
	14.3.1 Information-Theoretic Criteria
	14.3.2 Bayesian Decision-Making
	14.4 Further Reading
	14.5 Postscript

Part V Appendices

Software Sources				
Backgro	und Material			
B.1	Probability Theory			
B.2	Linear Algebra			
B.3	Square Matrices			
Referen	ces			
Index				

Basics

Introduction

1.1 Definition

The subject of this book is *multi-sensor data fusion* which we define as "the theory, techniques and tools which are used for combining sensor data, or data derived from sensory data, into a common representational format". In performing sensor fusion our aim is to improve the quality of the information, so that it is, in some sense, *better* than would be possible if the data sources were used individually.

The above definition implies that the sensor data, or the data derived from the sensory data, consists of multiple measurements which have to be combined. The multiple measurements may, of course, be produced by multiple sensors. However, the definition also includes multiple measurements, produced at different time instants, by a single sensor.

The general concept of multi-sensor data fusion is analogous to the manner in which humans and animals use a combination of multiple senses, experience and the ability to reason to improve their chances of survival.

The basic problem of multi-sensor data fusion is one of determining the best procedure for combining the multi-sensor data inputs. The view adopted in this book is that combining multiple sources of information with *a priori* information is best handled within a statistical framework. The main advantage of a statistical approach is that explicit probabilistic models are employed to describe the various relationships between sensors and sources of information taking into account the underlying uncertainties. In particular we restrict ourselves to the Bayesian methodology which provides us with a useful way to formulate the multi-sensor data fusion problem in mathematical terms and which yields an assessment of the uncertainty in all unknowns in the problem.

1.2 Synergy

The principal motivation for multi-sensor data fusion is to improve the quality of the information output in a process known as *synergy*. Strictly speaking, synergy does not require the use of multiple sensors. The reason being that the synergistic effect may be obtained on a temporal sequence of data generated by a single sensor. However, employing more than one sensor may enhance the synergistic effect in several ways, including: increased spatial and temporal coverage, increased robustness to sensor and algorithmic failures, better noise suppression and increased estimation accuracy.

Example 1.1. Multi-Modal Biometric Systems [226]. Biometric systems that rely on a single biometric trait for recognition are often characterized by high error rates. This is due to the lack of completeness or $universality^{[1]}$ in most biometric traits. For example, fingerprints are not truly universal since it is not possible to obtain a good quality fingerprint from people with hand-related disabilities, manual workers with many cuts and bruises on their fingertips or people with very oily or very dry fingers. Multi-modal biometric sensor systems solve the problem of non-universality by fusing the evidence obtained from multiple traits.

Example 1.2. Multiple Camera Surveillance Systems [142]. The increasing demand for security by society has led to a growing need for surveillance activities in many environments. For example, the surveillance of a wide-area urban site may be provided by periodically scanning the area with a single narrow field-of-view camera. The temporal coverage is, however, limited by the time required for the camera to execute one scan. By using multiple cameras we reduce the mean time between scans and thereby increase the temporal coverage. \Box

Broadly speaking, multi-sensor data fusion may improve the performance of the system in four different ways [18]:

- Representation. The information obtained during, or at the end, of the fusion process has an abstract level, or a granuality, higher than each input data set. The new abstract level or the new granuality provides a richer semantic on the data than each initial source of information
- Certainty. If V is the sensor data before fusion and p(V) is the *a priori* probability of the data before fusion, then the gain in certainty is the growth in p(V) after fusion. If V_F denotes data after fusion, then we expect $p(V_F) > p(V)$.
- Accuracy. The standard deviation on the data after the fusion process is smaller than the standard deviation provided directly by the sources. If data is noisy or erroneous, the fusion process tries to reduce or eliminate

¹ A biometric trait is said to be universal if every individual in the target population is able to present the trait for recognition.

noise and errors. In general, the gain in accuracy and the gain in certainty are correlated.

Completeness. Bringing new information to the current knowledge on an environment allows a more complete the view on this environment. In general, if the information is redundant and concordant, we could also have a gain in accuracy

Example 1.3. Multi-Modal Medical Imaging: Gain in Completeness. We consider images obtained through Magnetic Resonance Imaging (MRI), Computed Tomography (CT) and Positron Emission Tomography (PET). The multi-sensor data fusion of all three images allows a surgeon to view "soft tissue" information (MRI) in the context of "bone" (CT) and in the context of "functional" or "physiological information" (PET).

1.3 Multi-Sensor Data Fusion Strategies

As the above examples show, multi-sensor data fusion is a wide-ranging subject with many different facets. In order to understand it better, and to become familiar with its terminology, we shall consider it from three different points of view as suggested by Boudjemaa and Forbes [27], Durrant-Whyte [69] and Dasarathy [56]. Other points of view will be developed in succeeding chapters of the book.

1.3.1 Fusion Type

Boudjemaa and Forbes [27] classify a multi-sensor data fusion system according to what aspect of the system is fused:

- Fusion across sensors. In this situation, a number of sensors nominally measure the same property, as, for example, a number of temperature sensors measuring the temperature of an object.
- Fusion across attributes. In this situation, a number of sensors measure different quantities associated with the same experimental situation, as, for example, in the measurement of air temperature, pressure and humidity to determine air refractive index.
- Fusion across domains. In this situation, a number of sensors measure the same attribute over a number of different ranges or domains. This arises, for example, in the definition of a temperature scale.
- Fusion across time. In this situation, current measurements are fused with historical information, for example, from an earlier calibration. Often the current information is not sufficient to determine the system accurately and historical information has to be incorporated to determine the system accurately.

Example 1.4. Flood Forecasting [302]. Water companies are under constant pressure to reduce the frequency of combined sewer overflows to natural water courses from urban drainage systems (UDS). The management of stormwater through the UDS and similar applications require accurate real-time estimates of the rainfall. One way water companies have done this is to fuse measurements of the rainfall made by ground-based rain gauges and a weather radar system. This is "fusion across sensors" since, nominally, the two sensors measure the same property. \Box

1.3.2 Sensor Configuration

Durrant-Whyte [69] classifies a multi-sensor data fusion system according to its sensor configuration. There are three basic configurations:

- Complementary. A sensor configuration is called complementary if the sensors do not directly depend on each other, but can be combined in order to give a more complete image of the phenomenom under observation. Complementary sensors help resolve the problem of *incompleteness*.
- Competitive. A sensor configuration is competitive if each sensor delivers an independent measurement of the same property. The aim of competitive fusion is to reduce the effects of *uncertain* and *erroneous* measurements.
- Cooperative. A cooperative sensor configuration network uses the information provided by two, or more, independent sensors to derive information that would not be available from the single sensors.

Example 1.5. Triangulation: Cooperative Fusion [272]. The location (x, y) of an object O is found by cooperatively fusing two bearings, θ_1 and θ_2 , as measured by the direction-finding sensors, S_1 and S_2 . Let (X_1, Y_1) and (X_2, Y_2) denote the locations of the two sensors S_1 and S_2 , then by simple geometry (see Fig. 1.1), we find the location (x, y) of O, where

$$x = \frac{Y_2 - Y_1 + (X_1 \tan \theta_1 - X_2 \tan \theta_2)}{\tan \theta_1 - \tan \theta_2} ,$$

$$y = \frac{Y_2 \tan \theta_1 - Y_1 \tan \theta_2 + \tan \theta_1 \tan \theta_2 (X_1 - X_2)}{\tan \theta_1 - \tan \theta_2}$$

Note: See Ex. 4.2 for a Bayesian analysis of this problem.

1.3.3 Input/Output Characteristics

Dasarathy (1994) [56] classifies a multi-sensor data fusion system according to its joint input/output characteristics. Table 1.1 illustrates the types of inputs/outputs considered in this scheme.



Fig. 1.1. Shows the calculation of the location (x, y) of the object *O* by triangulation of the angles θ_1 and θ_2 .

Symbol	Name	Description/Example
DaI-DaO	Data Input/Data	Input data is smoothed/filtered.
DaI-FeO	Data Input/Feature	Features are generated from the input data e.g. edge detection in an image
FeI-FeO	Feature Input/Feature Output	Input features are reduced in number or new features are generated by fusing input fea-
FeI-DeO	Feature Input/Decision	tures. Input features are fused together to give out-
DeI-DeO	Output Decision Input/Decision Output	put decision. Multiple input decisions are fused together to give a final output decision.

Table 1.1. Dasarathy's Input/Output Data Fusion Model.

It is sometimes useful to divide the DeI-DeO fusion model into two sub-models: a "soft" decision-input model (denoted as DsI-DeO) in which each input decision is accompanied by a degree-of-support value and a "hard" decision-input model (denoted as DhI-DeO) in which the input decisions are not accompanied by any degree-of-support values.

1.4 Formal Framework

Multi-sensor data fusion systems are often of a size and a complexity that requires the use of a formal framework [234] around which we may organize our knowledge. The framework adopted in this book is that of a distributed network of autonomous modules, in which each module represents a separate function in the data fusion system. Apart from providing us with a structured framework, the modular design decreases the complexity involved in designing a data fusion system by compartmentalizing the design problem. Modularity also helps in the construction and maintenance of an operational multi-sensor data fusion system.

In analyzing a multi-sensor data fusion system it is useful to divide the system into three parts: the physical, information and cognitive domains and to determine the flow of data between these parts [72, 73, 76, 234] (Fig. 1.2).

- Physical Domain: Hardware. The physical domain contains the sensor modules each of which represents a sensor which physically interacts with the external environment. Each module contains a *sensor model* which provides us with a description of the measurements made by the sensor and of the local environment within which the sensor operates. In some applications we may wish to physically change the external environment. In this case the physical domain will also contain actuators which are able to modify the external environment.
- Information Domain: Software. The information domain constitutes the heart of a multi-sensor data fusion system. It contains three blocks which are responsible for data fusion, control application/resource management and human-machine interface (HMI). The data fusion block is constructed as an autonomous network of "fusion" modules. This network is responsible for combining all the sensor data into a unified view of the environment in the form of an "environmental image". The control application/resource management block is constructed as autonomous networks of "control"



Fig. 1.2. Shows the division of a generic multi-sensor data fusion system into three parts or domains: physical, informative and cognitive and the flow between the parts. The figure is divided into three panels which correspond to the physical domain (left-most panel), information domain (middle panel) and cognitive domain (rightmost panel).

modules. This network is responsible for all decisions which are made on the basis of the environmental image. In many applications the decisions are fed back to the sensor block. In this case the process is known as "sensor management".

Cognitive Domain: Human User. In many multi-sensor data fusion applications the human user is the final arbiter or decision maker. In this case it is important to design the system in such a way that all the information which is transmitted to the human user is transformed into a form which is intuitively usable by the user for his decision-making process.

Example 1.6. Control Loop [72, 73, 76]. Fig. 1.3 shows a control loop containing four blocks: sensor, actuator, data fusion and control application. The environment is observed with one or more sensors. The corresponding sensor observations are then passed to the data fusion block where they are combined to form a unified view of the environment ("environmental image"). The environmental image is, in turn, passed to the control application block. The loop is closed by feeding the decisions made by the control application block back to the environment. \Box

1.4.1 Multi-Sensor Integration

In the data fusion block only a few of the autonomous modules perform "multi-sensor data fusion" as defined in Sect. 1.1, the remaining modules perform auxiliary functions. The modules which perform the data fusion will receive input data from the physical sensors S_1, S_2, \ldots, S_N and from other



Fig. 1.3. Shows a simple control loop built from four blocks: sensors, actuators, data fusion and control application. A feedback mechanism is included by allowing the control application to act on the external environment via an actuator block.

modules. In this case, we sometimes find it convenient to differentiate between multi-sensor data fusion performed by the individual modules and multi-sensor data integration performed by the entire data fusion block [187] (see Fig. 1.4).



Fig. 1.4. Shows two multi-sensor data fusion blocks F_1 and F_2 . F_1 performs data fusion on the output of sensors S_1 and S_2 and F_2 performs data fusion on the output of F_1 and the sensor S_3 . Together F_1 and F_2 perform "multi-sensor integration".

1.5 Catastrophic Fusion

^[2] The unsuspecting reader may conclude, on the basis of what has been presented so far, that a multi-sensor fusion system is *always* better than a single sensor system. This conclusion is, however, mistaken: Often the performance of a multi-sensor data fusion system is below that of the individual sensors. This phenomenom is known as *catastrophic fusion* and clearly should be avoided at all times.

Formally, catastrophic fusion [218] is said to occur when the performance of a multi-sensor data fusion system F is significantly lower than the performance of one, or more, of the individual sensors S_m . In general, each sensor $S_m, m \in \{1, 2, \ldots, M\}$, is designed to operate correctly only under specific conditions, or environment, C_m . Let C_F denote the corresponding condition

² This section contains advanced material. It assumes the reader has some familiarity with Bayesian statistics (see Chapts. 8-10). The section is not, however, essential for what follows and it may be left for a second reading.

for the correct operation of all the sensors in system F, then C_F is the "intersection" of the C_m which we may write symbolically as

$$C_F = C_1 \wedge C_2 \wedge \dots C_M . \tag{1.1}$$

Sometimes, however, F is used in an environment C which is inconsistent with one of the C_m , say C_{m^*} . When this happens, the signal from the corresponding sensor, S_{m^*} , may dominant the fused output with catastrophic results. To prevent this happening, multi-sensor fusion systems often employ *secondary classifiers* which monitor the performance of each sensor S_m .

Example 1.7. Automatic Speech Recognition: Preventing Catastrophic Fusion [218]. In ideal, or clean, conditions automatic speech recognition systems perform very well using a single audio sensor S_A . However, we often observe a substantial reduction in performance when background noise, channel distortion or reverberation are present. This has led to the development of automatic speech recognition systems which employ an audio sensor S_A and a visual sensor S_V (see Ex. 3.11). The two sensors are complementary: speech characteristics that are visually confusable are often acoustically distinct, and characteristics that are acoustically confusable are often visually distinct. The audio-visual system work as follows. Given a finite set of utterances $U_i, i \in \{1, 2, ..., N\}$, we identify an unknown utterance, U, as

$$U^* = \arg\max_{i} \left(P(U_i | y_A, y_V) \right), \qquad (1.2)$$

where $P(U_i|y_A, y_V)$ is the conditional probability of U_i given an audio-visual observation (y_A, y_V) . To a first approximation, the audio and visual features are conditionally independent. In this case, we may write $P(U_i|y_A, y_V)$ as

$$P(U_i|y_A, y_V) \propto P(U_i|y_A) \times P(U_i|y_V) , \qquad (1.3)$$

where the conditional probabilities $P(U_i|y_A)$ and $P(U_i|y_V)$ are calculated offline using a set of training samples D.

If the set of training samples and test samples are well matched, the solution represented by (1.2) and (1.3) is theoretically optimal: we automatically compensate for any noise or distortion and assign more importance to the classification which is more "certain". The underlying assumption is, however, that the conditional probabilities $P(U_i|y_A)$ and $P(U_i|y_V)$ generated during training, match the speech data that is being tested. When the speech data is contaminated with noise this assumption is no longer valid and the probability estimates are incorrect. In such cases, it is common to use a weighted integration scheme, where the influence of the noisy channel is attenuated. This weighting can be a function of the signal-to-noise ratio (SNR) in each channel, and can be implemented as follows:

$$P(U_i|y_A, y_V) \propto P^{\alpha_A}(U_i|y_A)P^{\alpha_V}(U_i|y_V)$$
,

where

$$\alpha_A + \alpha_V = 1 \; .$$

The purpose of the weights $\alpha_m, m \in \{A, V\}$, is to "neutralize" a sensor S_m , whenever the secondary classifiers determine that S_m is operating outside its specified operating conditions C_m . In the worst case, when environmental conditions make the sensor completely unreliable we set the corresponding weight to zero.

1.6 Organization

Although we shall discuss the physical, information and cognitive domains in a multi-sensor data fusion system, the emphasis will be on the data fusion block. The book is organized into five parts as follows.

- Part I: Introduction. This consists of Chapts. 1-3. Chapt. 1 is a general introduction to the subject of multi-sensor data fusion. In this chapter we provide an overview of the basic concepts used in multi-sensor data fusion. In Chapt. 2 we consider the sensors, which ultimately, are the source of all input data into a multi-sensor data fusion system. In Chapt. 3 we consider the different system architectures which may be employed in constructing a data fusion system.
- Part II: Common Representational Format. This consists of Chapts. 4-7. In Chapt. 4 we provide an overview of the basic concept of a common representational format and the different techniques used to create such a representation. The techniques may be broadly divided into three different types which are then discussed in turn in Chapts. 5, 6 and 7.
- Part III: Data Fusion. This consists of Chapts. 8-13. In Chapt. 8 we give an overview of the Bayesian approach to multi-sensor data fusion and the different techniques involved. This is followed by Chapts. 9-11 in which we deal with parameter estimation theory including sequential estimation theory. In Chapts. 12-13 we consider decision fusion.
- Part IV. Sensor Management. This consists of a single chapter (Chapt. 14) in which we consider sensor management. Specifically we consider how the decisions made by the data fusion block may, if required, be fed back to the sensors.
- Part V. Appendices. This consists of Appendices A and B. Appendix A is a list of relevant software written in matlab which is available on the world wide web. Sufficient information is provided in the table so that the software may be easily found on the world wide web using a simple search machine ^[3]. Appendix B is a summary of elementary results in probability theory, linear algebra and matrix theory with which the reader should be familiar.

³ The internet addresses themselves are not given for the simple reason that internet adresses tend to have a very short timelife.

1.7 Further Reading

General overviews on multi-sensor data fusion are [3, 6, 7, 108, 109, 131, 154, 187, 318]. For an extended discussion regarding the issues involved in defining multi-sensor data fusion and related terms, see [316, 236].

Sensors

2.1 Introduction

Sensors are devices which interact *directly* with the environment and which are ultimately the source of all the input data in a multi-sensor data fusion system [87]. The physical element which interacts with the environment is known as the *sensor element* and may be any device which is capable of perceiving a physical property, or environmental attribute, such as heat, light, sound, pressure, magnetism or motion. To be useful, the sensor must map the value of the property or attribute to a quantitative measurement in a *consistent* and *predictable* manner.

In Chapt. 1 we introduced a formal framework in which we represented a sensor fusion system as a distributed system of autonomous modules. To support such a scheme, the sensors must not only measure a physical property, but must also perform additional functions. These functions can be described in terms of compensation, information processing, communication and integration:

- Compensation. This refers to the ability of a sensor to detect and respond to changes in the environment through self-diagnostic tests, self-calibration and adaption.
- Information Processing. This refers to processes such as signal conditioning, data reduction, event detection and decision-making, which enhance the information content of the raw sensor measurements.
- Communications. This refers to the use of a standardized interface and a standardized communication protocol for the transmission of information between the sensor and the outside world.

Integration. This refers to the coupling of the sensing and computation processes on the same silicon chip. Often this is implemented using microelectro-mechanical systems (MEMS) technology.

A practical implementation of such a sensor is known as a smart, or intelligent, sensor [77].

2.2 Smart Sensor

A *smart sensor* is a hardware/software device that comprises in a compact small unit a sensor element, a micro-controller, a communication controller and the associated software for signal conditioning, calibration, diagnostic tests and communication. The smart sensor transforms the raw sensor signal to a standardized digital representation, checks and calibrates the signal, and transmits this digital signal to the outside world via a standardized interface using a standardized communication protocol.

Fig. 2.1 shows the measurement of a physical property by a smart sensor.

The transfer of information between a smart sensor and the outside world is achieved by reading (writing) the information from (to) an interface-file system (IFS) which is encapsulated in the smart sensor as shown in Fig. 2.2 [75].



Fig. 2.1. The sensor element measures the physical property and outputs an analog signal which is amplified, filtered and then converted to a digital signal by the analog-to-digital, or A/D, unit. The digital signal is processed by the microprocessor, μP , where it is temporally stored before being transmitted by the transmitter/receiver. *Note*: The smart sensor may also receive command instructions via the transmitter/receiver unit.



Fig. 2.2. Shows a smart sensor with a sensor element and the encapsulated signal processing functions and the encapsulated Interface file system (IFS).

2.3 Logical Sensors

A logical sensor [116] is defined as any device which functions as a source of information for a multi-sensor data fusion node. Thus a logical sensor encompass both physical sensors and any fusion node whose output is subsequently fed into another fusion node. Unless stated otherwise, from now on the term "sensor" will refer to a logical sensor or a source-of-information.

2.4 Interface File System (IFS)

The IFS provides a structured (name) space which is used for communicating information between a smart sensor and the outside world [74]. The following example neatly illustrates the concept of an IFS.

Example 2.1. Brake Pedal in a Car: An Interface File System [160]. We consider a driver in a car. For the purposes of braking, the brake pedal acts as an IFS between the driver and the brakes. At this interface there are two relevant variables. The first variable is P: the pressure applied to the brake pedal by the driver and the second variable is R: the resistance provided by the brakes back to the driver. The relative position of the brake pedal uniquely identifies the values of P and R to both the driver and the brakes. The temporal association between sensing the information (e. g. by the driver) and receiving the information (e. g. by the brakes) is implicit, because of the mechanical connection between the driver and the brakes.

A record of the IFS can be accessed, both from the sensor and from the outside world (Fig. 2.2). Whenever one of the internal processes of a smart

sensor requires information from the outside world or produces information for the outside world, it accesses the appropriate records of the IFS and reads (writes) the information to (from) this record. The internal processes of a smart sensor are thus not visible to the outside world.

Often we implement the IFS so that it acts as a temporal firewall between the sensors and the outside world. In this case, the IFS uses *local* interface file systems that can be written by the sensor and read by the outside world, without having direct communication between the sensor and the outside world (Fig. 2.3).

2.4.1 Interface Types

In the smart sensor model we distinguish between three interface types: the real-time service interface, the diagnostic and maintenance interface and the configuration and planning interface. All information that is exchanged across these interfaces is stored in files of the IFS.

- Real-time Service (RS) Interface. This interface provides time sensitive information to the outside world. The communicated data is sensor observations made by a sensor on real-time entities in the environment. The RS accesses the real-time IFS files within the smart sensor which is usually time-critical.
- Diagnostic and Management (DM) Interface. This interface establishes a connection to each smart sensor. Most sensors need both parameterization and calibration at start-up and the periodic collection of diagnostic information to support maintenance activities. The DM interface accesses the diagnostic IFS files within the smart sensor which is not usually timecritical.
- Configuration and Planning (CP) Interface. This interface is used to configure a smart sensor for a given application. This includes the integration



Fig. 2.3. Shows the Interface file system (IFS) built as a temporal firewall. The IFS contains two separate local interface file systems: one is accessed by the sensor and the other is accessed by the outside world.