

Modeling and Simulation in Science, Engineering and Technology

Series Editor

Nicola Bellomo
Politecnico di Torino
Italy

Advisory Editorial Board

M. Avellaneda (Modeling in Economics)
Courant Institute of Mathematical Sciences
New York University
251 Mercer Street
New York, NY 10012, USA
avellaneda@cims.nyu.edu

K.J. Bathe (Solid Mechanics)
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
kjb@mit.edu

P. Degond (Semiconductor and Transport Modeling)
Mathematiques pour l'Industrie et la Physique
Universit P. Sabatier Toulouse 3
118 Route de Narbonne
31062 Toulouse Cedex, France
degond@mip.ups-tlse.fr

A. Deutsch (Complex Systems
in the Life Sciences)
Center for Information Services
and High Performance Computing
Technische Universitt Dresden
01062 Dresden, Germany
andreas.deutsch@tu-dresden.de

M.A. Herrero Garcia (Mathematical Methods)
Departamento de Matematica Aplicada
Universidad Complutense de Madrid
Avenida Complutense s/n
28040 Madrid, Spain
herrero@sunma4.mat.ucm.es

W. Kliemann (Stochastic Modeling)
Department of Mathematics
Iowa State University
400 Carver Hall
Ames, IA 50011, USA
kliemann@iastate.edu

H.G. Othmer (Mathematical Biology)
Department of Mathematics
University of Minnesota
270A Vincent Hall
Minneapolis, MN 55455, USA
othmer@math.umn.edu

L. Preziosi (Industrial Mathematics)
Dipartimento di Matematica
Politecnico di Torino
Corso Duca degli Abruzzi 24
10129 Torino, Italy
luigi.preziosi@polito.it

V. Protopopescu (Competitive Systems,
Epidemiology)
CSMD
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6363, USA
vvp@epmnas.epm.ornl.gov

K.R. Rajagopal (Multiphase Flows)
Department of Mechanical Engineering
Texas A&M University
College Station, TX 77843, USA
krajagopal@mengr.tamu.edu

Y. Sone (Fluid Dynamics in Engineering Sciences)
Professor Emeritus
Kyoto University
230-133 Iwakura-Nagatani-cho
Sakyo-ku Kyoto 606-0026, Japan
sone@yoshio.mbox.media.kyoto-u.ac.jp

Data Modeling for Metrology and Testing in Measurement Science

Franco Pavese
Alistair B. Forbes
Editors

Birkhäuser
Boston • Basel • Berlin

Franco Pavese
Istituto Nazionale di Ricerca
Metrologica, Torino, Italy
frpavese@tin.it
f.pavese@inrim.it

Alistair B. Forbes
National Physics
Laboratory, Middlesex, UK
alistair.forbes@npl.co.uk

ISBN 978-0-8176-4592-2 e-ISBN 978-0-8176-4804-6
DOI 10.1007/978-0-8176-4804-6

Library of Congress Control Number: 2008938171

Mathematics Subject Classification (2000): 60–80, 60K10, 60K40, 60K99, 62-XX, 65C50.

© Birkhäuser Boston, a part of Springer Science+Business Media, LLC 2009

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Birkhäuser Boston, c/o Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper.

www.birkhauser.com

Preface

The aim of this book is to provide, firstly, an introduction to probability and statistics especially directed to the metrology and testing fields and secondly, a comprehensive, newer set of modelling methods for data and uncertainty analysis that are generally not considered yet within mainstream methods. The book brings, for the first time, a coherent account of these newer methods and their computational implementation. They are potentially important because they address problems in application fields where the usual hypotheses that are at the basis of most of the traditional statistical and probabilistic methods, for example, relating to normality of the probability distributions, are frequently not fulfilled to such an extent that an accurate treatment of the calibration or test data using standard approaches is not possible. Additionally, the methods can represent alternative ways of data analysis, allowing a deeper understanding of complex situations in measurement. The book lends itself as a possible textbook for undergraduate or postgraduate study in an area where existing texts focus mainly on the most common and well-known methods that do not encompass modern approaches to calibration and testing problems.

The book is structured in such a way to guide readers with only a general interest in measurement issues through a series of review papers, from an initial introduction to modelling principles in metrology and testing, to the basic principles of probability in metrology and statistical approaches to uncertainty assessment. Later chapters provide a survey of the newer methods, from an introduction to the alternative approach of interval mathematics to the latest developments in data analysis using least squares, FFT, wavelets, and fuzzy methods; from data fusion (including decision taking and risk analysis), to tools for combining data of complex statistical structure; and from uncertainty issues related to model imperfection, to those related to combining testing data. The book also includes chapters on modern computational issues related to measurement: a computer-assisted simplified rigorous approach to data evaluation, an analysis of the strategies to adopt for measurement software validation, an introduction to the virtual instrument approach, and an

overview of the main IT applications in metrology. The book does not concentrate on any particular field of application, because the applications in the frames of metrology and testing cover so broad a range that it would be difficult to make a ranking of their importance or even to attempt a grouping into categories with homogeneous needs. On the other hand, most of the various techniques illustrated in the chapters of the book can find application to many different issues related to these application fields.

A DVD is attached to the book, containing software for free use (under the specified conditions), ranging from tutorials to sample codes of the implementation of methods described in the book, to software packages with demos of methods and tools, allowing the reader to try to see especially the newer tools at work with the minimum effort, without the need of implementing his or her own code.

The authors are mainly selected from an international collaborative framework (<http://www.imeko-tc21.org>, <http://www.imeko.org>), established in the early 1990s as ‘AMCTM’ (<http://www.amctm.org>), that has allowed a community of metrologists, mathematicians, statisticians, and software/IT engineers to work together, so creating a common understanding of the issues discussed in this book.

F. Pavese, Istituto Nazionale di Ricerca Metrologica
A B Forbes, National Physical Laboratory
Torino
June 2008

Contents

Preface	V
List of Contributors	XV
An Introduction to Data Modelling Principles in Metrology and Testing	
<i>Franco Pavese</i>	1
1 Introduction	1
2 Uncertainty components of the measurement process: Repeatability, reproducibility, accuracy	2
2.1 Basic nomenclature	3
3 The GUM approach to the measurement process: Type A and Type B components of uncertainty	8
3.1 Basic nomenclature	9
4 Other approaches to errors in the measurement process	11
4.1 The total error and its shared and specific components	12
4.2 A distinction between measurement and measurand	12
5 Data modelling in metrology and in testing for <i>intralaboratory</i> measurements	13
5.1 Repeated measurements	13
5.2 Nonrepeated measurements on the same standard within each laboratory	14
6 Data modelling in metrology and in testing for <i>interlaboratory</i> measurements (<i>intercomparisons</i>); comparison specific problems	15
6.1 Data modelling	16
6.2 Specific problems and outcomes of the <i>intercomparisons</i> , namely the MRA key comparisons	19
References	24

Probability in Metrology

<i>Giovanni B. Rossi</i>	31
1 Probability, statistics, and measurement – An historical perspective	31
1.1 The origins: Gauss, Laplace, and the theory of errors	31
1.2 Orthodox statistics	36
1.3 The Guide to the Expression of Uncertainty in Measurement	42
1.4 Issues in the contemporary debate on measurement uncertainty	44
2 Towards a probabilistic theory of measurement	55
2.1 Origin and early development of the formal theory of measurement	55
2.2 The representational theory of measurement	59
2.3 A probabilistic theory of measurement	63
3 Final remarks	67
References	69

Three Statistical Paradigms for the Assessment and Interpretation of Measurement Uncertainty

<i>William F. Guthrie, Hung-kung Liu, Andrew L. Rukhin, Blaza Toman, Jack C. M. Wang, Nien-fan Zhang</i>	71
1 Introduction	71
1.1 Notation	74
1.2 Statistical paradigms	74
1.3 Examples	76
2 Frequentist approach to uncertainty assessment	77
2.1 Basic method	77
2.2 Example 1	84
3 Bayesian paradigm for uncertainty assessment	87
3.1 Basic method	87
3.2 Example 1	89
4 Fiducial inference for uncertainty assessment	93
4.1 Basic method	93
4.2 Example 1	96
5 Example 2	100
5.1 Frequentist approach	101
5.2 Bayesian approach	104
5.3 Fiducial approach	107
6 Discussion	109
7 Chapter summary	112
References	114

Interval Computations and Interval-Related Statistical Techniques: Tools for Estimating Uncertainty of the Results of Data Processing and Indirect Measurements

Vladik Kreinovich 117

1 Importance of data processing and indirect measurements 117

2 Estimating uncertainty for the results of data processing and indirect measurements: An important metrological problem 119

3 Uncertainty of direct measurements: Brief description, limitations, need for overall error bounds (i.e., interval uncertainty) 119

4 Data processing and indirect measurements under interval uncertainty: The main problem of interval computations 120

5 Uniform distributions: Traditional engineering approach to interval uncertainty 121

6 Techniques for estimating the uncertainty of the results of indirect measurements in situations when the measurement errors of direct measurements are relatively small 123

7 Techniques for error estimation in the general case of interval uncertainty 127

8 Situations when, in addition to the upper bounds on the measurement error, we also have partial information about the probabilities of different error values 136

9 Final Remarks 143

References 144

Parameter Estimation Based on Least Squares

Methods

Alistair B. Forbes 147

1 Introduction 147

2 Model fitting in metrology 148

3 Linear least squares problems (LS) 148

 3.1 Orthogonal factorisation method to determine parameter estimates 149

 3.2 Minimum variance property of the least squares estimate 150

 3.3 Linear Gauss–Markov problem (GM) 151

 3.4 Generalised QR factorisation approach to the Gauss–Markov problem 152

 3.5 Linear least squares, maximum likelihood estimation, and the posterior distribution $p(\mathbf{a}|\mathbf{y})$ 152

4 Nonlinear least squares (NLS) 154

 4.1 The Gauss–Newton algorithm for nonlinear least squares 154

 4.2 Approximate uncertainty matrix associated with \mathbf{a}_{NLS} 155

 4.3 Nonlinear Gauss–Markov problem (NGM) 156

 4.4 Approximate uncertainty matrix associated with \mathbf{a}_{NGM} 156

 4.5 Nonlinear least squares, maximum likelihood estimation, and the posterior distribution $p(\mathbf{a}|\mathbf{y})$ 156

5	Exploiting structure in the uncertainty matrix	158
5.1	Structure due to common random effects, linear case	158
5.2	Structure due to common random effects, nonlinear case	160
6	Generalised distance regression (GDR)	161
6.1	Algorithms for generalised distance regression	162
6.2	Sequential quadratic programming for the footpoint parameters	164
6.3	Example application: Response and evaluation calibration curves	166
6.4	Example: GDR line	167
7	Generalised Gauss–Markov regression (GGM)	167
7.1	Structured generalised Gauss–Markov problems	168
8	Robust least squares (RLS)	169
8.1	Empirical implementation of RLS	171
8.2	One-sided RLS	171
8.3	RLS and the Huber M-estimator	172
8.4	Algorithms for robust least squares	173
9	Summary and concluding remarks	174
	References	174

**Frequency and Time–Frequency Domain Analysis Tools
in Measurement**

	<i>Pedro M. Ramos, Raul C. Martins, Sergio Rapuano,</i> <i>Pasquale Daponte</i>	177
1	Introduction	177
2	Fourier Analysis	179
3	How to use the FFT	181
3.1	Aliasing	182
3.2	Spectral leakage	182
3.3	Windowing	183
4	Example of spectral analysis application using FFT	184
5	Wavelet transform	187
6	The wavelet transform: Theoretical background and implementation	188
6.1	Continuous wavelet transform	189
6.2	Discrete time wavelet transform	193
7	Chirplet transform	199
8	Wavelet networks	201
	References	202

**Data Fusion, Decision-Making, and Risk Analysis:
Mathematical Tools and Techniques**

	<i>Pedro S. Girão, Octavian Postolache, José M. D. Pereira</i>	205
1	Data fusion	205

- 1.1 Definitions, concepts, and terms of reference (terminology).
Processing and topological issues 206
- 1.2 Data fusion models and architectures 207
- 1.3 Data fusion techniques and algorithms 208
- 1.4 Applications 219
- 1.5 Implementing data fusion 239
- 2 Decision-making 242
 - 2.1 Special decision-making tools and techniques 243
 - 2.2 Final notes on decision-making 247
- 3 Risk analysis 248
- References 250

Comparing Results of Chemical Measurements: Some Basic Questions from Practice

- Paul De Bièvre* 255
- 1 Introduction 255
- 2 Why do we compare measurement results? 256
 - 2.1 Comparison of two measurement results in chemistry 256
 - 2.2 Comparison of many measurement results in chemistry 265
- 3 Final remarks 272
- References 272

Modelling of Measurements, System Theory, and Uncertainty Evaluation

- Klaus-D. Sommer* 275
- 1 Introduction to the modelling of measurements 275
 - 1.1 Modelling tasks in uncertainty evaluation 275
 - 1.2 Basic categories of models 276
 - 1.3 Models for representing measurements and for evaluating the measurement uncertainty 277
- 2 Describing systems and signals 279
 - 2.1 Linear time-invariant systems 280
 - 2.2 Linearity of measuring systems 281
 - 2.3 Nonlinear systems 282
 - 2.4 Time-variant signals and systems 283
 - 2.5 Description of stochastic signals 286
- 3 Systematic approach to the modelling of measurements 288
 - 3.1 Standard modelling components 288
 - 3.2 Stepwise modelling procedure 290
- 4 Effects of imperfectly modelled measurements 292
 - 4.1 Disregarding nonlinearities 292
 - 4.2 Dynamic uncertainty contributions 293
- References 297

Approaches to Data Assessment and Uncertainty**Estimation in Testing**

<i>Wolfram Bremser, Werner Hässelbarth</i>	299
1 Introduction	299
2 Assessment and combination of quantitative results	300
2.1 Univariate case: Approaches and tools for two measurements	300
2.2 Univariate case: Approaches and tools for multiple measurements	304
2.3 Multivariate case: Approaches and tools	312
3 Assessment and combination of qualitative results	319
3.1 Approaches and tools	320
3.2 Uncertainty estimation	325
4 Some final remarks	325
References	326

Monte Carlo Modeling of Randomness

<i>Alan G. Steele, Robert J. Douglas</i>	329
1 Introduction	329
1.1 Strengths of Monte Carlo methods	331
2 Distributions for uncertainty claims	332
2.1 Gaussian or Normal distributions	332
2.2 Student- <i>t</i> distributions	333
2.3 Rectangular distribution	334
2.4 U-shaped distribution	334
3 General distributions	334
3.1 General distributions by moments	334
3.2 General distributions by tables	335
3.3 Skewed general distributions	336
4 Multivariate distributions	336
4.1 Covariance: Simplifications – Type A	337
4.2 Covariance: Simplifications – Type B	337
4.3 Covariances in building measurement equations	338
4.4 Example: Covariance matrix to perturbation vector	338
5 Pseudo-random number generators	339
5.1 Underlying pseudo-random number generators	340
5.2 Compact algorithms	340
5.3 Demonstrable randomness	341
5.4 Testing with pseudo-random number generators	341
5.5 Tests of pseudo-random number generators	342
5.6 Tests on your computer	342
5.7 Tests of dynamic range	343
6 Algorithms for generating specific distributions	345
6.1 Gaussian	345
6.2 Student- <i>t</i>	346
6.3 Uniform distribution	347

6.4 U-distribution 348
 6.5 General PRNG methods – for the U-distribution 349
 6.6 General PRNG Methods for Tabulated Distributions 350
 6.7 Choosing the number of simulated events 351
 7 Examples of Monte Carlo simulation 353
 7.1 Gaussian PDFs 353
 7.2 Non-Gaussian PDFs and measurement equations 356
 8 Monte Carlo simulation for hypothesis testing 363
 8.1 Chi-squared testing of the common mean hypothesis 364
 8.2 Chi-squared-like testing of the common mean hypothesis 365
 8.3 Chi-squared-like testing of the agreement hypothesis 366
 9 Final remarks 367
 References 368

Software Validation and Preventive Software Quality Assurance for Metrology

Norbert Greif, Dieter Richter 371
 1 Introduction: The risks of software in metrology 371
 1.1 What is software 372
 1.2 Software is different 373
 1.3 Tackling software issues in metrology 376
 2 Standards and guidelines related to software in metrology 376
 2.1 Approaches to software quality and systematics of standards . . 377
 2.2 Standards and guidelines concerning process quality 379
 2.3 Standards and guidelines concerning product quality 380
 2.4 Laboratory and metrology related standards and guidelines . . 383
 2.5 Future work 387
 3 Analytical and preventive software quality assurance 387
 3.1 Product testing versus preventive quality assurance 388
 3.2 Definition of testable requirements 389
 3.3 Software product testing 391
 3.4 Execution of software product tests 396
 3.5 Process assessment and improvement 396
 3.6 Future work 399
 4 Software validation in metrology 399
 4.1 Software-oriented interpretation of the standard
 ISO/IEC 17025 400
 4.2 Systematisation of software requirements 400
 4.3 Methods of software validation 407
 5 Outlook: Towards a GUM-like validation guideline 408
 References 409

Virtual Instrumentation

Octavian Postolache, Pedro S. Girão, José M. D. Pereira 413
 1 Introduction 413
 2 Hardware 416

2.1	Single unit based virtual instruments – Multifunction I/O boards.....	417
2.2	Multiunit based virtual instruments	422
3	Software	425
3.1	Textual programming languages	426
3.2	Graphical programming languages	433
3.3	VISA and IVI technologies	439
4	Virtual instrumentation operation	444
4.1	Virtual laboratories	444
	References	449

Internet-Enabled Metrology

	<i>Tanasko Tasić</i>	453
1	Introduction	453
2	Functionality of measuring instruments (distributed measuring systems)	455
2.1	Legal metrology applications	455
2.2	Remote operation of measuring instruments	458
3	Internet-supported metrological services	459
3.1	General metrology related services	459
3.2	Internet enabled/supported calibrations	460
3.3	Availability of specific metrological software validation services	464
4	Availability of the metrology-related data	466
5	Final Remarks	468
	References	469

DVD Content

	DVD contents	473
	Additional material to the chapters	473
	Reports and Guidelines	473
	Measurement Software	473
	NIST Engineering Statistical Handbook	474
	Monte Carlo Simulation	474
	Statistical Software	474

Index

	475
--	-------	-----

List of Contributors

Wolfram Bremser

Federal Institute for Materials
Research and Testing (BAM)
WG Gas Analysis, Metrology
Unter den Eichen 87
Berlin, Germany, D-12205
wolfram.bremser@bam.de

Pasquale Daponte

Università del Sannio, Department
of Engineering
Benevento, Italy, I-82100
daponte@unisannio.it

Paul De Bièvre

Independent Consultant
on Metrology in Chemistry
Duineneind 9
Kasterlee, Belgium, B-2460
paul.de.bievre@skynet.be

Robert J. Douglas

National Research Council
1200 Montreal Road
Ottawa, Canada, K1A 0R6
robert.douglas@nrc-cnrc.gc.ca

Alistair B. Forbes

National Physical Laboratory
Hampton Road Teddington,
Middlesex, UK TW11 0LW
alistair.forbes@npl.co.uk

Pedro S. Girão

Instituto de Telecomunicações,
Instituto Superior Técnico
Av. Rovisco Pais 1
Lisbon, Portugal, 1049-001
psgirao@ist.utl.pt

Norbert Greif

Physikalisch-Technische Bundes-
anstalt
Abbestraße 2-12
Berlin, Germany, D-10587
norbert.greif@ptb.de

William F. Guthrie

National Institute of Standards and
Technology (NIST)
100 Bureau Drive
Gaithersburg, MD, USA, 20899
guthrie@nist.gov

Werner Hässelbarth

Federal Institute for Materials
Research and Testing (BAM),
WG “Gas Analysis, Metrology”
Unter den Eichen 87
Berlin, Germany, D-12205
werner.haesselbarth@bam.de

Hung-kung Liu

National Institute of Standards and
Technology (NIST)
100 Bureau Drive
Gaithersburg, MD, USA, 20899
liu@nist.gov

Vladik Kreinovich

University of Texas at El Paso,
Department of Computer Science
500 W. University of Texas
El Paso, TX 79968, USA
vladik@utep.edu

Raul C. Martins

Instituto de Telecomunicações,
Instituto Superior Técnico, DEEC,
UTL
Av. Rovisco Pais 1
Lisbon, Portugal, 1049-001
rcmartins@ist.utl.pt

Franco Pavese

Istituto Nazionale di Ricerca
Metrologica (INRIM)
strada delle Cacce 73-91
Torino, Italy, I-10135
f.pavese@inrim.it
frpavese@tin.it

José M. D. Pereira

Instituto de Telecomunicações,
Escola Superior de Tecnologia, IPS
Setúbal, Portugal
joseper@est.ips.pt

Octavian Postolache

Instituto de Telecomunicações,
Instituto Superior Técnico
Av. Rovisco Pais 1
Lisbon, Portugal, 1049-001
octavian.postolache@ist.utl.pt

Pedro M. Ramos

Instituto de Telecomunicações,
Instituto Superior Técnico, DEEC,
UTL
Av. Rovisco Pais 1
Lisbon, Portugal, 1049-001
pedro.ramos@lx.it.pt

Sergio Rapuano

Università del Sannio,
Department of Engineering
Benevento, Italy, I-82100
rapuano@unisannio.it

Dieter Richter

Physikalisch-Technische Bundes-
anstalt
Abbestraße 2-12
Berlin, Germany, D-10587
dieter.richter@ptb.de

Giovanni B. Rossi

Università degli Studi di Genova
DIMEC
via all'Opera Pia 15A
Genova, Italy, I-16145
gb.rossi@dimec.unige.it

Andrew L. Rukhin

National Institute of Standards and
Technology (NIST)
100 Bureau Drive
Gaithersburg, MD, USA, 20899
andrew.rukhin@nist.gov

Klaus D. Sommer

Klaus-Dieter Sommer
Physikalisch-Technische Bundes-
anstalt
Chemical Physics and Explosion
Protection
Bundesallee 100
Braunschweig, Germany, D-38116
klaus-dieter.sommer@ptb.de

Alan G. Steele

National Research Council
1200 Montreal Road
Ottawa, Canada, K1A 0R6
alan.steele@nrc-cnrc.gc.ca

Tanasko Tasić

METREL d.d.
Ljubljanska cesta 77
Horjul, Slovenia, 1354
tanasko.tasic@metrel.si

Blaza Toman

National Institute of Standards and
Technology (NIST)
100 Bureau Drive
Gaithersburg, MD, USA, 20899
toman@nist.gov

Jack C. M. Wang

National Institute of Standards and
Technology (NIST)
325 Broadway
Boulder, CO, USA, 80305
jwang@boulder.nist.gov

Nien fang Zhang

National Institute of Standards and
Technology (NIST)
100 Bureau Drive
Gaithersburg, MD, USA, 20899
zhang@nist.gov

**and, Authors contributing
only to the contents
of the DVD**

(see Appendix at the end of the
book):

Patrizia Ciarlini

CNR –Istituto per le Applicazioni
del Calcolo “M. Picone”
viale del Policlinico 137
Roma, Italy, I-00161
p.ciarlini@iac.cnr.it

Francesco Crenna

Università degli Studi di Genova
DIMEC
via all’Opera Pia 15A
Genova, Italy, I-16145
crenna@dimec.unige.it

Andrzej Hetman

Technical University of Lodz
Stefanowskiego 18/22
Lodz, Poland, 90-924
jerzykor@p.lodz.pl
andrzejh@p.lodz.pl

Daniela Ichim

ISTAT
Roma, Italy, I-00100
ichim@istat.it

Rüdiger Kessel

Metrodata GmbH
Im Winkel 15-1
Weil am Rhein, Germany, D-79576
r.kessel@metrodata.de

Jerzy Korczynski

Technical University of Lodz
Stefanowskiego 18/22
Lodz, Poland, 90-924
jerzykor@p.lodz.pl

Giuseppe Regoliosi

CNR –Istituto per le Applicazioni
del Calcolo “M. Picone”
viale del Policlinico 137
Roma, Italy, I-00161
g.regoliosi@iac.cnr.it

An Introduction to Data Modelling Principles in Metrology and Testing

Franco Pavese

Istituto Nazionale di Ricerca Metrologica (INRIM), strada delle Cacce 73-91,
10135 Torino, Italy
f.pavese@inrim.it, frpavese@tin.it

Summary. The dispersion of the measured data introduces an uncertainty in the measure of the observed phenomena. Uncertainty associated with data is specified according to models that are different according to the underpinning assumptions, which must adequately match the characteristic of the observed phenomena or process. This chapter deals with the different types of description of the uncertainty components, with a wide selection of citations from reference international documents, and then with the different models corresponding to the different data characteristics. An extended bibliography is included.

Key words: Data modelling, uncertainty, repeatability, reproducibility, accuracy, Type A uncertainty component, Type B uncertainty component, mixture model

1 Introduction

The modern scientific method requires that experimental evidence be obtained whenever possible, for the validity of a theory to be considered substantiated.

Experimental evidence consists of a measure either quantitative or non-quantitative – often called ‘qualitative’ – of the observable quantities and this is obtained in general by means of measurement.

The degree of consistency of different measurement results, obtained by different independent experimenters or by the same experimenter at different times, is considered in general to provide a measure of the degree of reliability that can be associated to the results in representing the quantity under study, having taken into account the fact that experimental knowledge is always imperfect to some degree.

Consequently, replication of measurements¹ and the combination of observations are standard and essential practices in science.

¹ The term “replicated measurements” is used in this chapter to indicate, in a general way, the “determination of a value more than once” [ISO93]. The term ‘repeated’ has a specific statistical meaning and potential confusion should be avoided. In fact, replicated measurements can be either ‘repeated’ or

This requirement places the need in all branches of ‘exact’ science to evaluate the degree to which different observations can correctly or safely be compared (or combined) with each other, or, in other words, to assess the traceability of the measurements performed by different experimenters and at different times.

Metrology is that part of measurement science whose task is to provide a measure of traceability and is concerned with the measured values of the observed quantities (physical, chemical, biological, etc.); this is achieved by providing reference measurement methods and traceable standards for each quantity.

Testing is a contiguous area – in some respects overlapping with metrology – devoted to provide and apply reference methods in order to obtain a measure of the characteristics of ‘stuff’ (objects, devices, goods, foodstuffs, etc.), in general in order to allow one to take a decision, such as assessing conformity to specifications or to limits, often within a legal context.

An underpinning basic concept of science, and hence of measurement science, is that, due to imperfect knowledge of the observed phenomena, the numerical data² that are the outcomes of measurement are affected by errors. Irrespective of the reasons that are the causes of these errors, the resulting dispersion of the measured numerical values that is generally observed³ is interpreted as evidence of the imperfect knowledge.

Thus, the dispersion of the measured values introduces an uncertainty in the measure of the observed phenomena. Uncertainty associated with data is specified according to models that are different according to the underpinning assumptions, which must adequately match the characteristic of the observed phenomena or measurement process.

2 Uncertainty components of the measurement process: Repeatability, reproducibility, accuracy

In evaluating the uncertainty associated with measurement results in metrology (and testing), several steps can be enumerated, each characterised by the use of different methods to fulfil correspondingly different purposes [PF06].

For a measurement process entirely taking place only *within one laboratory*, the purpose of measurement replication:

‘nonrepeated’ depending on the conditions. See, for example, the statement “To verify control of precision, the laboratory may perform a number of replicate measurements under repeatability conditions” in [A2LA02].

² The data associated to an observation can also be nonnumerical, for example, lexical. In this chapter only numerical data are considered.

³ For a sufficiently high measurement sensitivity.

- (a) When performed on the same measurement standard, is primarily to obtain statistical information providing a measure of the *repeatability* of the measured values of the standard;
- (b) When performed on the same measurement standard, is secondly to evaluate the increase in the total uncertainty arising from the variability of the influence quantities affecting the standard, including those that have a dependence on time, that is, to have a measure of the *reproducibility* of the measured values of the standard;
- (c) When performed on several measurement standards of the laboratory, is finally to assess whether they have the same value or to provide a measure of the (systematic) differences between their measured values, and to evaluate the associated uncertainty, that is, to provide an estimate of the *accuracy* of the measured values of the laboratory standards. This step is called *intralaboratory* comparison.

When operation (c) is performed by directly comparing one (or more) measurement standards provided by different laboratories, so that it is part of a process taking place *between at least two laboratories*, it is then called an *interlaboratory* comparison.

When operation (c) is performed to assess “periodically the overall performance of a laboratory” [EA03], that is, to show that the laboratory can continue to demonstrate its ability to conduct correctly a certain type of measurement, it should be considered and used as a *proficiency test* (more details in [DP06]).

The term ‘repeated measurements’, involved in step (a), is long established in general statistics: it refers to homoscedastic data (recent uses can be found in [Dav02, LW06, LW06b]). In this respect, the latest statistical status of the measurements performed for the different purposes (a) to (c) listed above has been treated and defined in several reference documents in metrology and testing since the 1980s. In this respect, some evolution of the concepts can be observed. It is therefore useful to summarise the meanings of the terms used in different documents and in the literature, before introducing data modelling.

2.1 Basic nomenclature

Repeatability

According to the written standard ISO 3534 (−1 [ISO06] and −2 [ISO93, ISO06b]) (similarly in ISO 5725−1 [ISO94] and VIM until 2004 [VIM04]), “repeatability conditions” of measurement are “observation conditions where independent test/measurement results are obtained with the same method on identical test/measurement items in the same test or measurement facility by the same operator using the same equipment within short intervals of time”.

The VIM [VIM08] definition then changed somewhat: (2.11) “condition of measurement in a set of conditions that includes the same measurement procedure,

same operators, same measuring system, same operating conditions and same location, and replicate measurements on the same or similar objects over a short period of time” by adding “and replicate measurements on the same or similar objects”. In other words, the values of all the influence factors/parameters are assumed not to change during those intervals of time [PF06]. This situation is also said to indicate that all the measurements can be considered to occur at the same “experimental unit” [ISO94], and corresponds to the replication of measurements for the purpose (a) above. However, the methods used for purpose (a) do not provide any information about reproducibility and accuracy: purposes (b) and (c) for data assessment are therefore essential for deriving an uncertainty statement in metrology and testing. In particular purpose (c), which concerns *intra-* or *inter-* comparisons, has assumed a critical importance after the MRA [MRA99] introduced the use of “key” *inter-*comparisons for the definition of the “degree of equivalence” between metrological laboratories (*inter pares* comparisons: for a review of the problems arising from the needs prompted by MRA see [Pav06]).

Reproducibility

Reproducibility is defined by VIM [VIM08] as: (2.25) “measurement precision under reproducibility conditions of measurements”. A reproducibility condition is a “condition of measurement out of a set of conditions that includes different locations, operators, measuring systems, and replicate measurements on the same or similar objects” (2.24). Further more, it is noted that “the different measuring systems may use different measurement procedures”. For testing, ISO 5725–1 also includes the effects of calibration and time, states the same in (3.18), *except for the last condition related to different measurement procedures*, an important difference⁴.

Accuracy

In order to introduce the concept of accuracy, it is necessary to introduce first the concept of *systematic effects*. However, their meaning is *not* unequivocally specified in recent written standards and in the literature.

Consider first the written standards, which represent the substantial reference for the metrology and testing community. In VIM [VIM08], for example

⁴ There are also contrasting *reproducibility* definitions. For QUAM is instead the “variability obtained when different laboratories analyse the same sample”, whereas “intermediate precision relates to the variation in results observed when one or more factors, such as time, equipment and operator, are varied within a laboratory” [QUAM00] (the usual definition is given for repeatability: “variability observed within a laboratory, over a short time, using a single operator, item of equipment, etc.”)

(2.18), a systematic effect is a “component of measurement error that in replicate measurements remains constant or varies in a predictable way”⁵. This definition does not explicitly involve random variables. Similarly, in [QUAM00] the systematic error is defined as: (2.4.7) “a component of error which, in the course of a number of analyses of the same measurand, remains constant or varies in a predictable way. It is independent of the number of the measurements made and cannot therefore be reduced by increasing the number of analyses under constant measurement conditions”. ISO 21749 [ISO05] states “sometimes it is difficult to distinguish a systematic effect from random effects and it becomes a question of interpretation and the use of the related statistical models. In general, it is not possible to separate fully random and systematic effects”. According to the NIST Handbook, “the terms ‘bias’ and ‘systematic error’ have the same meaning in this handbook” [NIST06]. In ISO 5725 and, in general, in testing documents the term ‘bias’ is most commonly used, in fact, instead of ‘systematic error’. ISO 3534–2 [ISO06b] states: (3.3.2) “Note 1 Bias is the total systematic error as contrasted to random error. There may be one or more systematic error components contributing to the bias. A larger systematic difference from the true value is reflected by a larger bias value”.

The relevant literature, accordingly, reports a variety of interpretations. For example, in [LW06b] a systematic error is said to be “neither random nor observable” and “introduces a bias that may be estimated from other information, mostly in the form of enclosing limits”. According to [Gra05] “unknown systematic errors” are not treated statistically and “remain constant in time and unknown with respect to magnitude and sign” at least “during the time to pick up a series of repeated measurements”; they are restricted to a “confining interval” defined “by worst case estimations”, and are to be combined “arithmetically” with random errors. In [KDP03] the systematic error is referred to as “unknown bias” and considered as a random variable, with zero expectation⁶ in the “random laboratory-effect model” and with nonzero expectation in the “systematic laboratory-effect model”, where it is replaced with

⁵ But until the 2004 draft [VIM04], the VIM definition was the “mean that would result from an infinite number of measurements of the same measurand carried out under repeatability conditions minus a true value of the measurand” (i.e., a random variable that carries the very same uncertainty of the repeated measurements). *Reference in the text to VIM 2004 draft is done solely for the purpose to indicate recent significant changes in the definitions.*

⁶ According to [ISO06] ‘expectation’ is the “integral of a function of a random variable with respect to a probability measure over the sample space” (2.12). According to [Ruh07], ‘expectation’ is “a fuzzy term in metrology”, as it is said to have three different meaning: “synonym for the arithmetic mean value” (in [ISO06b] “expectation, i.e. the mean of a specified set of measurements” (3.2.7d)), “a result of the expectation operator” or, “the result of linear averaging operation in the statistic domain using infinite effort” (“the results are never expectations but always estimates.” See also Section 4.2).

the known “correction for bias”. In [Wil06b], there is “no distinction made between the so-called ‘random errors’ and ‘systematic errors’; both are regarded as being randomly distributed with respect to the set of errors that would arise under the universe of potential measurement systems and measurements”. In [FP06] the systematic error is considered a random variable with the characteristics illustrated in Section 6, model (4), and in [For06] they have a role in decision making.

Furthermore, DIN 1319–1 [DIN95] indicates that the total systematic error comprises two components:

- One covers the *known* systematic measurement error,
- The other one covers the *unknown* systematic measurement error.

Concerning the *known* systematic errors see Section 3 on GUM treatment.

Concerning the *unknown* systematic errors, the procedures defining reproducibility (within one laboratory) only allow an estimate of the increase in uncertainty due to the effect of the variability of the influence parameters to be evaluated, on the basis that the procedure was itself able to randomise them fully (an overly optimistic interpretation according to [Gra05]). In fact, no procedure suitable to obtain measurement precision within one laboratory can provide evidence on whether the expectation of the uncertainty component due the total variability of the influence factors is zero or not [PF06].

Two important consequences follow.

1. When the variability occurs during a “short time interval” or otherwise when the experimental conditions are such that most of the influence parameters are out of control, *the conceptual difference in ‘reproducibility’ and ‘repeatability’ vanishes*; ⁷
2. Should some influence factors remain constant or their mean value be different from zero⁸, the bias cannot be known within the laboratory, unless and until that standard is compared with another standard [PF06, VIM08, Wil06b, FP06]: *there is a time–scale dependence of the data model for reproducibility*.

Coming now to the term *accuracy*, VIM [VIM08]⁹ states: (2.13) “closeness of agreement between a measured quantity value and a true quantity value of a measurand”, noting: “1 – the concept ‘measurement accuracy’ is not a quantity and is not given a numerical quantity value. A measurement is said to be more accurate when it offers a smaller measurement error. 2 – the term ‘measurement accuracy’ should not be used for measurement trueness and the term measurement precision should not be used for ‘measurement accuracy’, which, however, is related

⁷ This issue is pointed out also in ISO 21749 [ISO05].

⁸ As in the definition of systematic error of [VIM08] and [QUAM00] reported above.

⁹ This is basically the definition up until the 2004 VIM draft [VIM04]. It is worth reporting that in the 2006 VIM draft, two alternate definitions had been

to both these concepts. 3 – ‘Measurement accuracy’ is sometimes understood as closeness of agreement between measured quantity values that are being attributed to the measurand.”

The *true value* of a quantity is: (2.11) the “quantity value consistent with the definition of a quantity”, noting that “in the Error Approach to describing measurement, a true quantity value is considered unique and, in practice, unknowable. The Uncertainty Approach is to recognize that, owing to the inherently incomplete amount of detail in the definition of a quantity, there is not a single true quantity value but rather a set of true quantity values consistent with the definition. However, this set of values is, in principle and in practice, unknowable”, and that “due to definitional measurement uncertainty, there is a distribution of true values consistent with the definition of a measurand . . . by nature unknowable”. In the NIST Handbook “accuracy is a qualitative term referring to whether there is agreement between a measurement made on an object and its true (target or reference) value” [NIST06].

Testing applications very often show an intrinsic difference with respect to most metrology applications, in the sense that a *true value* can be assigned to the measurand: in fact, ISO 5725–1 defines *trueness* (3.7) in an operational way as “the closeness of agreement between the average value obtained from a large series of test results and an accepted reference value” (ISO 5725–5 is dedicated to this issue). In testing, *bias*¹⁰ is defined as follows: “the difference between the expectation of the test result and an accepted reference value” [ISO93, A2LA02, ISO99], but as “the difference between the expectation of a test result or measurement result and a true value” in [ISO06b]; “the bias of a test method is usually determined by studying relevant reference materials or test samples. . . . The uncertainty associated with the measurement of the bias is an important component of the overall uncertainty” [EA03]; “where the bias is significant compared to the combined uncertainty, additional action is required” (i.e., eliminate, correct, report, or increase uncertainty) [QUAM00]. In testing there is also a specific term ‘method bias’. In general, according to [NIST06], “bias is

considered: (2.13) “<classical [error] approach> closeness of agreement between a measured quantity value and a true quantity value of the measurand” and noting “The concept ‘measurement accuracy’ is not given a numerical value, but a measurement is said to be more accurate when it offers a smaller measurement uncertainty. Measures of measurement accuracy are found in ISO 5725”; (2.14) “<uncertainty approach> closeness of agreement between measured quantity values that are being attributed to the measurand” and noting “the concept measurement accuracy is not given a numerical value, but a measurement is said to be more accurate when it offers a smaller measurement uncertainty”. Then, the VIM 3rd Edition [VIM08] adopted the “Uncertainty Approach”.

¹⁰ Actually, “consistent bias”, as indicated by NIST “bias that is significant and persists consistently over time for a specific instrument, operator, or configuration should be corrected if it can be reliably estimated from repeated measurements.” (2.5.3.3.2.) [NIST06] (notice the incorrect use of the term ‘repeated’). See also in this respect [ISO04].

a quantitative term describing the difference between the average of measurements made on the same object and its true value. In particular, for a measurement laboratory, bias is the difference (generally unknown) between a laboratory's average value (over time) for a test item and the average that would be achieved by the reference laboratory if it undertook the same measurements on the same test item", and "the terms 'bias' and 'systematic error' have the same meaning in this handbook".

Testing assumptions concerning a known true value are generally not possible in metrology as there is no authority higher in the hierarchy to determine a true value.

3 The GUM approach to the measurement process: Type A and Type B components of uncertainty

In order to arrive at an "expression of experimental uncertainties", a different classification of the uncertainty components to that arising from the three-step procedure discussed in Section 2 has been adopted since 1980 (Recommendation INC-1 (1980) [INC80, Kaa81, Gia82]). It defines two categories "according to the way in which their numerical value is estimated", respectively: A —"those which are evaluated by applying statistical methods to a series of repeated measurements"; and B —"those which are evaluated by other means". The GUM [GUM95] adopted in 1995 this classification for the evaluation of the standard uncertainty:

(Type A) "... method of evaluation of a standard uncertainty by the statistical analysis of a series of observations" (2.3.2);

(Type B) "... method of evaluation of a standard uncertainty by means other than the statistical analysis of a series of observations" (2.3.3).

With reference to the DIN distinction, GUM clearly considers only the *known* systematic errors (in contrast to EA-4/16 and A2LA Guide), that is, the *recognised* effects of *some* influence parameters. Its approach consists in randomising the *recognised* systematic effects¹¹. Therefore, it ensures only a partial compensation of the total systematic error, which consequently could not be declared to have zero expectation¹².

¹¹ Zero expectation after correction of 'known' systematic effects. The random components of uncertainty "include components of uncertainty associated with the corrections" [Gra05].

¹² Another viewpoint is that, should unrecognised systematic effects be postulated, the state of knowledge about these effects should be assumed to have (exactly) zero expectation. However, this seems to contradict the usual definitions of 'systematic effect' as the source of 'bias' (see Section 2.1). In addition, postulating that after the 'correction' of recognized systematic effects the expectation is zero requires knowledge of the 'true value', which is not always the case.

3.1 Basic nomenclature

Repeatability

According to VIM [VIM08], “Type A evaluation of measurement uncertainty” arises from the “evaluation of a component of measurement uncertainty by a statistical analysis of measured quantity values obtained under defined measurement conditions” (2.28), where the conditions can be “repeatability condition of measurement, intermediate precision condition of measurement, and reproducibility condition of measurement”. This definition represents a substantial change with respect to [VIM04] draft¹³ as it now includes *intermediate* and *reproducibility* conditions.

The GUM apparently dropped from the definition the term ‘repeated’ with respect to INC–1 (see above), and uses instead “founded on frequency distributions” (4.1.6), but it does in fact use ‘repeated’, for example, in (3.3.5): “. . . obtained from a Type A evaluation is calculated from a series of repeated observations . . .”¹⁴.

In other words, the values of all the *influence factors/parameters* are assumed not to change during those intervals of time [PF06]. This situation is that already said to indicate that all the measurements can be considered to occur at the same “experimental unit” [ISO99], and corresponds to the replication of measurements for the purpose (a) in Section 2. The current definition of Type A evaluation seems only to involve the replication of measurements for purpose (a).

Reproducibility

GUM prescribes: (3.2.4) “it is assumed that the results of a measurement have been corrected for all recognised significant systematic effects”, effects that are arising from “not being able to hold completely constant each influence quantity” [GUM95]¹⁵. The *recognised* systematic errors are assumed to be random variables with zero expectation after correction (GUM 3.2.3), but the assumption that *all* systematic effects have zero expectation is not always valid in

¹³ (2.13) “A statistical analysis of the quantity values obtained by measurements under repeatability conditions”.

¹⁴ But see Footnote 15.

¹⁵ However, GUM is referring this sentence to *repeated* measurements: (3.1.5) “variations in repeated observations are assumed to arise because influence quantities that can affect the measurement results are not held completely constant”, a concept that can be found also, for example, in (3.2.2). This is inconsistent according to the prevailing definition of ‘repeated measurements’. In (F.2.4.5) GUM considers the exceptional case when a correction is not applied and contributes to the uncertainty. An extension of this procedure for “uncorrected bias” can be found in [PEP97]. Methods for taking into account the uncertainty in the case of uncorrected bias can be found in [Letal00] and [ME08]. See also in this book Chapter 10, Section 3.2.1.

metrological experiments (or even, in some cases, can never be valid¹⁶), as explained before (but see also Footnote 12). Essentially, the GUM is “randomising systematic errors”, an effect obtained when “the [experimental] setup is frequently arbitrarily readjusted; ... many objects, each encumbered with a specific systematic error, are measured, ... all other systematic influences of any relevance are changed”¹⁷ [Gra05].

In this latter interpretation, the GUM prescription is typical of the replication of measurements for the purpose (b) in Section 2, that is, to get a measure of the *reproducibility*.

In addition, one should argue that the concept of ‘influence quantity’ variability is not a characteristic of *repeatability* (‘repeated observations’) but of *reproducibility* (and that of ‘systematic effects’ a characteristic of *accuracy*).

For an *influence quantity*, GUM means a “quantity that is not the measurand but that affects the result of the measurement” ([VIMII] 2.7), whereas in VIM [VIM08] it is defined as a “quantity that, in a direct measurement, does not affect the quantity that is actually measured, but affects the relation between the indication and the measurement result” (2.52).

It is essential to stress that, in fact, *the influence factors are all and the only potential sources of what are called ‘systematic errors’*.

GUM, as said, uses the reproducibility definition, but for “observations ... obtained under the same conditions of measurement” (4.2.1), and is not consistent with the prevailing definition of reproducibility.

On the other hand, “Type B evaluation” (“of measurement uncertainty components”) is defined by GUM as indicated above (and similarly in VIM up until 2004) and is “founded on *a priori* distributions”. VIM then [VIM08] changed the definition as follows: (2.29) “evaluation of a component of measurement uncertainty determined by means other than a Type A evaluation of measurement uncertainty”, with examples given: “evaluation based on information associated with authoritative published quantity values; associated with the quantity value of a certified reference material; obtained from a calibration certificate; about drift; obtained from the accuracy class of a verified measuring instrument; obtained from limits deduced through personal experience”. These examples of Type B evaluation seem basically involving only expert judgment.

Neither Type A nor Type B evaluations seem to fit unequivocally measurements performed with purpose (b), the assessment of *reproducibility*, as defined in Section 2: not Type B, because uses “method of evaluation of a standard uncertainty by means other than the statistical analysis”, nor Type A because the reproducibility assessment should be based on ‘not-repeated’ measurements.

¹⁶ For example, once a MRA key comparison is done and Draft A of the comparison Report is distributed, outlying data become an evidence of “known systematic effects that significantly influence the estimate”, so, according to GUM, should be “corrected”: this is not allowed by MRA.

¹⁷ Obviously these “influences” are only those under the control of the experimenter. *Time is almost never an influence quantity in itself, but in time the influence quantities can show a variability.*

Accuracy

For GUM, the situation has been illustrated above. In it “the term ‘true value’ is not used” because it is “viewed as equivalent” to the term “value of a measurand” (B.2.3), and it cannot be determined. However, for a definition of accuracy, the VIM one (3.5, II edition) [VIMII] is adopted, though it resorts to the term of ‘true value’.

EA-4/16 [EA03] indicates “this document interprets the GUM as based on corrections included in the model to account for systematic effects; such corrections are essential to achieve traceability”. A2LA Guide [A2LA02], uses the term ‘bias’, adding: “the method assumes that *all* significant systematic effects have been identified and either eliminated or else compensated for by allocation of suitable corrections”.

In conclusion, the effect of the variability of the influence factors can be summarised as follows.

- If the variability of (some of) the influence factors can be estimated, a correction can be performed for the expectation according to the GUM, and the associated variability shall be included in the random error.
- If their variability cannot be estimated and the values of some of the influence factors vary, no correction is possible and the whole effect shall be included in the random error.
- GUM assumes that after these operations, observed data are associated with random variables whose expectations are zero.

4 Other approaches to errors in the measurement process

A comprehensive discussion of contrasting viewpoints concerning the theory of errors, namely in respect to the approach treated in Section 3 (called “randomic theory of errors”) compared with an approach similar to the one treated in Section 2, in the period between the initial recommendation [INC80] and its adoption in [GUM95], can be found in [Col87]. See also [Eis83]. A comparison between the two approaches of Sections 2 and 3 can be found in [ISO04]. A treatment of bias in the context of the approach of Section 3 (GUM) can be found in [ISO05].

Other approaches have been proposed. In Section 4.1 (and 6.1.3) one of them is summarised. Another approach can be found in [SSC06], where also a comparison with other approaches is performed, “based on the level of the analytical information used to estimate the measurement uncertainty (e.g., *supralaboratory* or *intralaboratory* information), instead of the direction of information flow (‘bottom-up’ or ‘top-down’) towards the level of information where the test is performed”.

4.1 The total error and its shared and specific components

In [Wil06b] a different approach to the components of the measurement error in metrology has been proposed. It is based on the concept of “*interoperability*”, that is, of the “irrelevance, for the task in question¹⁸, of the identity of the standard to which the appropriate measurement result is traceable”. It comes from the vision of “an approach to uncertainty analysis emphasizing a natural expectation of a client” [Wil06d].

For this purpose, the guiding principle is that the uncertainty is related to a “total error”, linking the measurement results and the unknown measurand value, the latter being “the value taken by some appropriate random variable”. “No distinction [is] made between the so-called ‘random errors’ and ‘systematic errors’ that make up the total error; both are regarded as being randomly distributed with respect to the set of errors that would arise under the universe of potential measurement systems and measurements”.

A distinction is made, conversely, between “a [error] component incurred by each laboratory” and “a [error] component incurred by a laboratory independently of the results of the other laboratories and of the first component”, that is, specific to each laboratory considered, when the measurements are performed within each laboratory.

If the measurement results of at least two laboratories are considered, “the rationale surrounding a comparison involves the possible existence of an extra error with a magnitude that is sufficient to indicate inconsistency with other results”. Also for this additional error component, it is supposed “the existence of shared and individual components [of the error] . . . (possibly zero) arising from unforeseen or incompletely assessed sources of error”. See Section 6 about the corresponding proposed data model (6).

A distinction between shared and specific errors can also be found in this book, Chapter 5, Section 5.1.

4.2 A distinction between measurement and measurand

Measurement results aim at accurately representing the value of the measurand. However, although the measurement result is conceptually different from the value of the measurement, “for any scalar quantity, subtleties arise only when its uncertainty is not symmetric about that quantity’s reported value [DSWH05], so usually no distinction needs to be made between the distribution of measurements (or *gedanken* measurements) and the distribution of the measurand (to which formal uncertainty distributions refer). The formal justification for this is facilitated by the standard practice by metrologists of using the *same value* to be the best representation both of the (fully corrected) measurement, and of the measurand. This ‘fiducial

¹⁸ The task in question is that concerning the “operability” of a metrological institution (NMI), defined as “the ability of the NMI to make measurements that are sufficiently accurate for the task in question” [Wil06b].

value’ [see this book, Chapter 3, Section 4 and [WI06]] simplifies, and in our view strengthens, the fiducial argument”. [citation from this book, Chapter 11]

5 Data modelling in metrology and in testing for *intralaboratory* measurements

Intralaboratory operations on a single standard comprise consideration of both repeated and nonrepeated measurements. Though there are some basic differences in modelling for metrology and testing, it is useful to discuss the two frames together, in order to show and discuss in a simple way the similarities and dissimilarities.

5.1 Repeated measurements

When J replicated measurements are performed in a laboratory under *repeatability* conditions, within each i th ‘experimental unit’¹⁹ [ISO99], the data models²⁰ used for their results can be different for the metrology and for the testing contexts.

In calibration, the model (called “nonexistent laboratory effect model” in [KDP03]) is written:

$$y_{ij} = a + \varepsilon_{ij} \quad i = 1, \dots, I \quad j = 1, \dots, J^{21}, \quad (1)$$

where y , the estimate of the measurand value, is drawn from a random variable $Y = f(X_1, \dots, X_N)$, X_n the “measurable quantities” [EA03], a is the value of the measurand – always unknown by definition – and ε_j is the zero-mean random error occurring at the j th measurement. The replication of the measurements allows knowledge to be gained about the statistical properties of Y . By increasing the number of repeated measurements, the standard deviation associated with the estimate of the measurand value can be reduced.²²

¹⁹ Here, I distinct ‘experimental units’, or groups of measurements, are performed in a *single laboratory* at different times. Under certain special conditions, as discussed in Section 6, one can consider instead the case that I *laboratories* are performing the measurements, each on their own standard, and they pertain to the same ‘experimental unit’: in this case, the subscript i refers to the i th laboratory and any reference to ‘group’ should be changed to ‘laboratory’.

²⁰ “The statistical conclusions are conditional on the assumed model. Therefore, the conclusions are justified only to the extent that the assumed model is justified” [Kak04].

²¹ If same J for all i .

²² This model is basically the one underlying the GUM. It is also typical of the Bayesian approach: in fact, the “beliefs prior to making observations” require zero-mean probability distributions, because they “cannot reflect uncertainty arising from an unknown and unobservable nonzero distribution mean” and “the sign of a systematic error is typically unknown” [Col87].

In testing, the model is written differently:

$$y_{ij} = m + b_i + \varepsilon_{ij} \quad i = 1, \dots, I \quad j = 1, \dots, J, \quad (2)$$

where “ m is the general mean (expectation); b is the value of a bias component under repeatability conditions; ε is the random error occurring under repeatability conditions” [ISO94]²³. Commonly, in testing m is known, assessed by a hierarchically higher rank of laboratories or stipulated by consensus.

Thus, the metrology and testing models are based on different assumptions. This becomes evident by considering the different answers, in the two cases, to the simple question: can the concept of ‘repeated measurements’ be extended *interlaboratories*?

In testing the answer is yes, because: (i) a standard method is used; (ii) a default uncertainty is associated to the method; (iii) the ‘true value’ is known (‘reference value’); (iv) each laboratory must work under ‘repeatability conditions’.

In calibration the answer is generally no, because: (i) the use of a standard method is not required; (ii) uncertainties in different laboratories can span more than a factor ten; (iii) the ‘true value’ is, by definition, unknown because there is no higher-ranked laboratory to assess it; (iv) each laboratory assumes to be working under ‘reproducibility conditions’, though this assumption cannot be fully tested using only a *withinlaboratory* knowledge.

5.2 Non-repeated measurements on the same standard within each laboratory

Before introducing the models for reproducibility, it is worthwhile remembering what has been illustrated in Section 2.1, that the nature of *systematic effects* and their variability is *not* unequivocally specified in written standards and in the literature.

Methods for the purpose (b) (Section 2) resort to model (1) that becomes, for these *non-repeated* measurements²⁴ *before any comparison takes place* (i.e., according to the *withinlaboratory* knowledge; see the next Section 6 and Footnote 27 for the effect on the model of *betweenlaboratories* knowledge):

$$y_{ij} = a + \varepsilon_{ij} + \eta_{ij} \quad i = 1, \dots, I \quad j = 1, \dots, J, \quad (3)$$

where ε is the part of the zero-mean random error occurring at every i repeated measurement and η is the one arising from the additional non-repeated measurements obtained by checking with a suitable procedure – possibly augmented by expert judgement – the effect of the variability of the influence

²³ Model (2) can also be written by explicitly specifying the components of bias b_i and of the random error ε_{ij} [ISO94, ISO04, ISO05, PF06].

²⁴ They are generally performed in I groups or series, corresponding to I ‘experimental units’ [CCCP04, CCPR04, CCPR05, CPR02, Fil06, Pav00, Pav04, PC90, PF06, PICBC03].

factors: $\eta_{ij} = \sum \eta_{ijk}$ over k influence factors. Actually, only $(\varepsilon_{ij} + \eta_{ij})$ are measured. Should randomisation of systematic effects truly occur, both ε_{ij} and η_{ij} would be really zero-mean errors.

In [KDP03, Kak04] (see also [For06b]) this model is called the “random laboratory-effects model”, written as model (2), where (adapting to the chapter notation) $b_i = (X_i - Y_j)$ is the laboratory effect in x_{ij} (‘bias’ in the NIST terminology) and $\varepsilon_{ij} = (x_{ij} - X_i)$ is the *intralaboratory* error in x_{ij} . Capital letters indicate the random variables from which the samples, written in lowercase letters, are drawn. “The laboratory biases b_i are regarded as random variables having the same normal sampling distribution with expected value zero and variance $\sigma_i^2 \geq 0$, called *interlaboratory* variance” [KDP03].

As a matter of fact, in most cases, the data *supplied* by each laboratory *to its users* as a calibration, or *to a comparison* in the case of *intercomparisons*, should be considered representing the *typical* capability of the laboratory, that is, as samples from the population of *repeated* measurements performed on the specific laboratory standard. The laboratory value is supplied as a representative value of its (i.e., local) population; it is *not* a summary statistic of only the (few) specific measurement results obtained for the calibration, or for the comparison. It is assumed, instead, to be consistent with the expectation of the local standard, as currently maintained. In one word, it is *not* a ‘special’ value nor specific of the calibration, or of the comparison. Similarly the associated uncertainty is the laboratory typical level of capability in realising the standard, including the reproducibility component (η) in model (3)²⁵, *not* the uncertainty associated to the (generally few) specific measurement results obtained for the calibration or for the comparison [Pav06].

As already stated, no *intralaboratory* experimental procedure on a *single* measurement standard can provide any evidence that the uncertainty component due to the total variability of the influence factors, η , is zero-mean or not [PF06].

In general, *intralaboratory* comparisons of measurement standards add limited knowledge in this respect; however, they fully change the perspective of the measurements, as illustrated in Section 6.

6 Data modelling in metrology and in testing for *interlaboratory* measurements (*intercomparisons*); comparison specific problems

In order to evaluate accuracy, replicated measurements are gathered for purpose (c) in Section 2 in a comparison, to perform an evaluation that includes the analysis of the differences found between the expectations assigned to

²⁵ Notice again that, the weaker the control on the influence factors, the less the conceptual difference between the concepts of ‘repeatability’ and ‘reproducibility’.

their samples/devices/standards by the laboratory (in *intracomparisons*) or by each laboratory (in *intercomparisons*).

In other words, common past experience suggests that one should assume, *as prior knowledge*, that the comparisons (both *intra* and *inter*) are the only means that metrologists have to obtain evidence about the *differences* between the measured values of the samples/devices/standards assigned by each laboratory. In this respect, the aim of an *intra* or of an *inter* laboratory comparison operation, in general (and of a KC in particular), is *not* to increase the number of repeated measurements for decreasing the standard deviation associated with the measured value of the standards, but to obtain a measure of the differences between the measured values of the participant standards; that is, for MRA, to obtain “a quantitative measure of the degree of equivalence of national measurement standards” [MRA99].

As said in Section 5, *intracomparisons* are generally of limited value for this purpose. In fact, even if every laboratory provides values “corrected for all *known* systematic effects”, this knowledge is in general insufficient to obtain an uncertainty budget within each laboratory able to determine a reliable assessment of accuracy without external assessment.

It is necessary to stress immediately that, contrary to precision evaluation, *intercomparisons* do *not* evaluate the detailed uncertainty information *internal* to each laboratory, but instead evaluate only *input data provided to the comparison by each laboratory*, in general a single value for each standard and an associated uncertainty estimate. It is very seldom that comprehensive statistical information such as a probability density function (pdf) (whose moments provide the mean, standard deviation, etc.) is supplied.

6.1 Data modelling

The models illustrated in the following apply to the *input data* for an *intercomparison*, irrespective to the intrinsic nature of the measurand. The measurands are not all of the same nature in comparison operations. In fact, two broad classes [Pav05] or types [KDP04] of measurands have been identified. Measurands of different nature require careful consideration when the statistical treatment of the input data has to be decided upon and the outcomes computed, for example, the KCRV for the MRA Key Comparisons (see Section 6.2).

There are many approaches in the literature to model comparison data²⁶. However, one can basically summarise them into the two main viewpoints labelled as “Approach A” and “Approach B” in the following applying to

²⁶ For example, [CFH07, DP06, FP06, Gra05, IWM04, Kak04, KDP03, KDP04, KTH06, LW98, LW06, PM82, RV98, SE91, SoSi06, Whi00, Whi04, WI06, Wil06b, Wil06c] and references therein.

most relevant cases, with some exceptions summarised below in “Other Approaches”.

Approach A

The *prior* knowledge is used, consisting of the evidence that for comparisons in general, “when the i -th participant repeats the comparison j times, then its results can be distributed about an expectation value differing from the measurand value a by an amount b_i with standard deviation s_i ” [Whi00], where b_i has the same meaning as in model (2), and is normally called the ‘laboratory component of bias’ (with risk of confusion with the meaning of this word in testing).

In other words, the basic model for a comparison operation (e.g., [Bal01, CCCP04, Due07, KDP02, KDP03, Kak04, Pav04, Pav05, Pav06b, PF06, Whi00, Whi04, WI06b]) is the following²⁷:

$$y_i = a + b_i + \varepsilon_i \quad i = 1, \dots, I \tag{4}$$

where the subscript i refers to the standard of the i th laboratory and b_i are random variables. Because each laboratory generally supplies a single value to the comparison, index j is omitted, and consequently the input data to the comparison are a set of I *non-repeated* measurement results. The random variables B_i from which the values b_i are drawn are assumed to *not* have expectation zero.

In [KDP03] this model is called the “the systematic laboratory-effect model” and the use of the ‘corrections’ recommended by GUM is preferred to the use of the ‘bias’, because “the bias ($X_{\text{UCR}} - Y$) is an unknown constant but the correction for bias, denoted by C , is a variable with a state-of-knowledge probability distribution”, where X_{UCR} denotes the uncorrected X and the “expected value and standard deviation ... [are] denoted by c and $u(c)$, respectively”, with c not necessarily null.

In [FP06], model (4) has been used in an apparently different way. In it, “ a represents the measurand associated with the artefact, b_i a systematic [laboratory] effect present in all measurement results from that laboratory, and ε_i a random effect for that particular result: As part of the uncertainty budgeting activity, the laboratory estimates that $e_i \sim N(0, \sigma_i^2)$ and assigns a distribution $B_i \sim N(0, \rho_i^2)$ for the systematic effect parameter”, where B_i are the random variables from which the values b_i are drawn. The authors “regard ρ and σ as known *a priori*”²⁸ and

²⁷ One has to note that, applied to sets of *intralaboratory* data, *model (4) should also become the model to be used instead of model (3) for the non-repeated measurements performed for obtaining a measure of reproducibility when withinlaboratory knowledge is supplemented by the betweenlaboratories knowledge* arising, e.g., from a comparison operation. In model (4) ε_i becomes $\varepsilon_i + \eta_i$.

²⁸ In “the ‘standard’ model of ISO 5725 it is assumed that ρ and σ are constant over i ” [FP06] because a standard method is used, but that does not, in general, apply to metrology.

this justifies the zero-expectation assumption, as indicated here in Section 5, whereas, obviously, the related uncertainty ρ is not omitted, essentially the situation described in Section 5 when model (2) applies. Then the comparison is performed, and some knowledge is gained about the B_i . Therefore, the treatment in [FP06] is simply taking into account the *time scale* indicated in Section 5 concerning the effect of the variability of the influence factors (see Footnote 27). In [FP06] only the typical case of testing is discussed, where eventually “the reference value a is published”: in this case, each b_i can be obtained individually. In the most common case in metrology, instead, a remains unknowable, so the situation is that depicted in Footnote 30. However, a can be published in some cases, such as the KCRV of a key comparison (see Section 6.2).

Approach B

Model (1) is used in this approach (e.g., [CC03, Cox02, EWC05, Lir06, LW06, LW06b, Nie00]), disregarding the fact that comparison operations *in metrology* necessarily involve non-repeated measurements (see Footnote 22). It does *not* include a laboratory ‘bias’ term, and, according to Footnote 19, the *intercomparison* index i indicates the i th laboratory:

$$y_i = a + \varepsilon_i \quad i = 1, \dots, I, \quad (5)$$

where $Y \sim N(0, \sigma^2)$. Because each laboratory generally supplies a single value to the comparison, index j is omitted, and consequently the *repeated* measurement results are a set of I observation values, irrespective of the fact that they are taken one for each laboratory²⁹, where y_i is drawn from the random variable Y_i and ε_i is the zero-mean error associated to it.

Other approaches

Shared and individual components of error

In [Wil06b], as a consequence of the assumptions reported in Section 4.1, the data model is written:

$$y_i = a + \varepsilon + \varepsilon_i + \varepsilon' + \varepsilon'_i, \quad (6)$$

where:

- a is the value of the measurand, that is, “the quantity intended to be measured” according to VIM definition, “by nature unobtainable”;

²⁹ Because these are repeated calibration measurements performed in different laboratories, it may be difficult to apply them to the current ‘repeatability condition’ definition.

- The first pair of subsequent terms are estimates performed *within* each laboratory: ε is the component of the total error that is assumed to be common to all laboratories, therefore not depending on the subscript i ; ε_i is the zero-mean and with variance depending on the laboratory;
- The second pair of terms, “(possibly zero) arising from unforeseen or incompletely assessed sources of error”, additionally holds specifically in the case of comparisons: ε' is “the shared offset of which the laboratories are unaware, be it a reflection of the state of the art, an offset due to incomplete understanding of the measurement model or . . . , for instance”, which is *not* determined by the comparison; ε'_i are “the individual offset” for each laboratory arising from a source unrecognised by the laboratory.

The model, according to its author’s assumptions, can be simplified:

- ε'_i values “constitute a sample drawn randomly from a normal distribution with unknown variance σ'^2 , possibly contaminated by outliers”.
- $E + E' = \varepsilon + \varepsilon'$.

Thus, the expectation $\mu = a + \varepsilon + \varepsilon'$ and $E_i + E \sim N(0, \sigma_i^2 + \sigma'^2)$, with μ and σ'^2 being unknown parameters.

Nonprobabilistic systematic errors

Opposite to the previous approach and to GUM, in [Gra01,Gra05] the systematic errors are preferred to be treated as nonprobabilistic, “to make allowance for them by introducing biases and worst–case estimations”.

The plurality of “unknown systematic errors” are supposed to bring to an overall effect, a constant-in-time unknown systematic error f , whose value is supposed to fall into a confining interval $f_1 \leq f \leq f_2$. This interval is defined as the bound for the values of f , is not a confidence interval, and its width is matter of expert (metrologist) judgement only. The corresponding model is written as model (4), but in this case $b \equiv f$ is not a random variable.

A full exploitation of the concept of using bounding intervals, particularly useful in the treatment of systematic errors, is developed in Chapter 4 of this book. Interval statistics, in turn, can be seen as a particular case of fuzzy statistics ([ABF04] and this book Chapter 7). For a fuzzy approach to the theory of measurement inexactness see, e.g., [UW03].

6.2 Specific problems and outcomes of the *intercomparisons*, namely the MRA key comparisons

In Figure 1 the process, from planning to outcomes, of a comparison is summarised. As can be found in more detail in [Pav06,Pav07], the basic steps of this process are the following.

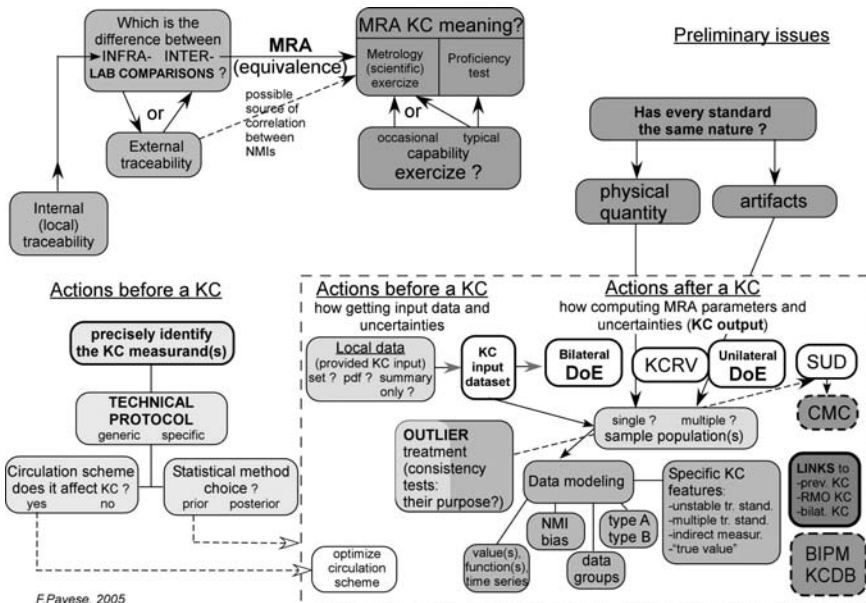


Fig. 1. Roadmap of a (key-)comparison [MRA99], from planning to outcomes. Path starts from top left to right, then bottom left to right

Step 1: Before planning a comparison

Preliminary issues have to be understood and decided:

- Whether a specific comparison, namely a ‘key comparison’, has to be considered a proficiency test [Pav06]
- Which is the intrinsic nature of the standards to be compared (‘classes’ or ‘types’ of standards) [CC03, Pav00, CP02, CPR02, CCCP04, CCPR05, Due07, Pav05, KDP04, SD05]

Step 2: Identifying the measurand and writing the protocol

Actions have to be taken before starting the exercise to identify correctly the measurand subject to the comparison [Pav07] and to write a clear protocol allowing a correct execution of the comparison and of the subsequent data analysis [Pav07].

Step 3: Modelling comparison input data and computing the comparison outcomes

Input data to the comparison: The circulation scheme of the standards needs to be carefully planned and clearly identified and so needs the type and number of input data to the comparison required from each participant [Pav07].

Modelling of the input data is generally performed depending upon the decisions of the consortium of participating laboratories, normally according either Approach A or Approach B.

Approach A. After estimates of the differences between values b_i are obtained³⁰, a check for the *compatibility*³¹ of the b_i with each other may be performed. Compatibility test failure for some ($b_h - b_k$) shall indicate that the hypothesis that these values are not significantly different from zero is false.

Approach B. A test is performed to check for *consistency* of the data with the assumption of repeated measurements. Usually, a χ^2 test is proposed for this purpose, or the use of ‘normalised errors’ (“metrological ratio”) or ‘z-score’ (e.g., see [SD06] and references therein). Test failure shall not pass the hypothesis that the measurements are repeated. Test acceptance shall not change the fact that they intrinsically are non-repeated measurement.

*Output summaries of the (key-)comparison*³². It is vital to compute correctly the summary statistics and, in the case of a MRA key comparison (KC) [MRA99, CIPM99], the specific *output* parameters summarising the results obtained:

- The key comparison reference value (KCRV, when necessary [Pav07])
- The degrees of equivalence (DoE, both bilateral and with respect to the KCRV)
- Any possible ‘significant unresolved deviation’ (SUD)
- The links between different KCs, for their repository in the BIPM database [KCDB] and their use for the ‘calibration measurement capabilities’ (CMC).

³⁰ In fact, in metrology b_i remain as unknown as a is, only the differences ($b_h - b_k$) of pairs of laboratories are measured.

³¹ According to VIM [VIM08] the definition of “metrological compatibility” is (2.47) as “absolute value of the difference of any pair of measured quantity values from two different measurement results is smaller than some chosen multiple of the standard measurement uncertainty of that difference”, also noting that “metrological compatibility of measurement results replaces the traditional concept of ‘staying within the error’, as it represents the criterion for deciding whether two measurement results refer to the same measurand or not. If in a set of measurements of a measurand, thought to be constant, a measurement result is not compatible with the others, either the measurement was not correct (e.g. its measurement uncertainty was assessed as being too small) or the measured quantity changed between measurements”. Until 2004 [VIM04], it was “property satisfied by all the measurement results of the same quantity, characterised by an adequate overlap of their corresponding sets of quantity values”.

³² For example, [Bal01, CCPR04, Cox99, CHW05, DSWH05, DLCS06, DBCSD06, Due07b, EWC05, IWM04, Kak04, KDP02, KDP03, KDP04, Letal00, Lir06, Mul00, Nie00, Pav05, Pav06, Pav07, PF06, Rat05, SD06, SD06b, SHD02, SHD04, SWD05, SWD06, Sut04, Tom05, vC03, WD98, WI04, Wil02, Wil03, Wil06a, Wog05, ZLSS04, ZSLS05].

Consistency and compatibility

It is worth noting that the issue of *consistency*³³ and *compatibility* should not be confused with that of *repeated* or *non-repeated* measurements. In fact, the latter does solely involve the assessment whether a set of measurements is homoscedastic: in particular for Approach B, the variance of the random variables Y_i should be the same for each and all laboratories. The ‘consistency’ test introduces, instead, the possibility that different random variables Y_i can have different expectations and different variances and performs a probabilistic evaluation of the differences in expectations, to which a criterion is generally associated for the definition of *outliers*, to be excluded from the computation of certain summary statistics. The purpose of a ‘compatibility’ test in the frame of *Approach A* is similar³⁴.

The lack of input data consistency cannot directly affect the outcomes of a key comparison, because, according to the MRA [MRA99], the differences between laboratory data have the meaning of “degrees of equivalence” (DoE), a nonhierarchical concept used in the MRA instead of the concept of ‘traceability’, so preventing the use of the concept of outlier, a concept common in testing and in comparisons performed for other purposes.

For this reason, several authors [Bei02, Bei03, Cox02, Cox07, DS06, DSD07, IWV04, SD06, SD06b, WI04, WI05] have recently proposed consistency tests.

This is a mandatory exercise in testing, and is also extremely useful in calibration, when possible. In the context of MRA key comparison, however, inconsistent data are proposed to be screened out only for the computation of the KCRV, limiting its computation to a subset of the dataset, often called the ‘maximum consistent set’. It is believed that this would lead to a ‘better’ KCRV without violating the MRA prescriptions indicated in Footnote 16,

³³ A term used in MRA but not defined by VIM. However, one can say that a set of compatible pairs of observation values is a consistent set.

³⁴ The test hypotheses are generally based on confidence levels or intervals. In metrology, the indication of a threshold for the definition of ‘outlier’ would appear less arbitrary if, instead, a risk level was used. In fact, assessing the level of the risk of a failure (e.g., a wrong value in a certificate of calibration or of test) by indicating how critical (risky) is that value – consider, e.g., a medical or contaminant analysis – is much closer to the intended use. Correct, safe results may be obtained only by deriving proper information concerning acceptable risk from the real situation, and evaluating accordingly the boundaries of the relevant confidence interval. When considering a set of experimental data obtained from a population described in terms of a statistical distribution, a result may fall in a low probability tail owing to chance only, or to the occurrence of an exceptional phenomenon, or a combination of both. No matter which is the real cause, an outlier is produced; should the existence of a perturbing phenomenon be ruled out, chance is left as the only explanation of an unlikely occurrence.

concerning the (unilateral) degrees of equivalence, $(y_h - \text{KCRV})$ for the h th participant, where the KCRV is obtained as a summary statistic of the y_h .³⁵

However, this is not true according to [Pav07b]. In fact, if the degree of consistency could determine how the KCRV is calculated, it will therefore influence the KCRV and the DoEs. However, MRA prescribes that the reported results cannot be adjusted because the DoEs are to be defined in terms of the (unadjusted) data.

Mixture model

A factor influencing the estimate of compatibility or consistency of the input comparison data is the nature of the standards being compared (cfr. *Step 1*), which, according to [CP02, CCPR04, CCPR05, PF06], can suggest the need to introduce a ‘mixture model’ (see, e.g., [Eve81, McL00]).

The probabilistic model, called the mixture model, can provide a characterisation of the overall stochastic variability of the population of a comparison involving Class 2 standards, that is, when not involving artefacts (Class 1) [CCPR04].

In fact, the mixture density function represents the total variability of a super-population comprising several populations, each one identified by a specific pdf. It is often used in the statistical quality control of an industrial process, for example, when nominally identical electronic components have various lifetime constants. Mixture densities are also used to model experimental situations in nonnormal cases: mixtures of normal densities have been used for robust estimation of parameters. The latter are also known as a ‘contaminated’ normal family when used to model a population that follows a normal density except for those occasions when a peculiar observation is recorded. These peculiar observations are not viewed as outliers, but as the effect of an admissible variability.

Consider a finite mixture of densities as a linear superposition of N component densities (named a mixture density in the following). In an intercomparison, suppose that, for $i = 1, \dots, N$, a density function $f_i(x; \mathbf{A}^{(i)})$ is provided for participant i , where $\mathbf{A}^{(i)}$ is the (row) vector of A_i , say, parameters of f_i . Then the mixture density

$$g(x; \mathbf{A}) = \sum_{i=1}^N \pi_i f_i(x; \mathbf{A}^{(i)}), \tag{7}$$

where $(\mathbf{A}, \boldsymbol{\pi})$ are the mixture parameters, g is assumed to characterise the total data variability related to the output pdf. In model (7), $\mathbf{A} = (A^{(1)}, \dots, A^{(N)})$ contains the parameters of the N participants’ pdfs and the π_i are positive *proportions* summing to unity. It may often be appropriate to take these

³⁵ Actually in most cases only of the $(y_h - y_r)$, where y_r is the value attributed to a reference standard, generally a transfer standard.

proportions to be equal; that is, $\pi_i = 1/N$, $i = 1, \dots, N$. Specific metrological reasoning may indicate unequal proportions.

In key comparisons, the KCRV is given as the expectation value of the density function. The moments of the mixture density are given by the weighted sum of the moments of the component densities according to the π_i . Other authors prefer to use other mixture distribution parameters to define the KCRV [Due07].

References

- [A2LA02] American Association for Laboratory Accreditation (A2LA) 2002 *Guide for the estimation of measurement uncertainty in testing*
- [Bal01] Ballico M 2001 Calculation of key comparison reference values in the presence of non-zero mean uncertainty distributions, using the maximum likelihood technique *Metrologia* **38** 155–159
- [ABF04] Allevard T, Benoit E, and Foulloy L 2004 Signal processing on fuzzy nominal scales, in *Proceedings 10th IMEKO TC7 International Symposium*, Saint-Petersburg, Russia, on CD-ROM
- [Bei02] Beissner K 2002 On a measure of consistency in comparison measurements *Metrologia* **39** 59–63
- [Bei03] Beissner K 2003 On a measure of consistency in comparison measurements: II. Using effective degrees of freedom *Metrologia* **40** 31–35
- [CC03] Chunovkina A and Cox M G 2003 A model-based approach to key comparison data evaluation, in *Proceedings 17th IMEKO World Congress*, Cavtat, Slovenia, on CD-ROM
- [CCCP04] Chunovkina A, Ciarlini P, Cox M, and Pavese F 2004 Handling and treatment of measurement data from different groups or sources, in *Proceedings 10th IMEKO TC7 International Symposium*, on CD-ROM
- [CCPR04] Ciarlini P, Cox M G, Pavese F, and Regoliosi G 2004 The use of a mixture of probability distributions in temperature interlaboratory comparisons *Metrologia* **41** 116–121
- [CCPR05] Ciarlini P, Chunovkina A, Pavese F, and Regoliosi G 2005 Mixture distribution models in temperature metrology and the assessment of the degree of equivalence, in *Proceedings TEMPMEKO 2004*, LPM, Zabreb, Croatia, pp 1003–1008
- [CFH07] Cox M G, Forbes A B, and Harris P 2007 Distributions associated with measurands, in *PTB-BIPM Workshop on the Impact of Information Technology in Metrology*, Berlin, Germany. Available on <http://www.npl.co.uk>
- [CIPM99] BIPM 1999 *Guidelines for CIPM Key Comparisons*, Bureau International des Poids et Mesures, Sèvres
- [Col87] Colclough A R 1987 Two theories of experimental error *J Res Natl Bur Stand* **92** 167–185
- [Cox99] Cox M G 1999 *A discussion of approaches for determining a reference value in the analysis of key-comparison data*, NPL Report CISE 42/99. Available on <http://www.npl.co.uk/ssfm>

- [Cox02] Cox M G 2002 The evaluation of key comparison data *Metrologia* **39** 589–595
- [Cox07] Cox M G 2007 Experience gained with the use of the largest consistent subset of results in a key comparison, in *PTB–BIPM Workshop on the Impact of Information Technology in Metrology*, Berlin, Germany. Available on <http://www.bipm.org>
- [CHW05] Cox M G, Harris P M, and Woolliams E 2005 Data evaluation of key comparisons involving linked bilateral measurements and multiple artefacts, in *NCSL International Workshop and Symposium*, on CD–ROM
- [CP02] Ciarlini P and Pavese F 2002 Some basic questions in the treatment of intercomparison data *Izmeritel'naya Tekhnika*, April 2003, in russian) **4** 8–71 (*Measurement Technique* **46** 421–426, from *Workshop on Mathematics, Statistics and Computation to Support Measurement Quality* 2002, St Petersburg, Russia)
- [CPR02] Ciarlini P, Pavese F, and Regoliosi G 2002 A bootstrap algorithm for mixture models and interval data in inter-comparisons, in *Algorithms for Approximation IV*, I J Anderson et al. eds., Huddersfield, University of Huddersfield pp 138–145
- [DBCSD06] Decker J E, Brown N, Cox M G, Steele A G, and Douglas R J 2006 Recent recommendations of the Consultative Committee for Length (CCL) regarding strategies for evaluating key comparison data *Metrologia* **43** L51–L55
- [Dav02] Davis C S 2002 *Statistical Methods for the Analysis of Repeated Measurements*, Springer, New York
- [DIN95] DIN 1319–1 1995 *Fundamentals of metrology —Part I: Basic terminology*, Berlin, Beuth
- [DLCSD06] Decker J E, Lewis A J, Cox M G, Steele A G, and Douglas R J 2006 Evaluating results of international comparisons: worked example of CCL–K2 comparison of long gauge block calibration, in *Proceedings IMEKO 18th World Congress*, Sociedade Brasileira de Metrologia, Rio de Janeiro, on CD–ROM
- [DP06] Desanfant M and Priel M 2006 Road map for measurement uncertainty evaluation *Measurement* **39** 841–848
- [DS06] Douglas R J and Steele A G 2006 Pair difference chi–squared statistics for key comparisons *Metrologia* **43** 89–97
- [DSD07] Decker J E, Steele A G, and Douglas R J 2007 Aspects of evaluating a consistent subset in comparison analysis 2007, in *PTB–BIPM Workshop on the Impact of Information Technology in Metrology*, Berlin, Germany. Available on <http://www.bipm.org>
- [DSWH05] Douglas R J, Steele A G, Wood B M, and Hill K D 2005 A useful reflection *Metrologia* **42** L35–L39
- [Due07] Duewer D L 2007 How to combine results having stated uncertainties, in *Combining and Reporting Analytical Results* ed A Fajgelj et al, RSC, London pp 127–142
- [Due07b] Duewer D L et al. 2007 *CCQM K21 Final Report*, BIPM KCDB. Available on <http://kcdb.bipm.org>
- [EA03] European Accreditation 2003 *EA guidelines on the expression of uncertainty in quantitative testing*, EA–4/16, rev00

- [Eis83] Eisenhart C, Ku H H, and Collé H 1983 Expression of the Uncertainties of Final Measurement Results: Reprints, in *NBS Special Publication 644*, National Bureau of Standards, Washington, DC
- [Eve81] Everitt B S 1981 *Finite Mixture Distributions* Chapman and Hall, London
- [EWC05] Elster C, Wöger W, and Cox M G 2005 Analysis of key comparison data: unstable travelling standards *Measur Techn* **48** 883–893
- [Fil06] Filipe E 2006 Evaluation of standard uncertainties in nested structures, in *Series on Advances in Mathematics for Applied Sciences 72, Advanced Mathematical and Computational Tools in Metrology VII (AMCTM VII)*, P Ciarlini et al. eds., World Scientific, Singapore pp 151–60
- [For06] Forbes A B 2006 Measurement system analysis and statistical process control, in *AMCTM VII (Series on Advances in Mathematics for Applied Sciences) 72* P Ciarlini et al. eds., World Scientific, Singapore pp 161–170
- [For06b] Forbes A B 2006 Measurement uncertainty and optimized conformance assessment *Measurement* **39** 808–814
- [FP06] Forbes A B and Perruchet C 2006 Measurement systems analysis: concepts and computational approaches, in *Proceedings IMEKO 18th World Congress*, Sociedade Brasileira de Metrologia, Rio de Janeiro, on CD-ROM
- [Gia82] Giacomo P 1982 News from the BIPM *Metrologia* **18** 41–44
- [Gra01] Grabe M 2001 Estimation of measurement uncertainties – an alternative to the ISO Guide *Metrologia* **38** 99–106
- [Gra05] Grabe M 2005 *Measurement uncertainties in science and technology*, Springer, Berlin
- [GUM95] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML 1995 *Guide to the expression of uncertainty in measurement (GUM)*, 2nd edn, International Organization for Standardization, Geneva
- [GUMs1] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML 2006 *Guide to the expression of uncertainty in measurement (GUM) – Supplement 1: Propagation of distributions using a Monte Carlo method*, International Organization for Standardization, Geneva
- [INC80] BIPM 1980 *Report on the BIPM Enquiry on Error Statements*, Rapport BIPM 80/3, Bureau International des Poids et Mesures, Sèvres
- [ISO93] ISO 3534–2 1993 2nd edn *Statistics – Vocabulary and symbols – Part 2: Applied statistics* International Organization for Standardization, Geneva
- [ISO94] ISO 5725 1994 2nd edn *Accuracy (trueness and precision) of measurement methods and results* International Organization for Standardization, Geneva, Switzerland
- [ISO99] ISO 3534–3 1999 *Statistics – Vocabulary and symbols – Part 3: Design of experiments* International Organization for Standardization, Geneva
- [ISO03] ISO 21749 2003 *Measurement uncertainty for metrological applications – Simple replication and nested experiments* International Organization for Standardization, Geneva

- [ISO04] ISO 21748 2004 *Guidance for the use of repeatability, reproducibility and trueness estimates in measurement uncertainty estimation* International Organization for Standardization, Geneva
- [ISO05] ISO 21749 2005 *Measurement and uncertainty for metrological applications – Repeated measurements and nested experiments* International Organization for Standardization, Geneva
- [ISO06] ISO 3534–1 2006 2nd edn *Statistics – Vocabulary and symbols – Part 1: General statistical terms and terms used in probability* International Organization for Standardization, Geneva
- [ISO06b] ISO 3534–2 2006 3rd edn *Statistics – Vocabulary and symbols – Part 2: Applied statistics* International Organization for Standardization, Geneva
- [IYW04] Iyer H K, Wang C M, and Vecchia D F 2004 Consistency tests for key comparison data *Metrologia* **41** 223–230
- [IWM04] Iyer H K, Wang C M, and Mathew T 2004 *J Am Stat Assoc* **99** 1060–1071
- [Kaa81] Kaarls R 1981 in *Proces Verbaux des séances du Comité International des Poids et Mesures* **49** A1–A12 (in French)
- [Kak04] Kakher R N 2004 Combining information from interlaboratory evaluations using a random effects model *Metrologia* **41** 132–136
- [KCDB] <http://dcd.bipm.org>
- [KDP02] Kakher R N, Datla R U, and Parr A C 2002 Combined result and associated uncertainty from interlaboratory evaluations based on the ISO Guide *Metrologia* **39** 279–293
- [KDP03] Kakher R N, Datla R U, and Parr A C 2003 Statistical interpretation of Key Comparison reference value and degrees of equivalence *J Res Natl Inst Stand Technol* **108** 439–446
- [KDP04] Kacker R N, Datla R U, and Parr A C 2004 Statistical analysis of CIPM key comparisons based on the ISO Guide *Metrologia* **41** 340–352
- [KTH06] Kakher R N, Toman B, and Huang D 2006 Comparison of ISO–GUM, draft GUM Supplement 1 and Bayesian statistics using simple linear calibration *Metrologia* **43** S167–S177
- [Lir06] Lira I 2006 Bayesian evaluation of comparison data *Metrologia* **43** S231–S234
- [LW98] Lira I and Wöger W 1998 The evaluation of the uncertainty in knowing a directly measured quantity *Meas Sci Technol* **9** 1167–1173
- [LW06] Lira I and Wöger W 2006 Evaluation of repeated measurements from the viewpoints of conventional and Bayesian statistics, in *AMCTM VII (Series on Advances in mathematics for applied sciences)* **72**, P Ciarlini et al. eds., World Scientific, Singapore pp 73–84
- [LW06b] Lira I and Wöger W 2006 Comparison between the conventional and Bayesian approaches to evaluate measurement data *Metrologia* **43** S249–S259
- [Letal00] Levenson M S, Banks D L, Eberhart K R, Gill L M, Guthrie W F, Liu H K, Vangel M G, Yen J H, and Zhang N F 2000 An approach to combining results from multiple methods motivated by the ISO GUM *J Res Natl Inst Stand Technol* **105** 571–579
- [McLP00] McLachan G and Peel D 2000 *Finite Mixture Models*, John Wiley, New York

- [ME08] Magnusson B and Ellison L R 2008 Treatment of uncorrected measurement bias in uncertainty estimation for chemical measurements *Anal Bioanal Chem* **390** 201–213
- [MRA99] CIPM 1999 *Mutual recognition of national measurement standards and of calibration and measurement certificates issued by national metrology institutes* Bureau International des Poids et Mesures, Sèvres
- [Mul00] Müller J W 2000 Possible advantages of a robust evaluation of comparisons *J Res Natl Inst Stand Technol* **105** 551–555
- [Nie00] Nielsen L 2000 *Evaluation of measurement intercomparisons by the method of least squares* Report DFM–99–R39, 3208 LN, Danish Institute of Fundamental Metrology, Lungby, Denmark
- [NIST06] National Institute of Standards and Technology (NIST) 2006 *Engineering statistics handbook (e-Handbook)*, NIST, USA. Available in the DVD attached to this book and online at <http://www.nist.gov/stat.stat.handbook>
- [Pav00] Pavese F 2000 Mathematical problems in the definition of standards based on scales: the case of temperature, in *AMCTM IV (Series on Advances in mathematics for applied sciences)* **53** P Ciarlini et al. eds., World Scientific, Singapore, pp 182–196
- [Pav04] Pavese F 2004 Compound-modelling of metrological data, in *AMCTM VI (Series on Advances in mathematics for applied sciences)* **66**, P Ciarlini et al. eds., World Scientific, Singapore, pp 240–250
- [Pav05] Pavese F 2005 Comments on ‘Statistical analysis of CIPM key comparisons based on the ISO Guide’ *Metrologia* **42** L10–L12
- [Pav06] Pavese F 2006 A metrologist viewpoint on some statistical issues concerning the comparison of non-repeated measurement data, namely MRA Key Comparisons *Measurement* **39** 821–828
- [Pav06b] Pavese F 2006 Some problems concerning the estimate of the degree of equivalence in MRA KC and its uncertainty, in *AMCTM VII (Series on Advances in mathematics for applied sciences)* **72**, P Ciarlini et al eds, World Scientific, Singapore, pp. 323–327
- [Pav07] Pavese F 2007 The definition of the measurand in key comparisons: lessons learnt with thermal standards *Metrologia* **44** 327–339
- [Pav07b] Pavese F 2007 Replicated observations in metrology and testing: modelling of repeated and non-repeated measurements *ACQUAL* **12** 525–534
- [PC90] Pavese F and Ciarlini P 1990 Accurate modelling of translational bias and its application to reduction of thermophysical data series *Metrologia* **27** 145–152
- [PEP97] Phillips S D, Eberhardt K R, and Parry B 1997 Guidelines for Expressing the Uncertainty of Measurement Results Containing Uncorrected Bias *J Res Natl Inst Stand Technol* **102** 577–585
- [PF06] Pavese F and Filipe E 2006 Some metrological considerations about replicated measurements on standards *Metrologia* **43** 419–425
- [PICBC03] Pavese F, Ichim D, Ciarlini P, Balle C, and Casas-Cubillos J 2003 Detection of thermometer clustering in the calibration of large batches of industrial thermometers for the LHC by automated data processing, in *Temperature, Its Measurement and Control in Science and Industry* **8**, D Ripple ed, AIP, New York, pp 429–433

- [PM82] Paule R C and Mandel J 1982 Consensus values and weighting factors *J Res Natl Bur Stand* **87** 377–385
- [QUAM00] Eurachem CITAC 2000 2nd edn *Guide CG4. Quantifying uncertainty in analytical measurements* (QUAM 2000.1)
- [Rat05] Ratel G 2005 Evaluation of the uncertainty of the degree of equivalence *Metrologia* **42** 140–144
- [Ruh07] Ruhm K H 2007 Expectation – A fuzzy term in metrology, in *Proceedings IEEE International Workshop on Advanced Methods of Uncertainty Estimation in Measurement* (AMUEM 2007)
- [RV98] Rukhin A L and Vangel M G 1998 *J Am Stat Assoc* **93** 303–308
- [SD05] Steele A G and Douglas R J 2005 Pair-difference statistics and data pooling: tools for comparison data and reference values, in *Proceedings TEMPMEKO 2004*, LPM, Zabreb, Croatia, pp 997–1002
- [SD06] Steele A G and Douglas R J 2006 Simplicity with advanced mathematical tools for metrology and testing *Measurement* **39** 795–807
- [SD06b] Steele A G and Douglas R J 2006 Extending E_n for measurement science *Metrologia* **43** S235–S243
- [SE91] Schiller S B and Eberhardt K R 1991 *Spectrochim Acta* **46B** 1607–1613
- [SHD02] Steele A G, Hill K D, and Douglas R J 2002 Data pooling and key comparison reference values *Metrologia* **39** 269–277
- [SHD04] Steele A G, Hill K D, and Douglas R J 2004 Pooled data distributions: graphical and statistical tools for examining comparison reference values, in *AMCTM VII (Series on Advances in mathematics for applied sciences)* **72**, P Ciarlini et al. eds., World Scientific, Singapore, pp 264–273
- [SSC06] da Silva R J N B, Santos J R, and Camoes M F G F C 2006 A new terminology for the approaches to the quantification of the measurement uncertainty *Accred Qual Assur* **10** 664–671
- [SWD05] Steele A G, Wood B M, and Douglas R J 2005 Linking key comparison data for Appendix B, in *Proceedings TEMPMEKO 2004*, LPM, Zabreb, Croatia, pp 1087–1092
- [SWD06] Steele A G, Wood B M, and Douglas R J 2006 The middle ground in key comparison analysis: revisiting the median, in *AMCTM VII (Series on Advances in mathematics for applied sciences)* **72**, P Ciarlini et al. eds., World Scientific, Singapore, pp 245–252
- [SoSi06] Sommer K D and Siebert B R L 2006 Systematic approach to the modelling of measurements for uncertainty evaluation *Metrologia* **43** S200–S210
- [Sut04] Sutton C M Analysis and liking of international measurement comparisons *Metrologia* **41** 272–277
- [Tom05] Toman B 2005 Bayesian approach to assensing uncertainty and calculating a reference value in a key comparison experiment *J Res Nat Inst Sci Technol* **110** 605–612
- [UW03] Urbanski M K and Wasowski J 2003 Fuzzy approach to the theory of measurement inexactness *Measurement* **34** 67–74
- [vC03] van der Veen A M H and Cox M C 2003 Error analysis in the evaluation of measurement uncertainty *Metrologia* **40** 42–50
- [VIMIII] BIPM 1993 2nd edn *International Vocabulary of Basic and General Terms in Metrology* (VIM), BIPM/ISO

- [VIM04] BIPM 2004 *International Vocabulary of Basic and General Terms in Metrology* (VIM) (3rd draft of April 2004)
- [VIM08] BIPM 2008 3rd edn *International Vocabulary of Basic and General Terms in Metrology* (VIM), BIPM/ISO
- [WD98] Wood B M and Douglas R J 1998 Confidence–interval interpretation of a measurement pair for quantifying a comparison *Metrologia* **35** 187–196
- [Whi00] White D R 2000 On the analysis and linking of comparison results, in *CPEM Conference digest*, CPEM 2000, Sydney, Australia, pp 325–326
- [Whi04] White D R 2004 *Metrologia* **41** 122–131
- [WI04] Wang C M and Iyer H K 2004 Consistency tests for key comparison data *Metrologia* **41** 223–230
- [WI05] Wang C M and Iyer H K 2005 Detection of influential observations in the determination of the weighted–mean KCRV *Metrologia* **42** 262–265
- [WI06] Wang C M and Iyer H K 2006 Uncertainty analysis for vector measurands using fiducial inference *Metrologia* **43** 486–494
- [WI06b] Wang C M and Iyer H K 2006 A generalized confidence interval for a measurand in the presence of type–A and type–B uncertainties *Measurement* **39** 856–863
- [Wil02] Willink R 2002 Statistical determination of a comparison reference value using hidden errors *Metrologia* **39** 343–354
- [Wil03] Willink R 2003 On the interpretation and analysis of a degree–of–equivalence *Metrologia* **40** 9–17
- [Wil05] Willink R 2005 A procedure for the evaluation of measurement uncertainty based on moments *Metrologia* **42** 329–343
- [Wil06a] Willink R 2006 Forming a comparison reference value from different distributions of belief *Metrologia* **43** 12–20
- [Wil06b] Willink R 2006 Meaning and models in key comparisons, with measures of operability and interoperability *Metrologia* **43** S220–S230
- [Wil06c] Willink R 2006 Principles of probability and statistics for metrology *Metrologia* **43** S211–S219
- [Wil06d] Willink R 2006 An approach to uncertainty analysis emphasizing a natural expectation of a client, in *AMCTM VII (Series on Advances in mathematics for applied sciences)* **72**, P Ciarlini et al. eds., World Scientific, Singapore, pp 344–349
- [Wog05] Wöger W 2005 Remarks on a key comparison reference value, in *NCSL International Workshop and Symposium*, on CD–ROM
- [ZLSS04] Zhang N F, Liu H K, Sedransk N, and Strawderman W E 2004 Statistical analysis of key comparisons with linear trends *Metrologia* **41** 231–237
- [ZSLS05] Zhang N F, Strawderman W E, Liu H K, and Sedransk N 2005 Statistical analysis for multiple artefact problem in key comparisons with linear trends *Metrologia* **42** 1–6

Probability in Metrology

Giovanni B. Rossi

Università degli Studi di Genova, DIMEC, Via All'Opera Pia 15 A, 16145 Genova, Italy
gb.rossi@dimec.unige.it

Summary. The relationship is investigated between probability and metrology, here intended as the science of measurement. Metrology is shown to have historically participated in the development of statistic–probabilistic disciplines, not only adopting principles and methods, but also contributing with new and influential ideas. Two mainstreams of studies are identified in the science of measurement. The former starts with the classical theory of errors and ends with the contemporary debate on uncertainty; the latter originates from the development of a formal theory of measurement and it has attained recent results that make a systematic use of probability as an appropriate logic for measurement. It is suggested that these two mainstreams may ultimately converge in a unique theory of measurement, formulated in a probabilistic language and applicable to all domains of science.

Key words: Metrology, measurement theory, probability, uncertainty, measurability

1 Probability, statistics, and measurement – An historical perspective

1.1 The origins: Gauss, Laplace, and the theory of errors

The Gauss problem

In his *Theoria motus corporum coelestium* (1809, [3]¹) Carl Friedrich Gauss (1777–1855) discusses how to obtain estimates of the parameters of the orbits of heavenly bodies on the basis of a set of observations. In the third section of the second book of the treatise he considers the case of any number of observations and formulates what we here call the Gauss problem: *given N observations that depend upon n unknown parameters, $n < N$, according to*

¹ Note that the references in the bibliography at the end of the chapter have been listed in chronological order, in order to provide an overview of the historical development of the subject.

a known functional relation, and that are affected by measurement errors, estimate the unknown parameters.

In modern notation, we write ²:

$$\mathbf{y} = \mathbf{f}(\mathbf{x}) + \mathbf{v}, \quad (1)$$

where \mathbf{y} is a vector of observations, \mathbf{x} is a vector of the unknown parameters, \mathbf{v} is the vector of measurement errors, and \mathbf{f} is a vector function. This framework applies to many problems in metrology and thus it provides a good introduction to the role of probability in metrology.

We have started with this problem in order to present, in this first part of the chapter, a mainstream of studies centered on the problem of measurement uncertainty, starting with Gauss and ending with the current state of play. We analyze the historical development of this subject, from the classical theory of errors, through the contribution of orthodox statistics, up to the development of the *Guide to the Expression of Uncertainty in Measurement* [26] and to some of the main issues of the contemporary debate. In this context we also present our own view, which is based on a general approach to the probabilistic modelling of the measurement process.

In the second part of the chapter, we deal with another area of studies that concern the foundations of measurement and we attempt to establish a formal theory. We consider the contributions of Helmholtz and Campbell and the debate on the possibility of the measurement of ‘sensory events’, promoted by the British Association for the Advancement of Science in the 1930s, which has had consequences up to the present day. Then we present the representational approach to measurement and some criticism of it. We discuss the role of the measuring instrument in a formal theory and the benefit of a probabilistic approach. This part concludes with a brief outline of a probabilistic theory of measurement that we have recently proposed and published and also with an attempt to make some previsions on the possible future role of probability in the science of measurement. We do not attempt to deal with other approaches such as those concerned with fuzzy sets or with the theory of evidence, however.

But now let us go back to Gauss. To confront his problem, he adopts the following estimation criterion: choose the value of \mathbf{x} that has maximum probability, given the observations \mathbf{y} . So the estimate, $\hat{\mathbf{x}}$, must be such that

$$p(\hat{\mathbf{x}}|\mathbf{y}) = \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}), \quad (2)$$

where $p(\mathbf{x}|\mathbf{y})$ is the distribution of the parameters, given the observations and, conversely, $p(\mathbf{y}|\mathbf{x})$ is the distribution of the observations, given the parameters.

² See Appendix for notation conventions and for a list of the main symbols used in this chapter.

Gauss uses a Bayesian argument to show that, assuming an indifference prior distribution for \mathbf{x} , this is equivalent to finding the value $\hat{\mathbf{x}}$ that maximises the probability of the observations, given the parameters; that is, $\hat{\mathbf{x}}$ may be equivalently characterised by the property

$$p(\mathbf{y}|\hat{\mathbf{x}}) = \max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x}). \quad (3)$$

If we now assume that all the errors are independent and equally distributed, and if $p_v(\cdot)$ is the distribution of each of them, we obtain

$$p(\mathbf{y}|\mathbf{x}) = \prod_i p_v(y_i - f_i(\mathbf{x})). \quad (4)$$

To proceed any further, it is necessary to adopt a proper distribution for the errors v_i . This is where the famous *normal* or *Gaussian* distribution comes into play.

Gauss's probabilistic model for measurement errors

Gauss distinguishes between *systematic* and *random* errors. This distinction, just mentioned in his *Theoria motus*, is more clearly expressed in the successive *Theoria combinationis observationum erroribus minimis obnoxiae* (1823 [5]). Due to the importance of this topic in metrology, it is worthwhile reading the original text.

‘Certain causes of error,’ he writes

are such that their effect on any one observation depends on varying circumstances that seem to have no essential connection with the observation itself. Errors arising in this way are called *irregular* or *random*. . . . On the other hand, other sources of error by their nature *have a constant effect on all observations of the same class*. Or if the effect is not absolutely constant, its size varies regularly with circumstances that are essentially connected with the observations. These errors are called *constant* or *regular*.

Gauss further observes that ‘this distinction is to some extents relative and *depends on how broadly we take the notion of observations of the same class*.’ He explicitly excludes the consideration of systematic (regular, in his terminology) errors in his investigation and warns that ‘of course, it is up to the observer to ferret out all sources of constant error and remove them’.

This choice of neglecting systematic errors characterises the classical theory of errors and may be its main limitation. We show later that the need to overcome this limitation has been the driving force behind the studies on uncertainty in the second half of the 20th century. But for now let us stay with Gauss's approach and appreciate its merits. We thus come back to the *Theoria motus* to see how he deals with random errors.

He considers a special, but very important case of the general problem (1), the measurement of a single constant quantity x by repeated observations. In this case the model reads

$$\mathbf{y} = x + \mathbf{v} \quad (5)$$

and the probability distribution for the observations, given x , is

$$p(\mathbf{y}|x) = \prod_i p_v(y_i - x). \quad (6)$$

At this point, Gauss needs an explicit expression for the distribution of the errors p_v , and thus assumes some properties that correspond to the common understanding of measurement errors. He assumes that p_v is symmetric, maximum in its origin, and decreasing on each side of the origin. It may be either defined on a finite support, allowing for a maximum error, or rapidly tending to zero as the argument tends to infinity. Yet these assumptions are not enough to fully define the distribution p_v . Here is where Gauss makes a simple and genial move: he assumes that the most probable value for x , once the observations \mathbf{y} have been acquired, is the arithmetic mean of the observed values, because

it has been customary certainly to regard as an axiom the hypothesis that if any quantity has been determined by several direct observations, made under the same circumstances and with equal care, the arithmetic mean of the observed values affords the most probable value, if not rigorously, yet very nearly at least, so that it is always safe to adhere to it.

This key assumption may be explicated in this way:

$$\hat{x} = \bar{y} \triangleq N^{-1} \sum_i y_i. \quad (7)$$

On the basis of this assumption, Gauss is able to derive his celebrated *normal distribution*, which, in modern notation, reads

$$p(v) = (\sqrt{2\pi}\sigma)^{-1} \exp\left(-\frac{1}{2} \frac{v^2}{\sigma^2}\right), \quad (8)$$

where σ is the standard deviation³.

To sum up, Gauss was able to derive a *probabilistic model for random errors in measurement* which still maintains its validity [15]. During the same period, a similar result was reached, using a different route, by Laplace.

³ Instead of considering the standard deviation, Gauss elicits a precision measure, $h = \sqrt{2}\sigma^{-1}$ and discusses how interpercentile ranges depend upon it.

Laplace's approach and the theory of errors

A near contemporary to the *Theoria motus* was Pierre-Simon Marquis de Laplace's (1749–1827) *Théorie analytique des probabilités*, published in 1812 [4]. He derived the normal distribution in another way [14]. Let us consider again the case of repeated measurement, as described by model (5). We still assume that the errors v_i are independent and equally distributed, and we also require that their distribution $p(v)$ is symmetric about the origin and has a finite support. Let $\hat{x} = \bar{y}$ be the selected estimate for x and

$$e = \hat{x} - x \quad (9)$$

the estimation error. Then Laplace shows that e is asymptotically normally distributed with a variance proportional to N^{-1} . So we find here another way of deriving the normal distribution: it is the distribution of the estimation error, suitable for long series of observations.

Still another viewpoint may be considered, offered by the central limit theorem [30], traceable again, in a basic formulation, to Laplace [14]. Informally, the basic idea is to consider the measurement error as resulting from the contribution of a large sum of small independent error sources; that is,

$$v = \sum_j w_j. \quad (10)$$

If none of them prevails over the others, the distribution of the resulting error tends to be normal as long as the number of the error sources tends to infinity.

In conclusion, the classical theory of measurement errors, which is due to the contributions of Gauss and Laplace in the main, concerns random errors only and results in a probabilistic model, the normal distribution, whose validity may be supported by different arguments.

- It results from assumptions about the nature of errors (symmetry about the origin, probability of large errors quickly decreasing) plus the axiom that the arithmetical mean of the observations provides the most reliable estimate, for a series of measurements ‘made under the same circumstances and with equal care’.
- It is asymptotically attained when estimating a quantity after a long series of observations of the same quality.
- It also results by assuming the error is the consequence of a large number of error sources, none of which prevails over the others.

We reconsider the theory of measurement errors later on and discuss its merits and limitations, and how to overcome them. But now we have to come back to the original Gauss problem and see how it can be solved.

The origins of the least squares method

In order to solve the Gauss problem we have to find the value of the parameters $\hat{\mathbf{x}}$ that maximise the probability of the observations (3). Formula (4) provides an explicit expression for $p(\mathbf{y}|\mathbf{x})$: if we substitute the normal distribution (8) in it, we obtain

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi)^{\frac{1}{2}N} \sigma^N} \exp\left\{-\frac{1}{2\sigma^2} \sum [y_i - f_i(\mathbf{x})]^2\right\}. \quad (11)$$

So the value of \mathbf{x} we are looking for is *the one that minimizes the sum of the squares of the errors*,

$$\sum [y_i - f_i(\mathbf{x})]^2 = [\mathbf{y} - \mathbf{f}(\mathbf{x})]^T [\mathbf{y} - \mathbf{f}(\mathbf{x})] \quad (12)$$

This is how the least squares method appears in the *Theoria motus*. At this point, Gauss considers the linear version of his problem, which in modern notation is

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{v} \quad (13)$$

and provides the solution. We do not examine the original development, choosing simply to recall that, in modern notation, the solution may be obtained by pseudo-inversion

$$\hat{\mathbf{x}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}. \quad (14)$$

Finally, if σ^2 is the variance of the errors, the variance of the estimate will be

$$\text{Var}(\hat{\mathbf{x}}) = \sigma^2 (\mathbf{A}^T \mathbf{A})^{-1}. \quad (15)$$

Gauss reconsidered the least squares method in much more detail in the *Theoria combinationis*. There he provided another rationale for its use, which is no longer based on assuming a normal distribution for the errors, but rather on the minimisation of the expected mean-square error. In modern terminology, we would say that he presented the method as a way of obtaining a minimum-variance estimate. For our purposes the original derivation is more significant, because it is grounded in probability.

The theory of errors and the method of the least squares provided a great start for the theory of measurement and were the major results of the 19th century. At the beginning of the 20th century new ideas and methods became available to experimenters thanks to the contribution of ‘orthodox’ statistics [36].

1.2 Orthodox statistics

Experiments in metrology

Orthodox or classic is the name given to the statistics developed in the first part of the 20th century and whose principal exponent was Ronald Aylmer

Fisher (1890–1962) [8, 11, 23, 36]⁴. He was a geneticist and had the merit of explicitly addressing some of the main problems that experimenters encounter in their work. This is probably a reason for the success his approach encountered among experimenters, including metrologists. On the other hand the effectiveness of some of his methods, for example, his approach to the design of experiments, may have led to an overestimation of the value of other methods of orthodox statistics, such as their approach to point or interval estimation. A book published in 1964 by John Mandel [12], a statistics consultant of the National Bureau of Standards (NBS), provides a nice synthesis of the statistical tools and instruments available to the metrologists near the middle 1900s, which mainly refer to this school. Orthodox statistics promoted the development of probabilistic–statistical models by providing a store of methods for their use in conjunction with experimentation. Such methods include

- Criteria for the design of experiments, in order to optimise the information obtainable in a finite number of trials
- Methods for the estimation of parameters involved in the models
- Criteria for assessing the validity of the models

We have no room here for dealing with the design of experiments, which anyway is less central to our subject. We do, however, use a very simple example to illustrate the other two points. Consider the measurement of a single constant quantity by a series of n repeated observations as described by model (5). This model assumes that systematic effects are negligible, as generally admitted in the classic theory of errors. Suppose now that we have a set of m measuring instruments of the same type, independently calibrated. If we want to apply model (5) to them, we should consider whether, for example, the residual calibration error, which could give rise to a systematic effect, is really negligible. So we may perform a simple experiment that consists in measuring the same fixed quantity x with all the instruments at our disposal, and repeating the measurement n times for each instrument, thus collecting a total of $N = n \cdot m$ observations. We may wish to estimate the variance σ_v^2 , to check whether the hypothesis of negligible systematic effect is justified and, if not, to provide a quantitative estimate of the systematic effect. In order to do that, we have to consider a more general model than (5), that is,

$$y_{ij} = x + \theta_i + v_{ij}, \quad (16)$$

where

- $i = 1, \dots, m$ is the index denoting the instruments,
- $j = 1, \dots, n$ is the index denoting the repetitions,
- θ_i is a random variable representing the residual calibration error of each instrument,
- v_{ij} is an array of random variables representing independent realizations of the same normal variable v , the random error.

⁴ We use the term ‘orthodox’, because we prefer to reserve the term ‘classic’ for authors such as Gauss and Laplace.

Note that, as already pointed out by Gauss, the same phenomenon, the residual calibration error θ , gives rise to a systematic error if we consider as ‘observations of the same class’ the indications of a single instrument (index i fixed to, say, i_0), whilst it becomes a random variation if we sample instruments from the class of all the instrument of the same type (index i varying from 1 to m). Consider the following averages.

- Grand average, $\bar{y} = 1/N \sum_{ij} y_{ij}$, which is an estimate of x .
- Average per instrument, $\bar{y}_i = 1/n \sum_j y_{ij}$.
- Instrument deviations, $(\bar{y}_i - \bar{y})$, which is an estimate of θ_i .

The variance of v may be estimated by

$$\hat{\sigma}_v^2 = \frac{1}{N - m} \sum_{ij} (y_{ij} - \bar{y}_i)^2. \quad (17)$$

We now want to check whether the influence of the calibration errors θ_i is negligible. To do so, suppose that all the θ_i are null: we call this the *null hypothesis* and denote it by H_0 . If H_0 is true, we may estimate the variance of v also by

$$\hat{\sigma}_v^{2'} = \frac{1}{N - 1} \sum_{ij} (y_{ij} - \bar{y})^2 \quad (18)$$

and the result will be, more or less, the same as obtained by (17). We may then check whether the difference between $\hat{\sigma}_v^2$ and $\hat{\sigma}_v^{2'}$ is *significant*. If H_0 is true, the difference between $\hat{\sigma}_v^2$ and $\hat{\sigma}_v^{2'}$ is only due to the different degrees of freedom of the two estimates, namely, $\nu_1 = N - m$ for the former and $\nu_2 = N - 1$ for the latter. It may be proved that the ratio $\rho = \hat{\sigma}_v^2 / \hat{\sigma}_v^{2'}$, considered as a random variable, has an *F-Fisher* distribution, with parameters ν_1 and ν_2 . So a *significance test* may be performed. To do this, we first divide the space of the possible values of the ratio ρ into two *regions*, one with high probability or highly likely and the other with low probability or unlikely. Note that this is possible because, if H_0 holds, the distribution of ρ is *known*. Then we compute the value $\hat{\rho}$ that is the outcome of the data and we check in which region it falls.

- If it falls in the ‘unlikely’ region, we reject H_0 and thus conclude that accounting for θ makes a difference and thus the calibration error is not negligible.
- If it falls in the ‘likely’ region, we conclude that the difference between the two estimates $\hat{\sigma}_v^2$ and $\hat{\sigma}_v^{2'}$ may be due to their different degrees of freedom and so we may neglect the calibration error θ and adopt model (5)⁵.

In the case of the calibration error not being negligible, we may quantitatively evaluate its influence by estimating its variance through

⁵ We have presented significance testing very informally here, but we discuss it more fully later on.

$$\hat{\sigma}_\theta^2 = \frac{1}{m-1} \sum_i (\bar{y}_i - \bar{y})^2. \quad (19)$$

In conclusion, we think that this example, extremely simple in our opinion, may give an idea of what may be achieved through orthodox statistics in the development of experiments apt to characterise measuring systems. A much more elaborate example is presented by Mandel under the heading ‘systematic evaluation of measuring process’ [12], to which the reader is referred for further details on this subject. Yet before trying to formulate some conclusions on the contribution of orthodox statistics to metrology, we have to discuss two additional key points, namely *estimation* and *statistical testing*.

Estimation in Fisher’s view

Consider again Gauss’s problem, for simplicity, in the special case of measurement based on repeated observations, as in model (5). Gauss looked for *the value of \hat{x} having maximum probability, given the (vector) observation \mathbf{y}* , that is, for \hat{x} such that

$$p(\hat{x}|\mathbf{y}) = \max_x p(x|\mathbf{y}). \quad (20)$$

For calculating \hat{x} he used a Bayesian argument that is essentially equivalent to the following considerations. Applying the Bayes–Laplace rule⁶, we see that

$$p(x|\mathbf{y}) \propto p(\mathbf{y}|x)p(x). \quad (21)$$

If we assume a uniform distribution for x ⁷, we obtain

$$p(x|\mathbf{y}) \propto p(\mathbf{y}|x). \quad (22)$$

So maximizing $p(x|\mathbf{y})$ with respect to x , in view of (22), is computationally equivalent to maximizing $p(\mathbf{y}|x)$, that is, to searching \hat{x} such that

$$p(\mathbf{y}|\hat{x}) = \max_x p(\mathbf{y}|x). \quad (23)$$

But, for Gauss, *if formula (23) may be used for computing \hat{x} , the meaning of \hat{x} is still established by formula (20)*. Fisher changes this perspective completely, although obtaining, in this case, the same final result. He argues that the Bayes–Laplace rule (21) can not be applied, unless it is possible to determine an ‘objective’ prior distribution, $p(x)$, for x . So he applies directly formula (23), without deriving it from formula (20), as Gauss does. Although, in this case, the result is the same, the interpretation of the estimate is different: \hat{x} is interpreted now as the *most likely* value for x , that is, the value that maximizes the function

⁶ We discuss in some detail the Gauss–Laplace rule in Section 1.4.

⁷ The reason for assuming a uniform distribution is that prior to making the measurement, all possible values of x may be considered equally likely.

$$l(x|\mathbf{y}) = p(\mathbf{y}|x), \quad (24)$$

now called the *likelihood function*. Note that the likelihood function is *not* a probability distribution (because, in general, it does not integrate to unity).

There is a major difference between Gauss's and Fisher's approaches because in the former a probabilistic statement is made for \hat{x} : it is the value having maximum probability, once that \mathbf{y} has been observed, whilst in the latter no such probabilistic statement is possible. In this sense we may say that maximum-likelihood estimation is not a probabilistic estimation. Orthodox statistics consider other estimation methods, that we can not review here, but to which similar arguments apply [23, 31, 36].

Epistemological aspects of statistical tests

If Fisher's position on estimation is, in our opinion, not fully convincing, much more interesting is his view of statistical testing. To introduce this subject, let us consider an example of *significance testing*, simple but of high metrological import.

Suppose that we assume, for some measurement process, that model (5) holds. As we have already noted, the potentially critical assumption with this model is the absence of any (noticeable) systematic effect. This hypothesis implies that, for each observation y_i ,

$$E(y_i) = x, \quad (25)$$

where E is the expectation operator. A straightforward way of checking the validity of this hypothesis is to apply the measurement system to a standard object, whose value is known, $x = x_0$, with negligible uncertainty. In these conditions, the measurement process is described by

$$\mathbf{y} = x_0 + \mathbf{v}, \quad (26)$$

where \mathbf{v} is a vector of N independent, zero-mean, normal variables, with variances all equal to an unknown value σ^2 . So, if we take the arithmetic mean of the observations \bar{y} , we expect that it is almost equal to x_0 . The question is *how much may we allow \bar{y} to differ from x_0 , while still maintaining our model?*

To answer this question we observe that it is possible to prove that the *scaled distance*

$$d = \sqrt{N-1} \frac{\bar{y} - x_0}{\hat{\sigma}}, \quad (27)$$

where $\hat{\sigma}^2 = (N-1)^{-1} \sum (y_i - \bar{y})^2$ is an estimate of the variance, has a known distribution,

$$p(d) = p_{t, N-1}(d), \quad (28)$$

where $p_{t, \nu}(\cdot)$ is a t -Student distribution, with ν degrees of freedom. Then the acceptance region, that is, the region where the difference d is not critical for

our model is a *wide enough* interval around the origin. How wide should it be?

We must fix a small probability α (typical values are 0.05, 0.01, or 0.001) and then we identify the points $t_{\alpha/2}$, $t_{1-\alpha/2}$, such that

$$\int_{-\infty}^{t_{\alpha/2}} p_{t,N-1}(d)dd = \frac{\alpha}{2}, \quad \int_{t_{1-\alpha/2}}^{+\infty} p_{t,N-1}(d)dd = \frac{\alpha}{2}. \quad (29)$$

Then the acceptance region is $A = [t_{\alpha/2}, t_{1-\alpha/2}]$: the *a priori* probability that \bar{y} falls into the acceptance region is high, and is equal to $1 - \alpha$, whilst the probability that it falls outside is small. So, once we have made our experiment and have calculated the actual value of \bar{y} , if it falls into the acceptance region we maintain our model; we say that the test *has corroborated it*; otherwise we reject it, because what we have observed is *highly unlikely* under that model, and we consider the possibility of some systematic effect.

Let us then summarise the logic underlying significance testing: in general,

- We assume a probabilistic model, relying on an H_0 hypothesis and calculate the probabilistic distribution of the observations (or of some function of them, such as the arithmetic mean \bar{y} just considered),
- We partition the space of the observations into two regions, an acceptance region, where the observations are likely to occur, and a rejection region, where the observations are unlikely to occur; this partitioning is based on the assumption of a conventional value α , called the level of significance of the test,
- Then we conduct the experiment and if the observations fall into the acceptance region, we maintain the model, otherwise we abandon it.

This logical process may be synthetically called a *hypothetic–deductive inference* [23, 36]: hypothetic, because it starts from assuming a probabilistic model, and deductive because from the model it deduces the probability distribution for the observations, on which the test is based.

Significance testing plays a fundamental role in statistics, as much as *it is the only way of checking statistical models*. Although this way of testing statistical hypotheses, in an embryonic form, may be traced from the very dawn of probability and statistics [23], it was undoubtedly developed by orthodox statisticians. Yet some of them considered another class of statistical tests, *hypothesis tests*. Although in many textbooks they are put together with significance tests, their epistemological status is very different and we think it is wise to keep them distinct. We cannot discuss hypothesis testing thoroughly here, so we just mention it as it differs from significance testing.

In significance testing we test a statistical hypothesis *with respect to its negation: the result may be interpreted as a validation of a statistical model*.

In hypothesis testing we compare two alternative hypotheses, H_0 and H_1 , and although the way we treat both of them is not symmetric, at the end

of such a test we cannot reach absolute conclusions about H_0 , because such conclusions depend on the alternative hypothesis we have chosen.

So the domains of application of these tests are quite different: the former are more suited for *scientific investigation*, the latter for addressing *decision making*. This is, very briefly, the core of the criticism that Fisher directed towards hypothesis testing, that was instead supported by two other orthodox statisticians, Neyman and Pearson.

On this point we agree with Fisher's position.

Final remarks on the contribution of orthodox statistics to metrology

Orthodox statistics has been and still is very influential to metrology. Its contribution is manifold, as we have seen. In our opinion we may elicit two main contributions, namely

- *Addressing the design and evaluation of experiments*, by providing valuable tools for the design (via the design-of-experiments approach) and the evaluation of the influence of the various factors (through the analysis-of-variance method)
- Providing an invaluable tool for *checking statistical models*, by significance testing

On the other hand, the orthodox approach to estimation is, in our opinion, not fully satisfactory and its limit is even more apparent in the following section.

In the second half of the 20th century, when orthodox methods had reached their systematisation and were very popular among many experimenters, including the metrologists, the metrology community felt the need for a critical revision of its entire approach to uncertainty.

1.3 The Guide to the Expression of Uncertainty in Measurement

In the late 1970s, the metrological community recognised the need of reaching an internationally agreed way of expressing uncertainty in measurement. It also recognised the need to accompany the reporting of the result of any measurement by some quantitative indication of its quality, not only in primary metrology, but also in everyday measurements. So, in 1978, the Bureau International des Poids et Mesures (BIPM) carried out an investigation on a large number of laboratories and prepared Recommendation INC-1 (1980), whose guidelines were adopted by the Conference International des Poids et Mesures CIPM. Then an international working group was instituted (ISO/TAG 4/WG 3) for the development of a technical guide. One of the major scientific problems to be faced was the composition of random and systematic effects causing uncertainty. This also required an evolution in the concept of uncertainty

itself. The work of the group was paralleled by intensive scientific debate on such themes. In 1993 an important result was attained with the publication of the *Guide to the Expression of Uncertainty in Measurement* (GUM) [26]. The document had a great impact both on the technical and the scientific side and further stimulated international debate on measurement uncertainty and related topics. Good introductions to the GUM are already available [34] and here we only want to highlight some points that are particularly relevant to our subject and to introduce some trends of the contemporary debate on uncertainty that are the object of the next section.

As we have mentioned, the main problem to be faced was the composition of systematic and random effects in the evaluation of uncertainty. To do this the GUM chose to adopt the paradigm of *indirect measurements*, in which ‘the value of the measurand is obtained by measurement of other quantities functionally related to the measurand. This may be expressed as

$$x = g(\mathbf{z}), \quad (30)$$

where x is the measurand, \mathbf{z} a vector of input quantities, and g is a function. We call this expression the (*GUM*) *evaluation model* or *formula*. Basically it allows us to propagate the uncertainties on the quantities \mathbf{z} to the measurand x . Such uncertainties, in turn, may be evaluated on the basis of different pieces of information, which the GUM classifies in two main categories: those coming from a series of observations (type A) and those coming from other sources, such as information provided by the instrument manufacturers, by calibration, by experience, and so on (type B). So the focus moved from the type of the uncertainty sources (systematic versus random) to the type of information on them (type A versus type B). Consequently, it was possible to support, on a pragmatic basis, a common treatment for both of them.

Let us now see how can we deal with direct measurement, that is, measurement which is obtained from the output of a measuring instrument or, more generally, from a measuring system (MS). We may interpret one of the z_i , for example, the first one, as the indication y of the MS, that is, $z_1 = y$, and the remaining z_i as ‘corrections’, that should be ideally applied to correct the effect of the various error sources. The (possible) spread of the indications is accounted for by considering the variability of the random variable y . The evaluation procedure for the standard uncertainty then proceeds as follows.

The variables that appear in the evaluation formula (30) are regarded as random. So, if $\hat{\mathbf{z}}$ is a ‘best estimate’ of \mathbf{z} (which usually means that it is its expected value, $\hat{\mathbf{z}} = \mathbf{E}(\mathbf{z})$), $\Sigma_{\mathbf{z}}$ the covariance of \mathbf{z} and \mathbf{b} the vector of the sensitivities of x with respect to \mathbf{z} , calculated for $\mathbf{z} = \hat{\mathbf{z}}$, that is,

$$b_i = \left. \frac{\partial g}{\partial z_i} \right|_{\mathbf{z}=\hat{\mathbf{z}}}, \quad (31)$$

then an estimate of x may be obtained as

$$\hat{x} = g(\hat{\mathbf{z}}), \quad (32)$$

and the standard uncertainty, u , to be associated to \hat{x} , is

$$u = \sqrt{\mathbf{b}^T \Sigma_{\mathbf{z}} \mathbf{b}}. \quad (33)$$

The generalisation in the case of a vector measurand \mathbf{x} is not given explicitly, but is simple to obtain.

The GUM allows substantial discretion for choosing a formal statistical inference framework and concentrates mainly on practical aspects.

The debate on uncertainty, stimulated by the GUM, has also involved theoretical aspects. Bayesian inference was rapidly recognised as a sensible approach to the problems considered by the GUM, in particular when dealing with a combination of different sources of information. Prior to entering into the debate, we briefly review, in the next section, some of the main ideas of Bayesian inference.

1.4 Issues in the contemporary debate on measurement uncertainty

Bayesian estimation

Consider an experiment in which we perform repeated trials, in each of which an event E may occur or not. Let p be the probability of its occurrence in a single trial. Then the probability that E occurs m times in N repeated trials is

$$P(n_N = m | p) = \binom{N}{m} p^m (1-p)^{N-m}, \quad (34)$$

where n_N is the number of occurrences of event E in N trials. This result was obtained by Jacob Bernoulli and was one of the earliest findings in the theory of probability. Reverend Thomas Bayes (1702–1761) in his *Essay* [1], published posthumously in 1763, considered the problem which is inverse to the above: suppose that we do not know p and that we perform N trials of the experiment and find that event E occurs m times: how can we estimate the probability p ?

To estimate the parameter p Bayes intends to assign a probability distribution to it, that is, to find a rule for calculating the probability that the value of p falls in any assigned interval $[a, b]$, with $0 \leq a < b \leq 1$. He obtains the following formula.

$$P(a \leq p \leq b \mid n_N = m) = \frac{\int_a^b p^m (1-p)^{N-m} dp}{\int_0^1 p^m (1-p)^{N-m} dp}. \quad (35)$$

This result comes from assuming a uniform prior distribution for p , over its range $[0,1]$. Bayes justifies this assumption by observing that it is the proper one when ‘concerning such an event I have no reason to think that, in a certain number of trials, it should rather happen any one possible number of

times than another'. This is a formulation of the principle of *indifference* or of *insufficient reason*: when there is not sufficient reason for treating different possible cases in a different way, they should be treated in the same way [23].

The key idea underlying Bayes' solution was further investigated and generalised, perhaps independently, by Laplace in his *Essay on the probability of causes* [2] and then in his *Analytic theory of probability* [4]. He formulated what is now known as the Bayes–Laplace rule as follows,

if an event can be produced by a number n of different causes, then the probabilities of these causes given the event are to each other as the probabilities of the event given the causes, and the probability of the existence of each of these is equal to the probability of the event given the cause, divided by the sum of all the probabilities of the event given each of the causes.

In symbols, if we denote the i th cause by A_i , we have

$$P(A_i|E) = \frac{P(E|A_i)}{\sum_i P(E|A_i)}. \quad (36)$$

As we have seen, Gauss used a similar, Bayesian, argument in his *Theoria motus* and Laplace's treatise on probability was highly influential throughout the 19th century. On the other hand, orthodox statisticians disagreed on the possibility of calculating the probability of causes and preferred a different approach to estimation, as we have seen. Bayesian estimation became popular again due to the works of de Finetti, Ramsey, Jeffreys, and others [36].

In modern terms, a Bayesian estimation problem may be formulated and solved as follows [22]. Consider a series of observations \mathbf{y} , depending upon a unobservable parameter x : then the parameter x may be estimated by assigning a probability distribution to it, conditioned by the observations:

$$p(x|\mathbf{y}) = \frac{p(\mathbf{y}|x)p(x)}{\int p(\mathbf{y}|x)p(x) dx} \propto p(\mathbf{y}|x)p(x). \quad (37)$$

The probability distribution $p(x)$ is called the *prior distribution* for the parameter x and may, in general, incorporate prior knowledge about it. A special, but very important case for us, is that in which the indifference principle is used for assigning the $p(x)$, which is also called, in this case, a *vague prior*. If we adopt a vague prior, the essence of Bayesian estimation may be summarised as follows. *It is a probabilistic inference aimed at assigning a probability distribution to some quantity x on the basis of a set of observations \mathbf{y} and of an hypothesis on a probabilistic relation, $p(\mathbf{y}|x)$, linking the quantity and the observations.* As such it may be called a *hypothetic inductive inference* [23, 36].

Bayesian inference has been applied to measurement problems in different ways so far, corresponding to different ways of formulating the core hypothesis. We review them in the following sections.

Furthermore, it is useful to compare the definition above with that which we provided for *significance tests*: we regarded them as *hypothetic deductive inferences* instead. This distinction is essential for understanding the logic of the measurement process and we return to it at the end of this section.

A Bayesian approach to the evaluation of measurement uncertainty

A first Bayesian approach to the evaluation of uncertainty in measurement is due to Weise and Wöger [24], and other authors, and is documented by several papers illustrating its application to measurement problems (see [34] for a bibliography). We now try to summarise it, considering the presentation by Lira [34] in particular, with the important case of the direct measurement of a quantity by a MS, subject to both random variations and an additive systematic effect, by repeated observations.

If we assume, for maximum simplicity, that the indications of the MS are already properly scaled so that its response function is unitary, we may model this process as

$$\mathbf{y} = x + \theta + \mathbf{v}, \quad (38)$$

where x is the measurand, \mathbf{y} is a vector of N indications of the MS, θ is a unknown constant systematic effect, and \mathbf{v} is a vector of random variations, that are independent realizations of a normal random variable v , with zero mean and known⁸ variance σ^2 . Note that the structure is similar to a Gauss problem, but it differs from it due to the presence of the systematic effect θ . Note also that the structure is similar to that of formula (16) but with an important difference: here we have a single instrument and so we have no way of directly observing the effects of θ .

In this approach we need to begin with an evaluation equation which we may state in the following way,

$$x = g(\tilde{y}, z) = \tilde{y} + z, \quad (39)$$

where \tilde{y} is the mean value of y and z is the (ideal) ‘correction’ of the systematic effect θ ; that is, very simply,

$$z = -\theta. \quad (40)$$

Note that formula (39) is a special case of the GUM evaluation formula (30). Then we may consider the vector of parameters $[x, \tilde{y}, z]$ ⁹ and apply the Bayes–Laplace rule to it,

⁸ The case of an unknown variance may also be treated, but our aim here is to keep the example as simple as possible.

⁹ We may be surprised by the apparent dishomogeneity between \tilde{y} on one side and x and z on the other, because \tilde{y} is a mean value. The reason is that y varies during the repeated observations, whilst x and θ do not. So, to combine them in the same expression, we have to consider the mean value of y instead of its individual realizations. Actually here there is a criticality because the motivation of formula (39) in this approach is essentially heuristic.

$$p(x, \tilde{y}, z|\mathbf{y}) \propto p(\mathbf{y}|x, \tilde{y}, z)p(x, \tilde{y}, z). \quad (41)$$

Because the indications \mathbf{y} depend on x and z only through the expected value \tilde{y} , the formula simplifies as

$$p(x, \tilde{y}, z|\mathbf{y}) \propto p(\mathbf{y}|\tilde{y})p(x, \tilde{y}, z).$$

The joint distribution $p(x, \tilde{y}, z)$ may be factorised as

$$p(x, \tilde{y}, z) = p(x|\tilde{y}, z)p(\tilde{y}, z) = p(x|\tilde{y}, z)p(\tilde{y})p(z),$$

having further assumed the independence of \tilde{y} from z . With the evaluation equation above in mind, we obtain

$$p(x|\tilde{y}, z) = \delta(x - g(\tilde{y}, z)) = \delta(x - \tilde{y} - z),$$

where δ is the Dirac-delta operator. If we also assume a uniform prior for \tilde{y} , we have

$$p(x, \tilde{y}, z|\mathbf{y}) \propto p(\mathbf{y}|\tilde{y})\delta(x - \tilde{y} - z)p(z).$$

To reach the final distribution, we integrate out \tilde{y} and z and we obtain the marginal distribution

$$p(x|\mathbf{y}) \propto \int \int p(\mathbf{y}|\tilde{y})\delta(x - \tilde{y} - z)p(z) d\tilde{y}dz. \quad (42)$$

In order to proceed with the analytical calculations, let us now assume, as anticipated, that v is normal with known variance σ^2 . Then

$$p(\mathbf{y}|\tilde{y}) \propto \exp\left(-\frac{1}{2} \frac{(\tilde{y}-\bar{y})^2}{\sigma^2/N}\right)$$

and, finally,

$$p(x|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2} \frac{(x - \bar{y} - z)^2}{\sigma^2/N}\right) p(z) dz. \quad (43)$$

A distribution is thus assigned to the measurand on the basis of the observations \mathbf{y} and of the following hypotheses.

- A probabilistic model for the observations, $p(\mathbf{y}|\tilde{y})$
- The evaluation equation $x = g(\tilde{y}, z)$
- A probability distribution for z , $p(z)$

This approach requires the previous assumption of an evaluation equation, which relies on an essentially heuristic basis (see Footnote 9). In the next section we consider a different approach: we first present a general probabilistic model of the measurement process and then we consider an alternative Bayesian approach, based on that model.

A probabilistic model of the measurement process

Recently a general probabilistic model of the measurement process (MP) has been proposed [35, 42]. It starts from the basic consideration that measurement is performed through a measuring system (MS) [25] and envisages a general functional description of it. The MS interacts with the measurand and produces an observable output, the indication, which is related to or, in other words, is caused by, the value of the measurand. So it is quite natural to describe the behaviour of the MS by an input–output model, whose input is the value of the measurand and whose output is the instrument indication. Such an input–output relationship may be experimentally determined by calibration. When we perform a measurement, we get an indication on the basis of which we are able to identify, within the uncertainty limitations, the value of the measurand, because we know in advance, thanks to the calibration, the cause–effect relation linking the two. The measurement process may thus be broken down into two subprocesses, namely:

- *Observation*, the process of producing an observable output that is caused by the measurand and depends on its value
- *Restitution*, the process of identifying the value of the measurand from the indication(s) of the MS

Consequently, *measurement* may be viewed as the process resulting from the chaining of observation and restitution and allowing a value (the measurement value) to be assigned to the measurand. Observation is always performed by the MS, whilst restitution may either be embedded in the MS or performed off-line, depending upon the technology. In any case it seems conceptually correct to distinguish between the two, because the former is a chain of physical transformations, whilst the latter is a kind of information processing. We also show how this distinction is practical, as much as it permits the development of a systematic approach to the modelling of measurement processes, which allows the final result of measurement to be expressed as a probability distribution over the set of the possible values of the measurand. Let us now see how can we describe all this in deterministic terms first. This may be seen as the description of the ideal MP and paves the way to the presentation of the probabilistic model. In a deterministic framework, *observation* may be described by a function that expresses the cause–effect relationship holding between the value of the measurand x and the indication y ; that is,

$$y = f(x). \tag{44}$$

The function f may be called the *response characteristic* of the MS, because it expresses the input–output behaviour of the MS, or also *calibration function*, because it may be experimentally determined by calibration [25]. For example, if x is the temperature, t , of an object, and the MS is a measuring chain made of a thermocouple, an amplifier, and a voltmeter, then y is the voltage reading, V , from the voltmeter and f includes the (direct) thermo-electric function for the thermocouple, f' , and the gain of the amplifier, A . So we have

$$V = Af'(t).$$

This deterministic description is ‘ideal’ in that we assume that the measuring system behaves exactly according to its response function f and that no other quantity influences the measurement.

Restitution, on the other hand, may be viewed as the inversion of observation, because for any indication y we provide the measurement value \hat{x} by

$$\hat{x} = f^{-1}(y). \quad (45)$$

In our example, we have

$$\hat{t} = A^{-1}f'^{-1}(V/A).$$

Note that for standard thermocouples both f' and f'^{-1} are standardised functions (polynomials).

In a traditional environment, restitution may be performed manually, whilst in a computerised measuring process it is performed automatically. Anyway, irrespective of the technology, the concept is the same.

Finally, *measurement* is the concatenation of the two transformations,

$$\hat{x} = f^{-1}[f(x)] = x, \quad (46)$$

and results in a unitary transformation, due to the fact that the deterministic model provides a description of an ideal MP. The meaning of this last equation is the following. If the MS behaved exactly according to its characteristic function f , and no other uncertainty cause applied, then we would obtain the exact value of the measurand. Of course this is not the case in real measurements, but this ideal scheme allows us to introduce the probabilistic framework, which instead properly represents an uncertain environment.

The results thus far obtained are summarised in Table 1, under the ‘deterministic model’ heading. We provide additional arguments in support of this model later on, in Section 2.2.

The probabilistic model may be obtained by translating what we have thus far exposed in probabilistic terms. The natural probabilistic description of the observation process, that is, the natural counterpart of the calibration function, is provided by the conditional probability distribution

$$p(y|x),$$

Table 1. Comparison between the deterministic model and the probabilistic one.

Process/ Subprocess	Deterministic Model	Probabilistic Model
Observation	$y = f(x)$	$p(\mathbf{y} x, \boldsymbol{\theta})$
Restitution	$\hat{x} = f^{-1}(y)$	$p(x \mathbf{y}) = \int_{\Theta} p(\mathbf{y} x, \boldsymbol{\theta}) [\int_X p(\mathbf{y} x, \boldsymbol{\theta}) dx]^{-1} p(\boldsymbol{\theta}) d\boldsymbol{\theta}$
Measurement	$\hat{x} = f^{-1}[f(x)] = x$	$p(\hat{x} x) = \int_Y \delta[\hat{x} - E(x \mathbf{y})] [\int_{\Theta} p(\mathbf{y} x, \boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}] d\mathbf{y}$

where
 $x \in X$: value of the measurand
 $\mathbf{y} \in \mathbf{Y}$: N -dimensional observation vector
 $\boldsymbol{\theta} \in \Theta$: K -dimensional parameter vector
 $\hat{x} \in X$: measurement value

that is, the probability distribution of the indication y , for any given measurement value x . In other words, whilst in the deterministic model for each value of the measurand we get one and only one indication, in the probabilistic case we may obtain a plurality of indications, ruled by a probability distribution. Consequently, *restitution may be described as the probabilistic inversion of the transformation defining observation*. Such an inversion may be performed according to the Bayes–Laplace rule. If we assume a uniform prior for x , we obtain

$$p(x|y) = \frac{p(y|x)}{\int p(y|x) dx} \propto p(y|x).$$

In this way we account for random variations, but how can we deal with systematic effects? The systematic effect (of any type, additive, multiplicative, etc.,) of an influence quantity θ may be expressed *by allowing the distribution $p(y|x)$ to also be conditioned by θ* , thus becoming

$$p(y|x, \theta).$$

If we now apply the Bayes–Laplace rule, the result will be still conditioned by θ ; that is,

$$p(x|y, \theta) \propto p(y|x, \theta).$$

To attain the final distribution $p(x|y)$ it is sufficient to ‘decondition’ with respect to θ , by applying the principle of total probability; that is,

$$p(x|y) \propto \int p(y|x, \theta) p(\theta) d\theta.$$

The generalisation to a vector of observations \mathbf{y} and to a vector of influence parameters is immediate and yields, for observation

$$p(\mathbf{y}|x, \boldsymbol{\theta}) \quad (47)$$

and for restitution

$$p(x|\mathbf{y}) \propto \int p(\mathbf{y}|x, \boldsymbol{\theta})p(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (48)$$

Restitution yields a probabilistic distribution rather than a single value, as happened in the deterministic case. Yet even now it is possible to define a single measurement value as

$$\hat{x} = E(x|\mathbf{y}). \quad (49)$$

This formula *provides the most general definition of \hat{x}* as a function of the indications \mathbf{y} . When it is possible to make it explicit as a function of $\boldsymbol{\theta}$ also, we obtain

$$\hat{x} = h(\mathbf{y}, \boldsymbol{\theta}), \quad (50)$$

which is another, essentially equivalent, way of expressing the evaluation equation envisaged by the GUM, but with an important difference: such an equation *is now derived by the model*.

Finally, the overall measurement process may be described by combining (chaining) observation and restitution in order to obtain the distribution of the measurement value for each possible value of the measurand, that is, $p(\hat{x}|x)$. This may be done by observing that the measurement value \hat{x} is a function of the indications \mathbf{y} , which, in turn, in observation are regarded as a vector random variable, conditioned by x . So, applying a formula for the propagation of distributions, we obtain

$$p(\hat{x}|x) = \int_{\mathbf{Y}} \delta[\hat{x} - E(x|\mathbf{y})] \int_{\boldsymbol{\Theta}} [p(\mathbf{y}|x, \boldsymbol{\theta})p(\boldsymbol{\theta}) d\boldsymbol{\theta}] d\mathbf{y} \quad (51)$$

which describes the overall *measurement* process. This completes the set of formulas of the probabilistic model, collected in Table 1, where they are compared with the corresponding deterministic ones. Further generalisations are possible (e.g., considering a vector measurand [35]), but are not treated here, for the sake of simplicity.

Note the different meaning of the distributions $p(x|\mathbf{y})$ and $p(\hat{x}|x)$. The former, $p(x|\mathbf{y})$, is the distribution that describes restitution: whenever we observe the (vector, in general) indication \mathbf{y} we may assign to the measurand the distribution $p(x|\mathbf{y})$. This distribution is thus *the basis for providing the measurement value, $\hat{x} = E(x|\mathbf{y})$, and its uncertainty*. For example, the standard uncertainty may be defined as

$$u = \sqrt{Var(x|\mathbf{y})}, \quad (52)$$

and the expanded uncertainty, at a coverage level p_0 , as the value $U > 0$ such that

$$\int_{\hat{x}-U}^{\hat{x}+U} p(x|\mathbf{y}) dx = p_0. \quad (53)$$

Instead, $p(\hat{x}|x)$ is the distribution that describes the overall measurement process and relates the measurement value \hat{x} to the value of the measurand x . Then it may be used *for declaring the performance of the measuring system*.

In summary, we may say that the two distributions considered thus far, $p(x|\mathbf{y})$ and $p(\hat{x}|x)$, are *complementary in meaning and purpose*: the former is the basis for expressing the uncertainty in a specific measurement; the latter is useful for expressing the performance of the measuring system in its measuring range (spanned by x) [39,42]. We may say that the former is of primary interest for instrument users and the latter for instrument producers.

Let us now look at the application of this model to the example discussed in the previous section, that is, the measurement of a constant quantity x , by repeated observations from an instrument affected both by random variations and systematic additive deviations, according to the model (38)

$$\mathbf{y} = x + \theta + \mathbf{v}.$$

For observation, we obtain

$$p(\mathbf{y}|x, \theta) = \prod_i p_v(y_i - \theta - x) \quad (54)$$

and for restitution

$$p(x|\mathbf{y}) \propto \int \prod_i p_v(y_i - \theta - x) p(\theta) d\theta. \quad (55)$$

When the distribution of v is normal with known variance σ^2 , we have

$$p(x|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2} \frac{(x - \bar{y} + \theta)^2}{\sigma^2/N}\right) p(\theta) d\theta, \quad (56)$$

which is equivalent to formula (43), because $z = -\theta$.

The derivation is now more straightforward and there is no need to assume an evaluation equation, rather the results derive directly from the model (38).

Lastly, the distribution that characterises the overall measurement process is

$$p(\hat{x}|x) \propto \int \exp\left(-\frac{1}{2} \frac{(\hat{x} - x - \theta)^2}{\sigma^2/N}\right) p(\theta) d\theta. \quad (57)$$

In the accompanying DVD numerical examples of the application of this model to uncertainty evaluation and to risk analysis are provided, with the related software codes. Additional examples may be found in References [34, 39, 40].

Some notes on inference in measurement

We call *probabilistic* an *inference* that yields the assignment of a probability distribution to some parameter under investigation.

We have encountered two kinds of such inferences so far, namely

- (a) *Hypothetic deductive* inferences, in significance testing
- (b) *Hypothetic inductive* inferences, in Bayesian estimation

Let us briefly recall their logical structure. In a hypothetic deductive inference

- (a1) We hypothesise a probabilistic model,
- (a2) On the basis of which we deduce the probability distribution of the observation \mathbf{y} in a given experiment, which allows us
- (a3) To define an acceptance region A for the observation, which is a region in which the observation complies with the model;
- (a4) Then we perform the experiment and acquire \mathbf{y} :
 - If \mathbf{y} falls into the acceptance region, the model is corroborated,
 - Otherwise it is ‘falsified’ by the observation and we may consider abandoning it.

In a hypothetic inductive inference, instead, if we consider the most important case for us, that of assuming a noninformative prior,

- (b1) We hypothesise a probabilistic relation, in a given experiment, linking the observation \mathbf{y} to a parameter x , expressed as a conditional distribution $p(\mathbf{y}|x)$;
- (b2) We perform the experiment and acquire the observation \mathbf{y} ;
- (b3) On the basis of the observation and of the hypothesised probabilistic relation, we assign a probability distribution to x , induced through the observation.

Let us now considered the logical structure of the measurement process, as outlined in the previous section. It includes the following steps.

- (c1) Assume a probabilistic relation between the value of the measurand and the indications of the MS, parametrical in respect to some influence parameters: this relation is a model of the observation process;
- (c2) Assume a probability measure over the space of the influence parameters;
- (c3) Perform observation and acquire the indications of the MS;
- (c4) Apply, in the restitution phase, the Bayes–Laplace rule and obtain a probability distribution for the measurand, still conditioned upon the influence parameters;
- (c5) Decondition the probability distribution with respect to the influence parameters, which concludes the restitution phase and the overall measurement process.

If we analyse this procedure in the light of what we have so far exposed, we recognise in steps c1, c3, and c4 a Bayesian inference, so that we may say that the measurement process *embeds* a Bayesian inference.

On the other hand, we also note that steps c2 and c5 are not typical of a Bayesian inference. They include the assumption of a probability distribution for some parameters (step c2) and their use according to the rules of the calculus of probability (step c5). We say that these two steps form a *hypothetic-deductive process*¹⁰: so we conclude that in general *in a measurement process we have the combination of a hypothetic-inductive inference and of a hypothetic-deductive process*.

This conclusion does not apply only to the approach based on the probabilistic model of the MS, but also to the previous approach, based on formula (39). Even in that case no inference is made on the influence parameters giving rise to systematic effects: actually no inference is possible because their effects are not observable via the indications.

We thus now have a new way of posing the problem of systematic effects. Because influence parameters giving rise to systematic effects must be treated via a hypothetic deductive process, what guarantees the validity of the final measurement result?

This question is a special case of the general requirement for scientific statements: they must be *controllable*, as it must be possible to design and perform experiments whose results may falsify such theories. This principle, the *falsifiability* of scientific theories, is central to Popper's epistemology and widely accepted [36].

So what can we do in the case of measurement?

The answer, from what we have seen so far, is simple and straightforward: the validity of the measurement process, which includes a hypothetic-deductive treatment of systematic effects, may be controlled by a significance test, that is, by a hypothetic-deductive inference.

Let us briefly see how this inference can be stated. Consider a measurement process described by $p(\hat{x}|x)$. Remember that this distribution accounts for systematic effects too. Suppose that we dispose of a standard whose value, x_0 , is known with uncertainty negligible for our purpose. Then we can measure the standard through the measurement process under consideration and perform a significance test on the difference $\hat{x}_0 - x_0$, where \hat{x}_0 is the measurement value obtained after measuring the standard. For a significance level α , the acceptance region will be $A = [-a, +a]$, such that

$$\int_{-a}^{+a} p(\hat{x}_0 - x_0|x_0) d\hat{x} = \alpha. \quad (58)$$

¹⁰ We distinguish between a hypothetic-deductive process and a hypothetic-deductive inference: in the latter we learn from experience, whilst in the former we do not. We show how to apply a hypothetic deductive inference to measurement in a moment.

This procedure formalises what is done in the practice of metrology, for example, in the verification of the calibration of a MS or in the control of a measurement process by check standards.

2 Towards a probabilistic theory of measurement

2.1 Origin and early development of the formal theory of measurement

Helmholtz

So far we have considered a mainstream of studies, centered on the problem of measurement uncertainty, from Gauss up to contemporary practice. In reality, there is another very important area of studies in the science of measurement that arose towards the end of the 19th century and concerns the problem of the foundations of measurement and the development of a formal theory for it. These two mainstreams, although conceptually related, have developed essentially in parallel and with few connections. The reason for this lack of connection is historical and we think that, at present, a merger of these two approaches is much needed and is a major challenge for metrology. We thus briefly review some of the main steps in the historical development of measurement theory and then we show why this theory also requires a probabilistic approach. We then overview what has been done so far and discuss what we may expect in the near future [41].

The beginning of the modern theory of measurement is usually traced to a genial work by Helmholtz, *‘Counting and Measuring from the Viewpoint of the Theory of Knowledge,’* published in 1887 [6]. In this essay he poses the problem of the foundation of measurement, because he investigates ‘the objective meaning of the fact that we express as quantities, through *concrete numbers*, situations of real objects’ and he discusses ‘under what circumstances we are allowed to do so.’ ‘Concrete numbers’, in his language, are those arising from the counting of real objects.

He finds a brilliant solution to the problem by establishing an analogy between measurement and counting. The key idea is that in many cases what we want to measure is literally a ‘quantity,’ in the sense that it is the amount of something, and thus it may be considered to be composed of the sum of a number of elementary parts, or units, of that something. In these cases measurement is equivalent to the counting of such units.

Counting is possible thanks to the properties of natural numbers which undergo an order based on the relation ‘greater than or equal to,’ denoted by \geq , and may be added to each other by an addition operation, denoted by $+$.

Similarly, measurement is possible and well founded whenever it is possible to identify the empirical counterparts of the order relation and of the addition operation for the objects carrying the characteristic of interest.

The main idea of Helmholtz, that measurement represents properties of objects by assigning numbers to them in such a way as to reproduce empirical relations in the numerical domain, has been the basis for the development of a theory of measurement.

Campbell

The first organic presentation of a theory of measurement was by Norman Campbell, in the second part of his book, *Physics: The Elements* [7], published in 1920. Like Helmholtz, he considers the problem of ‘Why can and do we measure some properties of bodies while we do not measure others’ and goes further in this investigation by asking, ‘What is the difference between the properties which determine the possibility or impossibility of measuring them.’ In order to answer this question, he distinguishes two main kinds of quantities, *fundamental*, such as mass or length, and *derived*, such as density, for example. Both of them require an empirical property of *order*, which is – as for Helmholtz – the basic requirement for measurement. But fundamental quantities allow for a *physical-addition* operation also. *Why is this operation so important?*

Because it is the key to permitting *the general procedure for fundamental measurement* to be applied. Such a procedure consists in constructing a *measurement scale*, that is, a *series of standards* with properly assigned numerical values, and then in comparing any unknown object, r , to it, in order to select the element in the series which is equivalent to it. Then it will be possible to assign to r the same number (measure) as the selected element.

The physical addition operation must satisfy – as Helmholtz had already pointed out – the logical properties of addition, that is, the associative and the commutative properties, and there must be experimental evidence of this.

On the other hand, derived quantities do not require a specific scale to be devised, because they may be measured thanks to a physical law relating them to other measurable quantities. In the case of density ρ , for example, we may define it as the ratio of mass to volume, that is, $\rho = m/V$, and thus reduce its measurement to one of mass and volume.

Campbell’s investigation contributed to a deeper understanding of the nature of measurement and his position has been, and still is, highly influential and he was deeply involved in a controversy that arose in the 1930s in the scientific community and that would strongly influence the development of measurement science.

The Committee of the British Association for the Advancement of Science

In the beginning of the 20th century measurement was popular not only in physics and engineering, but also in experimental psychology and in behavioural sciences. So it was quite natural for the scientific community

to consider whether the epistemological value of measurement in this new domain was well founded. With this aim, in 1932, the British Association for the Advancement of Science appointed a committee composed of physicists and psychologists, to consider and report upon the possibility of quantitative estimates of sensory events. The report of the committee, published only in 1939, after years of discussions, admitted that it had been impossible for the two sides to reach a common understanding of measurement [9]. The physicists, in particular, took a strong stance against the possibility of actually making measurements in the behavioural sciences.

Without entering into detail, the committee considered typical psychophysical experiments on ‘just perceptible differences’ and on ‘equal appearing intervals.’ The psychologists claimed that from those experiments and by assuming some feasible psychophysical law, such as Fechner’s law, it was possible to arrive at quantifying sensations. The physicists, instead, denied that, mainly because, in their opinion, direct estimation of sensations was not possible and additivity was inconceivable for them.

The report of the committee had an enormous influence in the following years and we may say that it led to an essentially parallel development of measurement science in physical science on one side and in behavioural sciences on the other, with consequences up to the present day. But let us see now some reactions from an outstanding psychologist, Stanley Stevens, who did not himself attend the committee, although his work on loudness was thoroughly discussed by them.

Stevens

Stevens, who was at Harvard and was dealing with problems similar to those considered by the committee in the same period, felt the need for a more general theory of measurement. This generalisation was aimed at enlarging the number of feasible measurement scales [10]. In Campbell’s view there was only one type of measurement scale, the one holding for quantities for which an empirical operation of addition was possible. Stevens instead proposed his famous fourfold classification of measurement scales, which is still in use and is summarised in Table 2. The classification is based on the notion of admissible transformations, that is, transformations that leave the scale form invariant. In the table we may see the scale types (column 3) and the groups of admissible transformations (column 5).¹¹ In doing so he shifted the focus from empirical relations, such as order, additivity..., to the *invariance properties* of the scales.

Nominal scales are involved in classification operations and numbers serve only to distinguish one class of objects from another. Any *biunivocal* transformation is permissible, because identification is still possible. Examples are colour measurements and pattern recognition techniques.

¹¹ The content of the other columns is presented later on.

Table 2. Summary of the main scales for fundamental measurement, as considered in the representational theory, based on the original classification by Stevens.

Empirical Structure	Empirical Relations	Scale Type	Representation	Admissible Transformations
Nominal	Equivalence among elements in each class	Nominal	$a \sim b \iff m(a) = m(b)$	Biunivocal
Order	Weak order among the objects	Ordinal	$a \succ b \iff m(a) \geq m(b)$	Monotone increasing
Difference	As above plus weak order among intervals	Interval	$\Delta_{ab} \succ \Delta_{cd} \iff m(a) - m(b) \geq m(c) - m(d)$	Linear positive $m' = \alpha m + \beta$ $\alpha > 0$
Extensive	As above plus a concatenation operation	Ratio	$a \sim b \circ c \iff m(a) = m(b) + m(c)$	Similarity $m' = \alpha m$ $\alpha > 0$

Ordinal scales permit a rank ordering of objects and remain invariant under *monotonic increasing* transformations. They include hardness of minerals and earthquake or wind intensity.

Interval scales entail a constant unit of measurement; that is, they introduce a metric, and so permit the calculation of differences between any two values. They remain invariant under *linear positive* transformations. Fahrenheit or Celsius temperatures are good examples, as well as position or time, intended as calendar.

Ratio scales also feature constant units of measurement, but, in addition, they allow the ratio of two values to be evaluated, because a true zero exists. They are invariant under any simply multiplicative transformation, or *similarity*. They include ‘extensive’ quantities, such as mass or length, but also, in Stevens’ view, perceptual quantities, such as loudness or brightness

So in order to overcome the position of the physicists, Stevens generalises the notion of measurement scale, already introduced by Campbell for fundamental measurements. Moreover, he argues that direct estimation of sensations is possible, as happens in *magnitude estimation*. Such a test may be performed, for example, by presenting a line of a given length and telling the observer to call it some number, say, 10. Then a line of some other length is presented and the subject is asked to assign it a number, considering that the first line was 10 and so forth. The important point is that, thanks to such tests, it is possible for Stevens to consider *equality between ratios, as the empirical relation for ratio scales, instead of addition*.

Summarising, Stevens proposes to overcome the severe limitation in measurability posed by the report of the British Association, by increasing the number of allowable measurement scales and by considering equality of ratio as an empirical relation. Yet his arguments did not convince the physicists.

With Stevens we have reached the second half of the 20th century. At that time a considerable body of results had been obtained in measurement theory and there was a need for a systematization, which was achieved with the representational theory of measurement.

2.2 The representational theory of measurement

The representational framework

A remarkable systematisation of the formal theory of measurement was achieved in the second half of the 20th century and referred to as representational theory. A comprehensive presentation is offered in the gigantic treatise, *Foundations of Measurement*, by Krantz, Luce, Suppes, and Tversky [13], as well as in other parallel works, such as those by Roberts [16] and Narens [20]. These studies share a common framework, which essentially may be seen as a combination of the viewpoints of Campbell and Stevens, that are seen as complementary rather than opposing. The main idea, traceable, as we have seen, to Helmholtz, is that the numbers we obtain through measurement represent empirical relations. This framework also applies to fundamental physical measurements as intended by Campbell, here called extensive. But now *extensive* is regarded as *a special*, though very important, kind of measurement, not as *the only* one worthy of this name. Consequently, the classification of scales proposed by Stevens may be retained and each scale is now characterized by

1. A *representation theorem*, showing how empirical relations are mapped into corresponding numerical relations
2. A *uniqueness theorem*, specifying which class of transformations maintain the properties of the scale

The uniqueness theorem allows the meaningfulness of statements concerning measurement to be addressed. In fact, we may say that a statement concerning the results of measurement on a given scale, is meaningful if its truth is unaffected by admissible transformations on that scale.

A summary of the representation framework has been presented in Table 2 above. We have already discussed the invariance properties of the scales, here called uniqueness conditions, when presenting Stevens' contribution. Let us now briefly comment on empirical structures (column 1), the associated empirical relations (column 2), and representation theorem (column 4).

In *nominal structures* we only have the *equivalence* of elements belonging to the same class, denoted by \sim and the result of operating on this scale is a classification.

Order structures are characterised by an empirical relation of *weak order* that we denote by the symbol \succsim . The relation of weak order plays a fundamental role in measurement and is satisfied also in the other scales to follow.

In the case of *difference structures* we are mainly concerned with *intervals* of objects and with a weak order relation among them. For example, if a, b are two elements of A , their interval will be denoted by Δ_{ab} (being positive if $a \succ b$) and $\Delta_{ab} \succsim \Delta_{cd}$ means that the first interval is, empirically, greater than or equal to the second one.

Finally, in *extensive structures* an empirical *concatenation* operation, or physical addition as Campbell named it, is present. In connection with this operation we also define a ternary relation $a \sim b \circ c$, meaning that the element a is equivalent to the empirical sum of b plus c .

The main representation theorem for the three structures we are dealing with thus reads as follows.

- For order structures:

$$a \succsim b \iff m(a) \geq m(b), \quad (59)$$

- For difference structures:

$$\Delta_{ab} \succsim \Delta_{cd} \iff m(a) - m(b) \geq m(c) - m(d), \quad (60)$$

- For extensive structures:

$$a \sim b \circ c \iff m(a) = m(b) + m(c). \quad (61)$$

Each structure includes the properties of the previous ones, so, for example, difference structures also satisfy the representation theorem that holds for order structures and so on.

The representational theory has been developed mainly in the field of behavioural sciences but has been brought to the attention of physicists and engineers since the 1970s, mainly by Finkelstein [19], and has received, afterwards, contributions also from that community. Such a theory was initially stated in an essentially algebraic fashion and, until very recently, it has only partially been treated in probabilistic terms. Moreover, due to its growth, especially in the field of behavioural science, little or no attention has been paid to the role of the measuring instrument (or system). We thus now briefly review some of the probabilistic developments proposed in the representational approach, leaving for the next section the presentation of a complete probabilistic theory of measurement that also accounts for the role of the measuring system.

Probabilistic issues

Because uncertainty is a constitutive aspect of measurement, several attempts have been made to include probabilistic issues in the representational

framework. A basic bibliography may be found in Roberts [16], Chapter 6], Krantz *et al.* [13, Vol. 2, Chapters 16–17], and in Luce and Suppes [33]. Although we cannot survey these references in detail, we may say that they mainly deal with a probabilistic treatment of comparison tests aimed at constructing order scales. The general problem that they consider may be formulated in the following terms. Given a probabilistic description of the empirical relations, under which conditions is it possible to arrive at a meaningful representation? This issue is also called *probabilistic consistency*.

To illustrate the problem, suppose we want to measure some perceived quantity, for example, the intensity of a class of sounds, and we want to construct an order scale for them [37]. Empirical relations are in this case defined by the responses of a class of subjects to the sounds under investigation. If we consider any two sounds, a , b , we cannot expect in general that a definite relation, for example, $a \succ b$, where \succ here means ‘louder than’, definitely holds for them, due to inter- and intrasubjects variability. We may rather expect that a probability may be attached to such a relation; that is, $\mathbb{P}(a \succ b)$, where \mathbb{P} denotes the probability of a relation.

What property should we require for this probability in order for representation to be possible? We have seen that for deterministic order scales the key property is transitivity; that is, for a , b , c belonging to A , if both $a \succ b$ and $b \succ c$ hold, $a \succ c$ should also hold. A suitable probabilistic replacement for this is the so-called *weak probabilistic transitivity*; that is,

$$\mathbb{P}(a \succ b) \geq \mathbb{P}(a \prec b) \quad \text{and} \quad \mathbb{P}(b \succ c) \geq \mathbb{P}(b \prec c) \quad \implies \mathbb{P}(a \succ c) \geq \mathbb{P}(a \prec c). \quad (62)$$

If this property holds, the following representation theorem may be proved [27].

$$m(a) \geq m(b) \iff \mathbb{P}(a \succ b) \geq \mathbb{P}(a \prec b). \quad (63)$$

Yet results such as this one, although important in some regards, are not completely satisfactory from a general standpoint. In fact, this means that in order to assign measures to the objects we must know the probability of the empirical relations. Empirical relations are no longer regarded as deterministic, but the measures may still be assigned exactly, once the required probabilities are known. Measurement is still, in some sense, a deterministic process, although it is no longer based *directly* on empirical relations but rather, *indirectly*, on their probabilities.

This is not yet what is needed for a completely satisfactory probabilistic theory of measurement. The need for a complete probabilistic formulation of the representational framework was clearly expressed by Leaning and Finkelstein [18], where a probabilistic framework was also proposed. Unfortunately that approach was not developed in detail in the following years. The most famous paper on the probabilistic approach is perhaps one from Falmagne [17], in which the author presented a probabilistic representation for extensive measurements, in the case that some special constitutive relations hold between

relational probabilities and measure values. The paper was a novelty at that time and had the merit of dealing with the quite complex structure of extensive measurement, but it dealt with special cases only.

Actually, there is an inherent difficulty in achieving a complete probabilistic formulation, partially documented in the studies concerning the so-called *random-utility model* (see Roberts [16], Section 6.2). Another difficulty is related to the notion of the *probability of a relation* (and of a relational system), which has been scarcely investigated in the past [36]. Only relatively recently has a major contribution been provided on this subject in two papers, by Regenwetter [28] and Regenwetter and Marley [32], which provide useful results. We account for them, later on, in proposing a full probabilistic theory, but, before that, we have to discuss the role of the measuring system in a formal theory of measurement.

The role of the measuring system in a formal theory of measurement

Although measuring instruments have been key players in the development of modern science, their role in the theory of measurement seems to have been underestimated [41]. Campbell, for example, as a physicist, was aware of their importance, yet he concentrated mainly on the problem of scale construction and considered the issues related to the measuring system as technical rather than theoretical ones. Lately, the formal theory of measurement has been developed mainly in the area of behavioural sciences, where the role of the measuring system is not felt as central. The need of explicitly accounting for the role of the measuring system in a theory of measurement has been pointed out by Gonella [21] and, more recently and with additional arguments, by Mari [29], who claims that measurement is essentially an evaluation performed by a calibrated MS. We essentially agree with this position, *provided that a proper (broad) definition of MS is adopted*. Let us then discuss this point in more detail.

In principle, measurement may be performed by selecting one object, a , and comparing it with a previously established measurement scale in order to identify a standard s *to which the element a is equivalent*. After that, we assign $m(a) = m(s)$. But how can we actually do that?

Measurement is in general performed *through the mediation of a measuring system*. So *what is the role and the behaviour of the MS precisely?*

In Section 1.4 we have observed that, because the value of the measurand is not directly observable, the function of the MS is to interact with the object and to produce, as a result of the interaction, an observable output, which is ‘caused’ by the measurand. From the observable output it is possible to go back to the cause and to properly assign a value to the measurand. We then propose to define the measuring system as an *empirical system* able to *interact* with any object carrying the quantity under investigation and to produce, as a result of the interaction, an *observable output*, on the basis of which it

is possible to assign a value to the measurand. Note that this definition is general enough to accomplish also measurements by a panel or a jury with respect to a previously defined reference scale, as often occurs in the case of perception [37]. Provided that this definition is accepted, how can we formally state this? In particular how can we characterise the interaction of the MS with the measurand object?

The solution is straightforward, because the property that we need is simply the following. *The output of the measuring system should depend only on the state of the quantity* (and thus should not depend on the specific object manifesting that state). The behaviour of the MS may thus be described and characterised by a function $\varphi : A \mapsto \mathbb{R}$ such that, for each $a, b \in A$,

$$a \sim b \Leftrightarrow \varphi(a) = \varphi(b). \quad (64)$$

The output of the MS does not depend on the specific object but only on the value of the measurand, thus another useful description is provided by the *calibration function* f , that we have informally introduced in Section 1.4. We may now provide a formal definition: let x be the value of the measurand and $f : X \mapsto \mathbb{R}$ the calibration function; then

$$\varphi(a) = f[m(a)]. \quad (65)$$

The deterministic description of the measurement process is then still provided by the formulae (44–46). All this holds in a deterministic framework and thus describes the ideal measurement. Later on we present the corresponding probabilistic model, which is also in agreement with what we have anticipated in Section 1.4.

2.3 A probabilistic theory of measurement

Probabilistic relations

We now briefly sketch a probabilistic theory of measurement that we have recently proposed and published [38].

The key point for attaining such a theory is the introduction of the notion *probability of a relation* [28, 32]. It is important to note first that the term ‘relation’ may be understood both in a general and in a specific meaning. For example, when we write $a \succsim b$, we mean that the relation \succsim holds for the couple (a, b) in A (this is the specific meaning). On the other hand, when we speak of the relation \succsim on A , we refer to the set of all the pairs of elements of A which satisfy it (this is the general meaning). Consider now this second, general, meaning. Then, a *probabilistic relation* of some kind, for example, a weak order, on a finite set A , may be defined by *assigning a probability distribution over the class of all possible relations of that kind on A* . This is illustrated by the simple example in Table 3, concerning a set with only three elements, $A = \{a, b, c\}$.

Table 3. An illustrative example of a probabilistic order structure on $A = \{a, b, c\}$.

Relational System \mathcal{A}_i	Weak Order Relations \succsim_i	x_a	x_b	x_c	$\mathbb{P}(\mathcal{A}_i)$ (Example)
\mathcal{A}_1	$a \succ b \succ c$	3	2	1	0.2
\mathcal{A}_2	$a \succ c \succ b$	3	1	2	0.2
\mathcal{A}_3	$b \succ a \succ c$	2	3	1	0.0
\mathcal{A}_4	$b \succ c \succ a$	1	3	2	0.0
\mathcal{A}_5	$c \succ a \succ b$	2	1	3	0.0
\mathcal{A}_6	$c \succ b \succ a$	1	2	3	0.0
\mathcal{A}_7	$a \sim b \succ c$	2	2	1	0.1
\mathcal{A}_8	$a \sim c \succ b$	2	1	2	0.1
\mathcal{A}_9	$b \sim c \succ a$	1	2	2	0.0
\mathcal{A}_{10}	$a \succ b \sim c$	2	1	1	0.3
\mathcal{A}_{11}	$b \succ a \sim c$	1	2	1	0.0
\mathcal{A}_{12}	$c \succ a \sim b$	1	1	2	0.0
\mathcal{A}_{13}	$a \sim b \sim c$	1	1	1	0.1

In the table all the possible weak orders on A are listed (column 2) and a probability is assigned to each of them (column 6). It may be that the probability of some of them is null, but it is necessary that no relation which is *not* a weak order has a nonnull probability. Note that when we assign a probability to a weak order, say \succsim_i , we also formally assign it to the *order structure* $\mathcal{A}_i = (A, \succsim_i)$.

If we now consider, in this example, a specific relation holding for a specific pair of elements, for example, $a \succ b$, we may note that it is verified in $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_5, \mathcal{A}_8$, and \mathcal{A}_{10} and consequently its probability is

$$\mathbb{P}(a \succ b) = \mathbb{P}(\mathcal{A}_1) + \mathbb{P}(\mathcal{A}_2) + \mathbb{P}(\mathcal{A}_5) + \mathbb{P}(\mathcal{A}_8) + \mathbb{P}(\mathcal{A}_{10}) = 0.8.$$

In general, for each couple of elements $a, b \in A$, we may calculate the probability of the empirical relations $a \succsim b$ as

$$\mathbb{P}(a \succsim b) = \sum_{a \succsim b \in \mathcal{A}_i} \mathbb{P}(\mathcal{A}_i). \tag{66}$$

What we have so far presented is a *probabilistic order structure*. In a similar way it is possible to define a probabilistic counterpart also for interval and extensive structures.

The measurement scale

Thanks to the notion of probabilistic relation it is possible to propose a *probabilistic counterpart of the representation theorem*. Consider a finite set of objects A and either

1. A probabilistic order structure
2. A probabilistic interval structure
3. A probabilistic extensive structure

Then it is possible to assign a discrete random variable x_a to each element $a \in A$ in such a way that (for each $a, b, c, d \in A$):

1. For a probabilistic order structure:

$$\mathbb{P}(a \succ b) = P(x_a \geq x_b); \quad (67)$$

2. For a probabilistic difference structure:

$$\mathbb{P}(\Delta_{ab} \succ \Delta_{cd}) = P(x_a - x_b \geq x_c - x_d); \quad (68)$$

3. For a probabilistic extensive structure:

$$\mathbb{P}(a \sim b \circ c) = P(x_a = x_b + x_c). \quad (69)$$

Proof of this probabilistic representation theorem is provided in Reference [38]. Let us now illustrate formula (67) with the example in Table 3. In the table, for each \mathcal{A}_i (column 1), a proper assignment of values to the random variables is presented (columns 3–5). For example, when \mathcal{A}_1 holds, $x_a = 3$, $x_b = 2$, and $x_c = 1$. So it is possible to calculate the probability distribution for each random variable: for example, because $x_a = 3$ in \mathcal{A}_1 and \mathcal{A}_2 ,

$$P(x_a = 3) = \mathbb{P}(\mathcal{A}_1) + \mathbb{P}(\mathcal{A}_2) = 0.4.$$

It is now easy to check that the \mathcal{A}_i s for which, say, $a \succ b$ holds, namely $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_5, \mathcal{A}_8$, and \mathcal{A}_{10} , are the same for which also $x_a > x_b$ holds and consequently

$$\mathbb{P}(a \succ b) = P(x_a > x_b) = 0.8.$$

In the accompanying DVD this example is studied in more detail and the related software is addressed.

The measurement process

In the deterministic description of the MS (Section 2.2) we have assumed that the output of the MS does not depend upon the specific object selected, but only on its state, and from that we have deduced that the observation transformation can be described by a function $y = f(x)$, defining a unique relation between each value of the measurand x and the corresponding indication y of the MS. Here we maintain the hypothesis that the output of the MS does not depend upon the specific object selected, but only on its state, that is, on the specific value that it manifests when we make the measurement, but we assume that, for each such value, a plurality of indications is possible, governed by a probabilistic distribution. Consequently, as we have seen in Section 1.4, a probabilistic description of the MS may be obtained by considering a

conditional probability distribution which describes the observation phase, that is,

$$P(y|x). \tag{70}$$

Then the restitution phase follows, described by

$$P(x|y) = P(y|x) \left[\sum_{x \in X} P(y|x) \right]^{-1}. \tag{71}$$

Lastly, the overall measurement process is characterised by

$$P(\hat{x}|x) = \sum_{y \in Y} \delta[\hat{x} - \mu(x|y)]P(y|x), \tag{72}$$

where μ is a position parameter appropriate for the scale that we are considering (i.e., μ is the median, if the scale is ordinal, or the expected value, if the scale is interval or ratio) and δ is the unitary discrete impulse. Note that the integrals appearing in Table 1 are now replaced by sums, because now we are dealing with discrete random variables. This hypothesis, anyway, is not critical, because it is essentially equivalent to requiring that the measuring system has a finite resolution (see [29, 38] for a discussion of this point). We omit, for the sake of brevity, the treatment of influence quantities that may be explicated as presented in Section 1.4. Lastly, the *calibration* of the MS may be now intended as the operation aiming at obtaining the conditional distribution $P(y|x)$. A summary of the main results of this probabilistic theory of measurement is provided in Table 4.

Table 4. Synopsis of the proposed theory: deterministic versus probabilistic approach.

The Measurement Scale		
Scale Type	Deterministic Approach	Probabilistic Approach
Order	$a \succ b \Leftrightarrow m(a) \geq m(b)$	$\mathbb{P}(a \succ b) = P(x_a \geq x_b)$
Interval	$\Delta_{ab} \succ \Delta_{cd} \Leftrightarrow$ $m(a) - m(b) \geq m(c) - m(d)$	$\mathbb{P}(\Delta_{ab} \succ \Delta_{cd}) =$ $P(x_a - x_b \geq x_c - x_d)$
Ratio	$a \sim b \circ c \Leftrightarrow m(a) = m(b) + m(c)$	$\mathbb{P}(a \sim b \circ c) = P(x_a = x_b + x_c)$
The Measuring Process		
Process	Deterministic Approach	Probabilistic Approach
Observation	$y = f(x)$	$P(y x)$
Restitution	$\hat{x} = f^{-1}(y)$	$P(x y) = P(y x) \left[\sum_{x \in X} P(y x) \right]^{-1};$ $\hat{x} = \mu(x y)$
Measurement	$\hat{x} = f^{-1}(f(x)) = x$	$P(\hat{x} x) = \sum_{y \in Y} \delta[\hat{x} - \mu(x y)]P(y x)$

3 Final remarks

Probability and metrology are two closely linked disciplines.

The science of measurement has taken advantage of the development of probability and statistics and has assumed methods and tools. But at the same time it has also greatly contributed to the development of these disciplines: the early theory of errors is an outstanding example of this.

We think that both in the present and in the future measurement problems may be best faced not simply by looking for existing statistical methods to adopt, but rather by considering probability as the natural tool for dealing with matters in which determinism is not appropriate for providing a satisfactory description and explanation of facts. In this way the relationship within the two disciplines may be rich and fruitful and the dialogue between the related scientific communities intense and enriching for both parts.

Moreover we have seen how measurement science has developed according to two distinct mainstreams, because of the division which arose between scientists in physics and engineering on the one hand and in psychology and behavioural science, in the first half of the last century. Such a division has been, in our opinion, detrimental in many respects, because the two approaches, one based on the study of the measurement process, the other on the problem of measurability and of the construction of the measurement scale, naturally complement each other and are both necessary to attain a satisfactory overall theory of measurement. We have also seen how a new and unconventional way of using probability, the probability of relations, has been recently proposed and is extremely promising as it paves the way to a better foundation for measurement. So we believe in the possibility of a new foundation for a unique science of measurement, valuable for all domains of knowledge in which measurement is seen as a necessary tool for reinforcing knowledge and for gathering information, and we consider probability as the natural logic for such a science.

Appendix: Symbols and Notation

Some of the main symbols are listed in Table 5. As a general criterion, we have tried to keep notation as lean as possible. Due to the broadness of the subject, some symbols are polysemantic: we have preferred to establish an easy connection between similar ideas, rather than resorting to a wide mass of difficult-to-relate symbols. We adopt the usual convention of denoting vectors and matrices by bold characters. We do not use special notation (such as capitals or bold) for random variables. So the same symbol may be used for denoting a random variable as well as its specific value. For example, the probability distribution of v may be denoted either as $p_v(\cdot)$ or, in a shorthand notation, as $p(v)$.

Table 5. List of the main symbols of Sections 1 and 2, respectively.

x, \mathbf{x}	parameter(s) to be estimated, measurand (either scalar or vector)
y, \mathbf{y}	observation(s), instrument indication(s)
f, \mathbf{f}	scalar function, vector function
v, \mathbf{v}	random errors affecting the indications of a measuring system
max	operator that calculates the maximum of a function
\mathbf{A}	matrix
E, Var	expectation operator, variance operator
σ, σ^2	standard deviation, variance
N, n, m	integers
$\theta, \boldsymbol{\theta}$	influence parameter(s) producing systematic effects
p	probability density function, also called probability distribution
P	probability function, discrete probability distribution
$p_{t,\nu}$	t -Student probability density function, with ν degrees of freedom
\hat{x}	the “hat” symbol indicates an estimator or an estimated value; if applied to the measurand it denotes the measurement value
\bar{y}, \tilde{y}	arithmetic mean of y , mean value of y
u, U	standard uncertainty, expanded uncertainty
g	function appearing in the GUM evaluation formula
z, \mathbf{z}	corrections of influence quantities
A	set of objects manifesting the characteristic x
m	measure function, $m : A \mapsto \mathbb{R}$
\succsim	empirical weak-order relation on A , empirical weak order relation between intervals
\succsim_i	i th empirical weak-order relation definable on A
Δ_{ab}	interval between elements a and b of A
\circ	binary empirical operation of concatenation (i.e., empirical sum) of elements of A
\sim	empirical equivalence relation on A defined, for $a, b \in A$ by $a \sim b \Leftrightarrow (a \succsim b) \text{ and } (b \succsim a)$
$\mathcal{A} = (A, \succsim)$	empirical order system on A
$\mathcal{A}_i = (A, \succsim_i)$	i th empirical order system on A
\mathbb{P}	probability function whose argument is a relation or a relational system (such as an order system)
X	set of the possible values of the measurand, image of A in \mathbb{R} , through the measure function m ; that is, $X = m(A)$
Y	set of the output values (indications) of the measuring system
μ	position parameter of a probability distribution (e.g., expected value or median)
φ, f	characteristic functions of the measuring system: $\varphi : A \mapsto Y, f : X \mapsto Y$; <ul style="list-style-type: none"> • the property characterising φ is: $\forall a, b \in A, a \sim b \Leftrightarrow \varphi(a) = \varphi(b)$ • the function f is defined by: $\forall a \in A, y = \varphi(a) = f[m(a)]$
g	measurement function, $g : X \mapsto X$, defined, for $a \in A, x = m(a)$, by: $\hat{x} = g(x) = f^{-1}[f(x)] = x$, and $m(a) = g[m(a)] = g(x) = x$

References

1. Bayes, T.: An essay towards solving a problem in the doctrine of chances. *Phil. Trans. R. Soc. London*, **53**, 370-418 (1763)
2. Laplace, P. S.: Mémoire sur la probabilité des causes par les événemens. *Mem. Acad. R. Sci.*, **6**, 621-656 (1774)
3. Gauss, C. F.: *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*. Hamburg (1809) English translation by Davis, C. H., Dover (1963) reprinted 2004
4. Laplace, P.S.: *Theorie analytique des probabilités*, Courcier, Paris, (1812) In: *Oeuvres Complètes de Laplace*. Gauthier-Villars, Paris, vol. VII
5. Gauss, C. F.: *Theoria combinationis observationum erroribus minimis obnoxiae*. Göttingen, (1823) English translation by Stewart, G.W., SIAM, Philadelphia (1995)
6. von Helmholtz, H.: *Zählen und Messen Erkenntnis – theoretisch betrachtet, Philosophische Aufsätze Eduard Zeller gewidmet*. Fuess, Leipzig (1887)
7. Campbell, N. R.: *Physics - The elements*. (1920) *Representations as Foundations of science*. Dover, New York (1957)
8. Fisher, R. A.: *Statistical methods for research workers*. Oliver and Boyd, Edinburgh (1925)
9. Ferguson, A., Myers, C. S., Bartlett, R. J.: *Qualitative estimates of sensory events*. Final Report – British Association for the Advancement of Science, **2**, 331-349 (1940)
10. Stevens, S. S.: On the theory of scales and measurement. *Science*, **103**, 667-680 (1946)
11. Fisher, R. A.: *Statistical methods and scientific inference*. Oliver and Boyd, Edinburgh (1956)
12. Mandel, J.: *The statistical analysis of experimental data*. Wiley (1964) Repr. Dover (1984)
13. Krantz, D.H., Luce, R.D., Suppes, P., Tversky, A.: *Foundations of Measurement*. Vol 1–3, Academic Press, New York (1971–1990)
14. Sheynin, O. B.: Laplace's theory of errors. *Archive for history of exact sciences*, **17**, 1–61 (1977)
15. Sheynin, O. B.: C. F. Gauss and the theory of errors. *Archive for history of exact sciences*, **20**, 21–72 (1979)
16. Roberts, F.S.: *Measurement theory*. Addison-Wesley, Reading, MA (1979)
17. Falmagne, J.C.: A probabilistic theory of extensive measurement. *Philosophy of Science*, **47**, 277–296 (1980)
18. Leaning M.S., Finkelstein, L.: A probabilistic treatment of measurement uncertainty in the formal theory of measurement. In: Streker, G. (ed) *ACTA IMEKO* 1979. Elsevier, Amsterdam (1980)
19. Finkelstein, L., Leaning, M.S.: A review of the fundamental concepts of measurement. *Measurement*, **2**, 25–34 (1984)
20. Narens, L.: *Abstract measurement theory*. MIT Press, Cambridge (1985)
21. Gonella, L.: *Measuring instruments and theory of measurement*. In: Proc. XI IMEKO World Congress, Houston, (1988)
22. Press, S.J.: *Bayesian statistics*. Wiley, New York (1989)
23. Costantini, D., Garibaldi, U., Penco, M. A.: *Introduzione alla statistica- I fondamenti dell'argomentazione incerta*. Muzzio, Padova (1992)

24. Weise, K., Wöger, W.: A Bayesian theory of measurement uncertainty. *Measurement Science and Technology*, **4**, 1–11 (1993)
25. BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML: International Vocabulary of Basic and general terms in Metrology. Second Edition (1994)
26. BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML: Guide to the Expression of Uncertainty in Measurement. ISO, Geneva (1995)
27. Michellini, R.C., Rossi, G.B.: Measurement uncertainty: a probabilistic theory for intensive entities. *Measurement*, **15**, 143–157 (1995)
28. Regenwetter, M.: Random utility representations of finite m-ary relations. *Journal of Mathematical Psychology*, **40**, 219–234 (1996)
29. Mari, L.: Beyond the representational viewpoint: a new formalization of measurement. *Measurement*, **27**, 71–84 (2000)
30. Monti, C. M., Pierobon, G.: *Teoria della probabilità*. Zanichelli, Bologna (2000)
31. Hacking, I.: An introduction to probability and inductive logic. Cambridge Press, Cambridge (2001) Italian edition: *Il Saggiatore*, Milano (2005)
32. Regenwetter, M., Marley, A.A.J.: Random relations, random utilities, and random functions. *Journal of Mathematical Psychology*, **45**, 864–912 (2001)
33. Luce, R.D., Suppes, P.: Representational measurement theory. In: Stevens' *Handbook of Experimental Psychophysics*. Vol. 4, Wiley, (2002)
34. Lira, I.: *Evaluating the measurement uncertainty*. IOP, Bristol (2002)
35. Rossi, G.B.: A probabilistic model for measurement processes. *Measurement*, **34**, 85–99 (2003)
36. Costantini, D.: *I fondamenti storico-filosofici delle discipline statistico probabilistiche*. Bollati Boringhieri, Torino (2004)
37. Rossi, G.B., Crenna, F., Panero, M.: Panel or jury testing methods in a metrological perspective. *Metrologia*, **42**, 97–109 (2005)
38. Rossi, G.B.: A probabilistic theory of measurement, *Measurement*, **39**, 34–50 (2006)
39. Rossi G.B., Crenna F., Cox M.G., Harris P.M.: Combining direct calculation and the Monte Carlo Method for the probabilistic expression of measurement results. In: Ciarlini, P., Filipe, E., Forbes, A.B., Pavese, F., Richter, D. (eds) *Advanced Mathematical and Computational Tools in Metrology VII*, World Scientific, Singapore (2006)
40. Rossi, G.B., Crenna F.: A probabilistic approach to measurement-based decisions. *Measurement*, **39**, 101–119 (2006)
41. Rossi, G.B.: Measurability. *Measurement*, **40**, 545–562 (2007)
42. Cox, M.G., Rossi, G.B., Harris, P.M., Forbes, A.: A probabilistic approach to the analysis of measurement processes, *Metrologia*, **45**, 493–502 (2008)

Three Statistical Paradigms for the Assessment and Interpretation of Measurement Uncertainty

William F. Guthrie,¹ Hung-kung Liu,² Andrew L. Rukhin,³ Blaza Toman,⁴
Jack C. M. Wang,⁵ Nien fan Zhang⁶

¹ NIST, 100 Bureau Dr Stop 8980, Gaithersburg, MD 20899 guthrie@nist.gov

² NIST, 100 Bureau Dr Stop 8980, Gaithersburg, MD 20899 liu@nist.gov

³ NIST, 100 Bureau Dr Stop 8980, Gaithersburg, MD 20899 and Dept. of Math &
Stat, UMBC, Baltimore, MD 21250 rukhin@nist.gov

⁴ NIST, 100 Bureau Dr Stop 8980, Gaithersburg, MD 20899 toman@nist.gov

⁵ NIST, 325 Broadway, Boulder, CO 80305 jwang@boulder.nist.gov

⁶ NIST, 100 Bureau Dr Stop 8980, Gaithersburg, MD 20899 zhang@nist.gov

Summary. Since its adoption, the ISO *Guide* has sparked a revolution in uncertainty analysis. Of course, even with all of the positive contributions from the development and adoption of the ISO *Guide*, there will always be a need to improve the assessment of uncertainty in particular applications and to extend it to cover new areas. Among other work along these lines, the International Committee on Weights and Measures, Joint Committee on Guides in Metrology is currently working on several supplements to the ISO *Guide*. Other authors have also recently made many important contributions to the theory and practice of uncertainty analysis. The goals of this chapter are to discuss different approaches to uncertainty assessment from a statistical point of view and to relate them to the methods that are currently being used in metrology or are being developed within the metrology community. The particular statistical paradigms under which different methods for uncertainty assessment are described include the frequentist, Bayesian, and fiducial paradigms. Each approach is illustrated using common examples and computer code for carrying out each analysis is illustrated using open-source software.

Key words: Uncertainty analysis, frequentist, Bayesian, fiducial, metrology, statistics

1 Introduction

The adoption of the ISO *Guide to the Expression of Uncertainty in Measurement* [19] by a wide array of scientific organizations in 1992 has led to an increasing recognition of the need to include uncertainty statements in measurement results. The trend toward laboratory accreditation based on

standards such as ISO 17025 has greatly accelerated this process in the last few years. Recognizing that uncertainty statements are required for effective decision making, metrologists in laboratories of all types, from national metrology institutes to commercial calibration laboratories, are currently putting considerable effort into the development of appropriate uncertainty analyses for different types of measurements using the methods outlined in the ISO *Guide* (also commonly referred to as the “GUM”).

Some of the strengths of the procedures outlined and popularized in the ISO *Guide* are its standardized approach to uncertainty analysis, its accommodation of sources of uncertainty that are evaluated either statistically (Type A) or nonstatistically (Type B), and its emphasis on reporting all sources of uncertainty that have been considered. The main approach to the propagation of uncertainty advocated by the ISO *Guide*, linear approximation of the formula used to obtain a measurement result, is simple to carry out and in many practical situations gives results that are surprisingly similar to those obtained using more formally justified statistical methods. In short, since its adoption, the ISO *Guide* has sparked a revolution in uncertainty analysis.

Of course, even with the positive contributions from the development and adoption of the ISO *Guide*, there will always be more work needed to improve the assessment of uncertainty in particular applications and to extend it to cover additional areas. Among other work along these lines the International Committee on Weights and Measures’ Joint Committee on Guides in Metrology, whose members authored the ISO *Guide*, is currently working on several supplements to the ISO *Guide* on special topics such as propagation of distributions using a Monte Carlo method, models with any number of output quantities, and modeling [1].

The utility of the ISO *Guide* is also widened by recent important work on the assessment of uncertainty in situations involving autocorrelated measurements. As indicated in Section 4.2.3 of the ISO *Guide*, use of the sample standard deviation to estimate the uncertainty in a set of replicate measurements is only appropriate when the measurements are independent. Section 4.2.7 notes “If the random variations in the observations of an input quantity are correlated, for example, in time, the mean and experimental standard deviation of the mean as given in 4.2.1 and 4.2.3 may be inappropriate estimators (C.2.25) of the desired statistics (C.2.23).” Use of the sample standard deviation for positively autocorrelated data from a stationary process will typically result in an underestimate of the standard uncertainty. Zhang [33] proposed an approach involving the autocorrelation function for calculating the standard uncertainty of the mean of stationary autocorrelated measurements. For some nonstationary processes, in particular those that exhibit long-term memory or long-range self-dependence, the Allan variance, which is mentioned in Section 4.2.7 of the ISO *Guide*, or the wavelet variance are useful tools for uncertainty analysis [26]. Measurements from other nonstationary processes, however, may require other methods for uncertainty assessment.

Finally, because it must apply to the widest possible range of measurement processes, the definition of measurement uncertainty as “a parameter characterizing the dispersion of the quantity values being attributed to a measurand” cannot reasonably be given at more than a relatively conceptual level in the *International Vocabulary of Metrology* [21]. As a result, defining and understanding the appropriate roles of different statistical quantities in uncertainty assessment, even for relatively well-understood measurement applications, is a topic of particular interest to both statisticians and metrologists.

Many authors have approached these topics from a metrological point of view in earlier investigations. Some authors have focused on characterizing the statistical properties of the procedures that are given in the *ISO Guide*. Gleser [14] indicates that these procedures, as originally presented, are not strictly consistent with any Bayesian or frequentist interpretation in general. Kacker and Jones [22] propose some minor modifications to the *ISO Guide* procedures that bring the results into closer agreement with a Bayesian interpretation in some situations. Another recent short communication in *Metrologia* by Elster, Wöger, and Cox [10] discusses the relationship between procedures for uncertainty analysis proposed in a draft supplement to the *ISO Guide* and the results of a Bayesian analysis for a particular class of models. Willink [32] also discusses different possible probabilistic interpretations of uncertainty intervals and recommends approximating the posterior distributions for this class of Bayesian analyses with probability distributions from the Pearson family of distributions.

Lira and Wöger [25] do compare frequentist (referred to as “conventional”) and Bayesian approaches to uncertainty analysis. However, they limit their comparison of Bayesian and frequentist approaches to measurement systems for which all sources of uncertainty can be evaluated using Type A methods. In contrast, measurement systems with sources of uncertainty evaluated using both Type A and Type B methods are treated in this chapter and are illustrated using several examples, including one of the examples from Annex H of the *ISO Guide*.

Due to their training, statisticians have historically placed an especially strong emphasis on using methods for uncertainty assessment that have a clearly delineated probabilistic justification or interpretation. Through their work over the years, often outside metrology, several different paradigms for statistical inference relevant to uncertainty assessment have been developed. The goals of this chapter are to present some of those approaches to uncertainty assessment from a statistical point of view and to relate them to the methods that are currently being used in metrology or are being developed within the metrology community. The particular statistical paradigms under which different methods for uncertainty assessment are described include the frequentist, Bayesian, and fiducial paradigms, which are discussed further after outlining the notational conventions that are needed to distinguish different types of quantities clearly.

1.1 Notation

In Section 4.1.1 of the ISO *Guide* it is noted that Latin letters are used to represent both physical quantities to be determined via measurement (i.e., measurands) as well as random variables that may take on different observed values of a physical quantity (potential measurement results). This overlapping use of the same symbols, whose different meanings are only indicated by context, can be difficult to interpret and sometimes leads to unnecessary ambiguities or misunderstandings. To mitigate this potential source of confusion, we revert to the more traditional notation often used in the statistical literature. In this notation Greek letters are used to represent parameters in a statistical model (or measurands in ISO *Guide* terminology), uppercase Latin letters to represent random variables that can take on different values of an observable quantity, and lowercase Latin letters to represent specific observed values of a quantity. Of course, because additional notation may be required to denote other physical, mathematical, or statistical concepts, there will still always be some possibility for ambiguity. In those cases the correct interpretation should be discernable from the context.

Using this traditional statistical notation, a measurement equation for a quantity θ is denoted

$$\theta = f(\mu_1, \mu_2, \dots, \mu_p), \quad (1)$$

where f is a known mathematical function of the quantities μ_i .

When discussing the values of standard uncertainties, we also distinguish between theoretical and estimated values of the standard uncertainties of different quantities. To do this we use notation such as σ_μ or σ_X to denote theoretical standard uncertainties and notation such as S_X and s_x to denote estimates of standard uncertainty before and after being observed, respectively.

1.2 Statistical paradigms

The first statistical paradigm we discuss, in which uncertainty can be probabilistically evaluated, is frequentist. It is based on the statistical theory that is probably most familiar to many readers. As a result, the frequentist approach is sometimes referred to as “classical” or “conventional” by some authors. However, due to the nature of uncertainty in metrology, these familiar methods must often be adapted to obtain frequentist uncertainty intervals under realistic conditions.

In the frequentist approach, the input parameters to the function f and its output θ all are modeled as unknown constants. Then data related to each input parameter μ_i are obtained and used to estimate the value of θ based on the measurement equation or equivalent statistical models. Finally exact or approximate, data-derived confidence limits for θ , with a nominal level of confidence specified by the analyst, are obtained using one of several mathematical principles or procedures, for example, least squares, maximum likelihood, and the bootstrap.

Because θ is traditionally treated as constant, the probabilistic statement associated with a confidence interval for θ is not a direct probability statement about its value. Instead, it is a probability statement about how frequently the procedure used to obtain the final uncertainty interval would encompass the value of θ with repeated use. Traditional frequentist uncertainty intervals provide a probability statement about the long-run properties of the estimation procedure used to construct the interval under the particular set of conditions assumed to apply to the measurement process.

In most practical metrological settings, on the other hand, uncertainty intervals must account for both the uncertainty in quantities estimated using data and the uncertainty in quantities whose values are based on expert knowledge. To obtain an uncertainty interval analogous to a confidence interval in this case, those measurands that are not observed must typically be treated as random variables with probability distributions for their values whereas measurands whose values can be estimated using statistical data are treated as unknown constants.

The various traditional frequentist procedures for the construction of confidence intervals must then be modified to attain the specified confidence level after averaging over the potential values of the quantities assessed using expert judgment [14]. Such modified uncertainty intervals provide long-run probability statements about the procedure used to obtain the interval given the probability model for the quantities that have not been observed, just as traditional confidence intervals do when all of the parameters are treated as constants.

The second paradigm is called the Bayesian approach. It is named after the fundamental theorem on which it is based, which was proved by the Reverend Thomas Bayes in the mid-1700s [3]. In this approach the analyst's knowledge about the measurands in Equation (1) is modeled as a set of random variables that follow a probability distribution on the joint parameter space of μ_1, \dots, μ_p and θ . Bayes' theorem then allows these probability distributions to be updated based on the observed data (also modeled using probability distributions) and the interrelationships of the parameters defined by the function f or equivalent statistical models. Then, one obtains a probability distribution describing one's knowledge of θ given the observed data. Uncertainty intervals that contain θ with any specified probability can then be obtained from this distribution. Because one's knowledge of the parameter values is described using probability distributions, Bayesian methods provide direct probabilistic statements about the value of θ , and the other parameters, using a definition of probability as a measure of belief.

The third, and last, of the statistical paradigms discussed in this section is called the fiducial approach. It was initially developed by R. A. Fisher in the 1930s [11]. In this approach a probability distribution, called the *fiducial distribution*, for θ conditional on the data is obtained based on the interrelationship of θ and the μ_i described by f and the distributional assumptions about the data used to estimate μ_i . Once obtained, the fiducial distribution

for θ can be used to obtain uncertainty intervals that contain θ with any specified probability.

The fiducial argument that justifies the process used to obtain the fiducial distribution is most easily understood using a simple example. Suppose the values taken on by a measurement Y can be described by the equation $Y = \mu + Z$, where μ is the measurand and Z is a standard normal random variable. If y is a realized value of Y corresponding to the realized value z of Z , then we have $\mu = y - z$. Of course Z is not observable. Nevertheless, the fact that we know the distribution from which z was generated helps us determine a set of values of μ that we consider plausible. We can use the probabilities associated with Z to infer the probabilities for μ . The process of transferring the relationship $\mu = y - z$ to the relationship $\mu = y - Z$ is what constitutes the fiducial argument. The fiducial distribution of μ is the distribution of the random variable $y - Z$ with y fixed.

When describing the different methods for uncertainty analysis under each of these statistical paradigms, we discuss their fundamental underlying assumptions, incorporation of uncertainties evaluated using Type A or Type B methods, and the probabilistic interpretation of the resulting uncertainty assessments. We also describe how the methods used in the ISO *Guide* relate to the frequentist, Bayesian, or fiducial results.

1.3 Examples

Each of the approaches is illustrated using two examples. In Example 1, the goal is to assess the value of a physical quantity θ based on measurements that must be corrected for a background interference β . Three different versions of Example 1 are presented to facilitate comparison of the different approaches. The notation used for each version is given in Table 1.

In the first version, denoted Example 1a, a random sample of five independent measurements of the signal plus background was observed. Each measurement is assumed to follow a Gaussian distribution with true mean $\gamma = \theta + \beta$ and true standard deviation σ_Y . The observed values of the signal plus background y are

$$3.738, 3.442, 2.994, 3.637, \text{ and } 3.874.$$

Table 1. Notation that is used for Example 1.

Quantity	Symbol
Physical quantity of interest (i.e., the measurand)	θ
Background interference	β
Quantity detected by measurement method when measuring samples (i.e., measurand plus background)	γ
Standard deviation of measurement method when measuring samples	σ_Y
Standard deviation of measurement method when measuring background	σ_B

These data have a sample mean of $\bar{y} = 3.537$ and a sample standard deviation of $s_y = 0.342$. Similarly, five independent measurements of the background are made. These measurements are also assumed to follow a Gaussian distribution, with true mean β and true standard deviation σ_B . The observed values of the background b are

1.410, 1.085, 1.306, 1.137, and 1.200.

Because there are statistical data for each quantity that is a source of uncertainty in Example 1a, this version of the example lends itself to straightforward statistical interpretation under each of the different paradigms.

The scenario for the second version of Example 1, Example 1b, is the same as for Example 1a with the exception that the assessment of the background is based on expert knowledge or past experience, rather than on fresh experimental data. In this case, the background β is believed to follow a uniform (or rectangular) distribution with endpoints 1.126 and 1.329. Because the background is being assessed using expert judgment in this scenario, its uncertainty will have to be evaluated using Type B methods. Thus this version is more like most real measurement scenarios than Example 1a.

The scenario for the final version of Example 1, Example 1c, is exactly the same as that for Example 1b except that the value of the signal θ is closer to the background. The data observed for the signal plus background in this case are

1.340, 1.078, 1.114, 1.256, and 1.192.

With the signal just above the background, this version illustrates how physical constraints can be incorporated in the assessment of uncertainty under each paradigm.

Example 2 is the calibration of the length of an end gauge taken from Annex H.1 of the ISO *Guide*. Because it is more complicated, it is introduced and discussed in Section 5, after the three methods for uncertainty assessments are discussed and illustrated using Example 1.

2 Frequentist approach to uncertainty assessment

2.1 Basic method

One of the main statistical techniques used in applications, in particular in metrology, consists of constructing confidence intervals for an unknown parameter θ . The $100(1 - \alpha)\%$ confidence interval (or uncertainty interval) is supposed to cover θ with probability $1 - \alpha$. Its midpoint gives a point estimate of the value of this parameter, and its half-width provides an idea about accuracy of the estimation. In metrology the half-width of the uncertainty interval is commonly termed the expanded uncertainty of the estimator.

In the frequentist context, the parameters are unknown constants. Following our convention to denote the random variables by uppercase letters and the observed values of random variables by lowercase letters, a confidence interval can be obtained from a *pivotal quantity* for θ (i.e., a function $W(\mathbf{Y}, \theta)$ of the (possibly multivariate) data \mathbf{Y} and the parameter θ), whose distribution under θ is parameter-free (provided such a function with a tractable distribution can be found). Then by determining the upper and lower percentiles ℓ_α and u_α of such a distribution, $P_\theta(\ell_\alpha \leq W(\mathbf{Y}, \theta) \leq u_\alpha) = 1 - \alpha$, one can solve the inequalities $\ell_\alpha \leq W(\mathbf{Y}, \theta) \leq u_\alpha$ in θ to obtain a $100(1 - \alpha)\%$ confidence interval for θ .

For example, let $\mathbf{Y} = (Y_1, \dots, Y_n)$ represent a random sample from $N(\mu, \sigma^2)$ with $\bar{Y} = \sum Y_i/n$. If the parameter of interest is μ , then for known σ ,

$$Z = \frac{\bar{Y} - \mu}{\sigma/\sqrt{n}} \sim N(0, 1),$$

is a pivotal quantity. The frequentist confidence interval for μ is given by

$$\bar{Y} \pm \frac{\sigma}{\sqrt{n}} z_{\alpha/2}. \quad (2)$$

If σ is not known, the sample standard deviation

$$S = \sqrt{\frac{\sum_{j=1}^n (Y_j - \bar{Y})^2}{n - 1}}$$

is used to estimate σ . Then the (exact) pivotal quantity for μ is obtained by replacing σ in (2) by S ,

$$\frac{\bar{Y} - \mu}{S/\sqrt{n}} \sim t(n - 1). \quad (3)$$

Thus, the Student- t distribution based confidence interval for μ is

$$\bar{Y} \pm \frac{S}{\sqrt{n}} t_{n-1, 1-\alpha/2},$$

where $t_{n-1, \beta}$ is the 100β percentile of the t -distribution with $n - 1$ degrees of freedom.

In lieu of *exact* pivotal quantities, which exist only in simple situations, approximate pivotal quantities are commonly employed in applications. For large samples the central limit theorem can be invoked to get approximate confidence intervals based on the normal distribution.

Further methods of obtaining confidence intervals (inverting a test statistic, pivoting a continuous cumulative distribution function, ordering the discrete sample values according to their probabilities, etc.) are discussed in [6]. Some of them are mentioned in Example 1. A computer-intensive method,

called the bootstrap, also can be used to construct a confidence interval for pivot quantities that have unknown distributions. The bootstrap procedure is discussed later in this section.

Although not explicitly given a frequentist justification from first principles, the procedures recommended in the ISO *Guide* sometimes can be used to obtain an approximate confidence interval for the measurand. Such confidence intervals are based on an approximate pivotal quantity with an assumed Student's- t distribution obtainable from the measurement equation (1), which is the centerpiece of the *Guide*. Under this procedure the estimate of the measurand θ is based on the *plug-in method*. More precisely, the unknown p quantities μ_1, \dots, μ_p are replaced by estimates x_1, \dots, x_p obtained from physical measurements or from other sources. Typically measurements x_i are sample means or other functions of data designed to estimate some of the quantities, say, $\mu_i, i = 1, \dots, m$. Their (Type A) uncertainty $u(x_i)$ is also estimated on the basis of the data by statistical methods, usually via the sample standard deviation or by other robust, rank-based procedures. Then the degrees of freedom $\nu_i = \nu_i(x_i)$ associated with $u(x_i)$ is determined from the sample size used to estimate μ_i .

As physical measurements are not always possible or feasible for some of the μ_j , subjective evaluations $x_j, j = m + 1, \dots, p$, are used along with associated uncertainty characteristics. Thus, nonstatistical types of information are used to estimate μ_{m+1}, \dots, μ_p using Type B methods including scientific judgment, manufacturer's specifications, or other indirectly related or incompletely specified information. There are situations where both Type A and Type B information is employed in estimating some of the μ_j . Commonly in this case these quantities are considered to have uncertainties whose evaluation is of Type B.

The ISO *Guide* (see also Section 4.5 in [23]) recommends that the same functional relationship that relates the value θ of the measurand to the input quantities μ_1, \dots, μ_p be used to calculate y from x_1, \dots, x_p . Thus, the measurement (or, in statistical nomenclature, the estimator) y of θ is obtained as

$$y = f(x_1, \dots, x_m, x_{m+1}, \dots, x_p);$$

that is, the evaluated magnitude of $y, y = f(x_1, \dots, x_p)$, is taken to be the measurement of θ .

The uncertainty $u(y)$ in the measurement result y must be evaluated. The suggested approach is to estimate $u(y)$ by the method of uncertainty propagation, that is, the uncertainties $u(x_1), \dots, u(x_p)$ in the quantities $\mathbf{x} = (x_1, \dots, x_p)$ are used in the Taylor series expansion of the function $f(x_1, \dots, x_p)$ at μ_1, \dots, μ_p whose first-order terms are

$$f(x_1, \dots, x_p) \approx f(\mu_1, \dots, \mu_p) + \sum_{i=1}^p c_i(x_i - \mu_i). \quad (4)$$

The partial derivatives here,

$$c_i = \left. \frac{\partial f}{\partial \mu_i} \right|_{\mu=\mathbf{x}},$$

are called *sensitivity coefficients*. Treating these partial derivatives as constants and applying the rule for the propagation of uncertainty gives the approximate combined standard uncertainty of y ,

$$u_c(y) \approx \sqrt{\sum_{i=1}^p c_i^2 u^2(x_i) + 2 \sum_{i<j} c_i c_j u(x_i, x_j)}, \quad (5)$$

where $u(x_i, x_j)$ is the estimated covariance between X_i and X_j . Validity of this formula is based on the implicit assumption of well-defined, nonvanishing sensitivity coefficients. In other words, the ISO *Guide* prescription works well when the underlying function f is *locally linear*. If some of the sensitivity coefficients vanish, a higher-order Taylor's expansion is needed, but then the formulas (5) through (9) do not hold.

To quantify the uncertainty of $u_c(y)$, the ISO *Guide* suggests computing the effective degrees of freedom using the Welch–Satterthwaite formula,

$$\nu_{\text{eff}} = \frac{u_c^4(y)}{\sum_{i=1}^p \frac{c_i^4 u^4(x_i)}{\nu(x_i)}}. \quad (6)$$

This formula is appropriate when all of the individual uncertainties in the denominator are estimated independently of one another and when all of the input variables are approximately normal and uncorrelated. This formula can give a reasonable answer only when the provided degrees of freedom is not too heterogeneous. Guthrie [15] discusses a counterintuitive property according to which in interlaboratory studies a confidence interval based on the Welch–Satterthwaite approximation may be shorter for a between-lab difference than for one of its components.

Finally, in order to construct a confidence interval for θ , the approximate pivotal quantity,

$$W(y, \theta) = \frac{y - \theta}{u_c(y)} \quad (7)$$

is employed. According to the ISO *Guide*,

$$W(Y, \theta) \sim t(\nu_{\text{eff}}); \quad (8)$$

that is, $W(Y, \theta)$ is an approximately pivotal quantity which has a t -distribution with ν_{eff} degrees of freedom. The $100(1 - \alpha)\%$ confidence interval for θ ,

$$y \pm u_c(y) t_{\nu_{\text{eff}}, 1-\alpha/2} \quad (9)$$

then can be recommended as the $1 - \alpha$ uncertainty interval for θ . The half-width of this interval, $t_{\nu_{\text{eff}}, 1-\alpha/2} u_c(y)$, is the expanded uncertainty of y .

This recommendation agrees with standard statistical practice when all uncertainties are evaluated using Type A methods, in which case, the most commonly used statistical estimate for the particular input μ is the sample mean of n observations. The traditional method for summarizing data to obtain the Type A standard uncertainty of this estimator is S/\sqrt{n} with $n - 1$ degrees of freedom. This fact is true for more general statistics of the form $Y = f(X_1, \dots, X_p)$, where the estimators $X_i, i = 1, \dots, p$, obey the central limit theorem for large sample sizes. Indeed in this situation, the standard deviation of Y can be approximated by (5) with $u(x_i, x_j)$ replaced by $\text{Cov}(X_i, X_j)$.

The ISO *Guide* method presents the collective wisdom of many metrologists, but it is not an exact mathematical theorem. It is restricted by assumptions of

- local linearity of the function f : ideally the sensitivity coefficients should not vary much and they do not vanish,
- normality of the distribution of point estimators: may not hold even approximately for small samples, and
- validity of the Welch–Satterthwaite formula (6): it may not work well in the heterogeneous case. In addition, the degrees of freedom for distributions unrelated to the chi-squared law are difficult to interpret, indeed, they are not being used in statistical theory.

To motivate (5) in the frequentist setting, one can employ the concepts of statistical decision theory and interpret the uncertainty $u^2(y)$ as the mean squared error of the statistical estimator of $f(x_1, x_2, \dots, x_p)$ provided that the quantities whose uncertainties are of Type B, x_{m+1}, \dots, x_p , are eliminated by integrating over their distributions [14]. Then (5) provides the first-order approximation of the mean squared error if f “is sufficiently close to being linear.”

The discussion in Example 1 gives another customary frequentist procedure in the context of confidence intervals.

Bootstrap uncertainty intervals

The key assumption used in constructing the ISO *Guide* interval is (8), which may not hold approximately even for simple problems. However, through the use of a statistical method called the bootstrap [8], we can obtain accurate confidence intervals without making assumptions such as (8). One way to get such intervals is the “bootstrap- t ” approach. In essence, this procedure generates an empirical distribution for the approximate pivot $W(Y, \theta)$ (to replace the Student- t distribution in (8)) that is appropriate for the dataset at hand. Of course, when (8) is in fact correct, the bootstrap- t distribution will reproduce the Student- t distribution. The empirical bootstrap- t distribution is then used to construct a confidence interval in exactly the same way that the Student- t distribution is used in constructing (9).

Here is an outline of the bootstrap- t method as applied to (1). Using the estimates x_i , the associated uncertainties $u(x_i)$, and the assumed distributions, generate samples x_i^* and $u(x_i^*)$, $i = 1, \dots, p$. In other words, we use x_i as the mean and $u(x_i)$ as the standard deviation of the assumed distribution of the data X_i to generate parametric bootstrap input estimate x_i^* and its uncertainty $u(x_i^*)$. In typical metrological problems, datasets are not big enough to ensure the validity of the nonparametric bootstrap approach.

Just as described above, the ISO *Guide* takes $(x_i, u(x_i))$, for $i = 1, \dots, p$, as its input to produce $y, u_c(y)$, and $W(y, \theta)$; we take $x_i^*, u(x_i^*)$, $i = 1, \dots, p$, as input, to compute $y^*, u_c(y^*)$, and

$$W^* = W(y^*, y) = \frac{y^* - y}{u_c(y^*)}. \quad (10)$$

To get a bootstrap distribution of $W(Y, \theta)$, generate B bootstrap samples $x_i^*(b), u(x_i^*(b))$, $i = 1, \dots, p$, and for each compute $W^*(b)$, $b = 1, \dots, B$. The 100 α th percentile of the bootstrap- t distribution of $W(Y, \theta)$ is then estimated by the value \hat{t}_α such that

$$|\{W^*(b) \leq \hat{t}_\alpha\}|/B = \alpha,$$

where $|A|$ is the number of elements of the set A . Finally, the 100(1 - α)% bootstrap- t confidence interval is

$$(y - \hat{t}_{1-\alpha/2} \cdot u_c(y), y - \hat{t}_{\alpha/2} \cdot u_c(y)). \quad (11)$$

The Student- t percentiles are symmetric about zero, and as a consequence, (9) is symmetric about y . In contrast, the bootstrap- t percentiles can be asymmetric about zero, leading to an asymmetric uncertainty interval about y . The details of this process in constructing a 95% uncertainty interval are shown in the following algorithm.

1. For $i = 1, \dots, p$, using x_i as the mean and $u(x_i)$ as the standard deviation of the assumed distribution of X_i , generate B bootstrap samples $(x_i^*(1), u(x_i^*(1))), \dots, (x_i^*(B), u(x_i^*(B)))$.
2. For each bootstrap sample $(x_i^*(b), u(x_i^*(b)))$, $i = 1, \dots, p$, compute $y^*(b)$, $u_c(y^*(b))$, and $W^*(b) = (y^*(b) - y)/u_c(y^*(b))$ following the ISO *Guide*.
3. Estimate the 100 α th percentile of the bootstrap- t distribution of $W(Y, \theta)$ by the value \hat{t}_α such that $|\{W^*(b) \leq \hat{t}_\alpha\}|/B = \alpha$.
4. The 95% bootstrap- t confidence interval is

$$(y - \hat{t}_{0.975} \cdot u_c(y), y - \hat{t}_{0.025} \cdot u_c(y)).$$

Bootstrap samples can also be used to replace $u_c(y)$ by estimating the standard deviation of Y , when the Taylor's approximation (4) is deemed to be inappropriate. To do so, for $i = 1, \dots, p$ and $b = 1, \dots, B$, only input estimates $x_i^*(b)$ are to be generated. For each bootstrap sample,

$y^*(b) = f(x_1^*(b), \dots, x_p^*(b))$ is evaluated. The bootstrap estimate of the standard uncertainty of y is the sample standard deviation of the B replicates,

$$u_c(y) = \sqrt{\sum_{b=1}^B [y^*(b) - y^*(\cdot)]^2 / (B - 1)},$$

where $y^*(\cdot) = \sum_{b=1}^B y^*(b) / B$.

Finally a nested bootstrap of $(B_1 \cdot B_2)$ bootstrap samples can be carried out to construct a bootstrap- t interval using the bootstrap standard deviation estimator. We generate B_1 bootstrap samples of input estimates and the corresponding y^* . For each bootstrap sample, $u_c(y^*)$ is computed by B_2 second-level bootstrap samples, and

$$\frac{y^* - y}{u_c(y^*)}$$

is evaluated. The collection of B_1 such ratios is then used to estimate percentiles of $W(Y, \theta)$, which leads to the construction of a bootstrap- t interval as in (11). The algorithm of this nested bootstrap in constructing a 95% uncertainty interval is as follows.

1. For $i = 1, \dots, p$, using x_i as the mean and $u(x_i)$ as the standard deviation of the assumed distribution of X_i , generate B_1 first-level bootstrap samples $x_i^*(1), \dots, x_i^*(B_1)$.
2. For each first-level bootstrap sample $x_i^*(b_1), i = 1, \dots, p$, compute $y^*(b_1) = f(x_1^*(b_1), \dots, x_p^*(b_1))$, and $W^*(b_1) = (y^*(b_1) - y) / u_c(y^*(b_1))$, where $u_c(y^*(b_1))$ is determined by a second-level bootstrap using the following algorithm.
 - (a) For $i = 1, \dots, p$, using $x_i^*(b_1)$ as the mean and $u(x_i)$ as the standard deviation of the assumed distribution, generate B_2 second-level bootstrap samples $x_i^*(1), \dots, x_i^*(B_2)$.
 - (b) For each second-level bootstrap sample, $y^*(b_2) = f(x_1^*(b_2), \dots, x_p^*(b_2))$ is evaluated.
 - (c) The bootstrap estimate of the standard uncertainty of $y^*(b_1)$ is the sample standard deviation of the B_2 replicates

$$u_c(y^*(b_1)) = \sqrt{\sum_{b_2=1}^{B_2} [y^*(b_2) - y^*(\cdot)]^2 / (B_2 - 1)},$$

where $y^*(\cdot) = \sum_{b_2=1}^{B_2} y^*(b_2) / B_2$.

3. Estimate the 100α th percentile of the bootstrap- t distribution of $W(Y, \theta)$ by the value \hat{t}_α such that $|\{W^*(b_1) \leq \hat{t}_\alpha\}| / B_1 = \alpha$, where $|A|$ is the number of elements of the set A .
4. The 95% “nested bootstrap- t ” confidence interval is

$$(y - \hat{t}_{0.975} \cdot u_c(y), y - \hat{t}_{0.025} \cdot u_c(y)).$$

Although it is a more general approach that does not rely on the quality of the Taylor's approximation, the nested bootstrap is computationally more involved and harder to implement. We have chosen the simpler bootstrap method to analyze all our examples.

2.2 Example 1

As an illustration, consider the statistical model given for Example 1 in the Introduction

$$Y_i = \theta + \beta + \epsilon_i, \quad i = 1, \dots, n, \quad (12)$$

where θ is the quantity of interest, β represents the background, and ϵ_i are independent $N(0, \sigma^2)$ errors. For a fixed value β , with γ denoting the mean of the data, the measurement equation for this model is $\theta = f(\beta, \gamma) = \gamma - \beta$.

If background β has a uniform distribution on the interval $(a - d, a + d)$, the interval for θ derived from the ISO *Guide* is

$$\bar{Y} - a \pm 2\sqrt{\frac{\sigma^2}{n} + \frac{d^2}{3}}.$$

Gleser [14] discusses the properties of such intervals and compares them to the interval

$$\bar{Y} - a \pm \left[2\sqrt{\frac{\sigma^2}{n} + d} \right], \quad (13)$$

which can be motivated as follows. Because the conditional distribution of \bar{Y} for given β is normal, $N(\theta + \beta, \sigma^2/n)$,

$$P\left(|\bar{Y} - \theta - \beta| \leq \frac{2\sigma}{\sqrt{n}}\right) \geq 0.95,$$

and

$$P(|a - \beta| \leq d) = 1.$$

It follows that the *Eisenhart interval* (13) recommended by C. Eisenhart [9] is conservative,

$$P\left(|\bar{Y} - a - \theta| \leq \frac{2\sigma}{\sqrt{n}} + d\right) \geq 0.95.$$

However, if $d > 12\sqrt{\sigma^2/n}$, the ISO *Guide* recommended interval contains (13), which demonstrates the difference between these two approaches.

The interval (13) is easily adjustable for a t -distributed ratio $\sqrt{n}(\bar{Y} - a - \beta)/S$, and for other distributions of the background (triangular, trapezoidal, etc.). Different frequentist methods for construction of confidence intervals are available in this situation. Indeed in the model (12), \bar{Y} subsumes all the information the data provide about θ (i.e., in statistical jargon, \bar{Y} is a sufficient statistic for θ) with probability density

$$\frac{\sqrt{n}}{2\sqrt{2\pi}\sigma d} \int_{a-d}^{a+d} e^{-0.5n(\bar{y}-\theta-\beta)^2/\sigma^2} d\beta.$$

The special form of this distribution allows for many other statistical techniques [6] to derive alternative confidence intervals (all centered at the maximum likelihood estimator, $\bar{Y} - a$, but of different lengths.)

Example 1a

The simple example introduced in Section 1.3 summarizes the measurements in the model (12), with $\bar{y} = 3.537$ and $u(\bar{y}) = 0.153$. The latter must substitute for σ/\sqrt{n} above and the factor 2 should be replaced by the percentile of t -distribution with 5.15 effective degrees of freedom. In Example 1a, the background β can be estimated from measurements following a normal distribution leading to $\bar{b} = 1.228$, and $u(\bar{b}) = 0.059$. Our estimate of θ is $\bar{y} - \bar{b} = 2.309$ with combined standard uncertainty $\sqrt{u^2(\bar{y}) + u^2(\bar{b})} = 0.164$. The ISO *Guide* interval is

$$2.309 \pm 2.548 \times 0.164 = 2.309 \pm 0.417 = (1.892, 2.727).$$

The $100(1 - \alpha)\%$ bootstrap- t confidence interval according to (11) is $(2.309 - \hat{t}_{1-\alpha/2} \cdot 0.164, 2.309 - \hat{t}_{\alpha/2} \cdot 0.164)$, where \hat{t}_β is the 100β th percentile of W^* of (10). An R program [27] for generating the $B = 10,000$ realizations of W^* is listed below.

```
B = 10000
y.star = rnorm(B, mean=3.537, sd=0.153)
u.y.star = 0.153 * sqrt(rchisq(B, df=4)/4)
b.star = rnorm(B, mean=1.228, sd=0.059)
u.b.star = 0.059 * sqrt(rchisq(B, df=4)/4)
w.star = ((y.star-b.star)-2.309)/sqrt(u.y.star^2+u.b.star^2)
```

The 95% bootstrap- t confidence interval based on the 0.025 and 0.975 quantiles of the simulated distribution is

```
2.309 - quantile(w.star, c(0.975, 0.025)) * 0.164
## 1.895754 2.728817
```

Example 1b

When there are no statistical data for the background β , instead, β is assumed to have a uniform distribution on the interval (1.126, 1.329), then the approximate interval derived from the ISO *Guide* is

$$3.537 - 1.228 \pm 2.533 \sqrt{\frac{0.342^2}{5} + \frac{0.102^2}{3}} = 2.310 \pm 0.415 = (1.895, 2.724).$$

The Eisenhart interval is wider,

$$3.537 - 1.228 \pm \left[2.776 \frac{0.342}{\sqrt{5}} + 0.102 \right] = 2.310 \pm 0.526 = (1.783, 2.836).$$

Similar to Example 1a, one can construct a bootstrap- t confidence interval for θ . For this example, the estimates and their uncertainties of γ , β , and θ are numerically the same as those in Example 1a, except that β is estimated by the Type B method. Therefore, the realizations of W^* are generated differently from Example 1a only in generating the bootstrap sample b^* and its uncertainty. The bootstrap sample b^* is now generated from the known uniform (1.126, 1.329) distribution with uncertainty 0.059. The R code for generating $B = 10,000$ realizations of W^* is as follows.

```
B = 10000
y.star = rnorm(B, mean=3.537, sd=0.153)
u.y.star = 0.153 * sqrt(rchisq(B, df=4)/4)
b.star = runif(B, min=1.126, max=1.329)
u.b.star = 0.059
w.star = ((y.star-b.star)-2.309)/sqrt(u.y.star^2+u.b.star^2)
```

The 95% bootstrap- t confidence interval based on the 0.025 and 0.975 quantiles of the simulated distribution is

```
2.309 - quantile(w.star, c(0.975, 0.025)) * 0.164
## 1.918643 2.699749
```

Example 1c

If $\bar{y} = 1.196$, $s_{\bar{y}} = 0.047$, both intervals have negative lower endpoints. If the mean θ is known to be positive, they are replaced by 0 leading to the ISO *Guide* recommended interval (0, 0.124) and to the Eisenhart interval (0, 0.202).

An R program for generating the $B = 10,000$ realizations of W^* to obtain the bootstrap interval is the same as Example 1b with $\bar{y} = 1.196$ and $u(\bar{y}) = 0.047$.

```
B = 10000
y.star = rnorm(B, mean=1.196, sd=0.047)
u.y.star = 0.047 * sqrt(rchisq(B, df=4)/4)
b.star = runif(B, min=1.126, max=1.329)
u.b.star = 0.059
w.star = ((y.star-b.star)+0.032)/sqrt(u.y.star^2+u.b.star^2)
```

The untruncated 95% bootstrap- t confidence interval is

```
-0.032 - quantile(w.star, c(0.975, 0.025)) * 0.075
## -0.1762648 0.1128422
```

As θ is known to be positive, the truncated 95% bootstrap- t confidence interval for θ is (0, 0.113).

3 Bayesian paradigm for uncertainty assessment

3.1 Basic method

In most metrological experiments, the measurand and the input variables of model (1) are physical quantities with fixed quantity values. Nevertheless, under the Bayesian paradigm, the corresponding parameters μ_i and θ are considered random variables in the sense that their probability distributions summarize our knowledge about these physical quantities.

The Bayesian framework uses a definition of probability that allows probability distributions to be defined without physical data, for example, using manufacturer specifications or other expert knowledge. In most measurement applications however, there will be physical measurements (data) that can be used for estimation of one or more of the input quantities. In such cases the corresponding probability density can be obtained via Bayes' theorem as follows. Let $p(\mu_i)$ be a probability density for μ_i as given before the physical data are obtained. This is called the prior density of μ_i . Let Y denote a random variable for which a realization y (data) exists. The probability density of Y , $p(y|\mu_i)$, is called the *statistical model*. Under the Bayesian framework, as μ_i is a random variable, the notation $|$ represents the fact that the probability density of Y is conditional on μ_i . For a particular realization of y , $p(y|\mu_i)$, viewed as a function of μ_i is called the *likelihood function*. Applying Bayes' theorem,

$$p(\mu_i|y) = \frac{p(y|\mu_i)p(\mu_i)}{\int p(y|\mu_i)p(\mu_i)d\mu_i} \quad (14)$$

is the posterior density of μ_i which summarizes our knowledge about it after the data y were observed.

When no prior knowledge of the μ_i exists, then a so-called noninformative prior distribution [4] is used. In cases when prior information does exist, it is represented by an informative probability distribution. This is one of the mechanisms, under the Bayesian paradigm, for including information obtained by Type B methods into the uncertainty analysis. The form of the likelihood function is usually selected based on knowledge of the process that generates the data.

It is clear that the form of the likelihood function and of the prior densities will determine the shape of the posterior density. It is therefore important to select them carefully and to perform sensitivity analysis of the results with respect to changes in these distributions. For the prior distributions this may mean using several different densities and comparing the results. Tests of appropriateness of the probability distribution corresponding to the likelihood function are called model fitting [18] and apply equally to Bayesian and frequentist models.

The definition of measurement uncertainty in metrology, given in the Introduction, can be interpreted in the context of Bayesian statistics as referring to the posterior probability distribution of the measurand θ ; that is, the *standard uncertainty* is the standard deviation of this probability distribution. To obtain it, it is necessary first to find the joint probability distribution of all of the μ_i , and then apply a change-of-variables formula [6] to derive the distribution of θ . Moments of this distribution can be obtained more simply as follows. For a function $h(\theta)$, we obtain the expected value $E(h(\theta)) = \int \cdots \int h(f(\mu_1, \dots, \mu_p)) p(\mu_1, \dots, \mu_p) d\mu_1 \cdots d\mu_p$. Note that the variance can be obtained as $\text{Var}(\theta) = E(\theta^2) - [E(\theta)]^2$. Often, the necessary integration is done using Monte Carlo methods [4].

In the case when the μ_i are independent random variables, their joint probability distribution is simply the product of the individual distributions. In many situations, however, the μ_i are not independent random variables. For example, consider the case when the probability distribution of Y is a function of μ_1 and μ_2 ; that is, $p(y|\mu_1, \mu_2)$ is the statistical model and $p(\mu_1, \mu_2) \neq p(\mu_1)p(\mu_2)$. Then the posterior density of (μ_1, μ_2) , must be obtained as

$$p(\mu_1, \mu_2|y) = \frac{p(y|\mu_1, \mu_2) p(\mu_1, \mu_2)}{\int p(y|\mu_1, \mu_2) p(\mu_1, \mu_2) d\mu_1 d\mu_2}.$$

A common situation that leads to such dependence is when the statistical model is a function of θ , as well as some of the μ_i . Both of the examples considered here fall into this category. This illustrates the point that under the Bayesian paradigm, whenever measurement data are available, the process of specifying the related probability distributions needs to be started by defining a statistical model. This will automatically lead to the likelihood functions needed for the application of Bayes' theorem and to the correct posterior densities. The process can be summarized as follows.

1. Identify all measurement data relevant to the physical quantities of interest (parameters).
2. Specify a statistical model (also called observational model) relating the data to the parameters; these could be the μ_i s or sometimes the measurand θ .
3. Specify prior distributions for all parameters involved.
4. Apply Bayes' theorem to obtain posterior distributions of the parameters.

5. Compute the posterior mean and posterior standard deviation of the measurand.
6. Perform sensitivity analysis of the results with respect to changes in the prior distributions.

In some situations, a Taylor series approximation may be used to avoid the numerical computations. Specifically, the Taylor series expansion of $f(\mu_1, \dots, \mu_p)$ about the expected values of the μ_i s can be used to state that $f(\mu_1, \dots, \mu_p)$ is approximately distributed as $N(f(E(\mu_1), \dots, E(\mu_p)), \omega^2)$ where

$$\omega = \sqrt{\sum_i c_i^2 \text{Var}(\mu_i) + 2 \sum_{i < j} c_i c_j \text{Cov}(\mu_i, \mu_j)},$$

$\text{Cov}(\mu_i, \mu_j)$ denote the covariance of the μ_i s, and the c_i are the partial derivatives of θ with respect to μ_i evaluated at the expected values of the μ_i s. (Similarly appearing formulas (4) and (5) are used in Section 2.1, but there the expansion is employed to find an estimate of the variance of the estimator of θ , not of θ itself.)

3.2 Example 1

This process is now illustrated on Example 1 given in the Introduction. The measurand in this example is denoted by θ . The measurement equation can be written as

$$\theta = \gamma - \beta. \quad (15)$$

Example 1a

Consider Example 1a first. There are two relevant sets of data, the five Gaussian measurements Y_i with mean $\gamma = \theta + \beta$, and standard deviation σ_Y , and the five Gaussian measurements B_i with mean β and standard deviation σ_B . Thus, the statistical model for Y_i is

$$Y_i | \theta, \beta, \sigma_Y^2 \sim N(\theta + \beta, \sigma_Y^2),$$

and because the five observations are independent,

$$p(y_1, \dots, y_5 | \theta, \beta, \sigma_Y) = \left(\frac{1}{\sigma_Y \sqrt{2\pi}} \right)^5 \exp \left\{ - \frac{\sum_{i=1}^5 (y_i - \theta - \beta)^2}{2\sigma_Y^2} \right\}.$$

For B_i , the statistical model is

$$B_i | \beta, \sigma_B^2 \sim N(\beta, \sigma_B^2);$$

that is,

$$p(b_1, \dots, b_5 | \beta, \sigma_B) = \left(\frac{1}{\sigma_B \sqrt{2\pi}} \right)^5 \exp \left\{ - \frac{\sum_{i=1}^5 (b_i - \beta)^2}{2\sigma_B^2} \right\}.$$

Because the two sets of observations are assumed to be mutually independent, the statistical model is

$$p(y, b | \theta, \beta, \sigma_Y, \sigma_B) = p(b_1, \dots, b_5 | \beta, \sigma_B) p(y_1, \dots, y_5 | \theta, \beta, \sigma_Y).$$

There are four parameters, θ , β , σ_Y , and σ_B that must be assigned prior distributions. In this example, there is no additional information about these parameters, other than that they are not negative, and thus the random variables will be assumed to be independent. It is desirable for the forms of the prior distributions to have minimal effect on the analysis. Such a result is obtained by the use of so-called *reference* priors [4]. For the parameters associated with the means, that is, θ and β , such a density can be approximated by

$$\begin{aligned} \theta &\sim \text{Uniform}(0, c), \\ \beta &\sim \text{Uniform}(0, c), \end{aligned}$$

with a large value for c . For the scale parameters, σ_Y and σ_B , the reference prior density

$$\begin{aligned} p(\sigma_Y) &= 1/\sigma_Y, \\ p(\sigma_B) &= 1/\sigma_B, \end{aligned}$$

is improper; that is, it does not integrate to 1. This can cause difficulties in numerical computation so a proper density such as

$$\sigma_Y \sim \text{Uniform}(0, c),$$

or

$$\sigma_Y \sim \text{Gamma}(c, c),$$

with large values of c is usually substituted. The notation $\text{Gamma}(\phi_1, \phi_2)$ represents a gamma distribution with parameters ϕ_1 and ϕ_2 ; that is, for a random variable X , this probability density is given by

$$p(x | \phi_1, \phi_2) = \frac{\phi_2^{\phi_1}}{\Gamma(\phi_1)} x^{\phi_1-1} e^{-x\phi_2}.$$

This completes the prior distribution specification.

Application of Bayes' theorem results in the joint posterior density of θ , β , σ_Y , and σ_B as follows.

$$p(\theta, \beta, \sigma_Y, \sigma_B | y, b) = \frac{p(y, b | \theta, \beta, \sigma_Y, \sigma_B) p(\theta) p(\beta) p(\sigma_Y) p(\sigma_B)}{\int p(y, b | \theta, \beta, \sigma_Y, \sigma_B) p(\theta) p(\beta) p(\sigma_Y) p(\sigma_B) d\theta d\beta d\sigma_Y d\sigma_B}.$$

The posterior density of the measurand θ is obtained by integration as

$$p(\theta | y, b) = \int p(\theta, \beta, \sigma_Y, \sigma_B | y, b) d\beta d\sigma_Y d\sigma_B.$$

This posterior distribution summarizes all of the information about θ available after the measurements were obtained. The mean of this distribution is taken as an estimate of the physical quantity. The uncertainty is the standard deviation of this distribution. It is straightforward to obtain a coverage interval for the measurand from this distribution. This is an interval of possible values for θ with a fixed probability; in Bayesian statistics this is called a *credible interval*. In many cases, numerical methods may be employed to accomplish the necessary integrations of Bayes' theorem. One possible method of obtaining a random sample from the posterior distribution is *Markov chain Monte Carlo* (MCMC) [13] using the software WinBUGS [2]. The code for this example, with the uniform prior distributions with $c = 100$, is as follows.

```
Example1a{
theta~dunif(0,1)
beta~dunif(0, 1)
gamma <- theta+beta

sigma.Y~dunif(0,100)
sigma.B~dunif(0,100)
tau.Y <- 1/(sigma.Y*sigma.Y)
tau.B <- 1/(sigma.B*sigma.B)

for(i in 1:n){
y[i]~dnorm(gamma,tau.Y)
b[i]~dnorm(beta,tau.B)}
}
```

With the data given in the Introduction, and $n = 5$, the program produces a posterior mean of θ of 2.315, and a posterior standard deviation of 0.235. A 95% credible interval for θ is (1.828, 2.798). A sensitivity analysis with respect to changes in the form of the four prior distributions can be carried out by varying the value of c , and by substituting the lines

```
tau.Y~dgamma(1.0E-5,1.0E-5)
tau.B~dgamma(1.0E-5,1.0E-5)
```

for the four lines

```
sigma.Y~dunif(0,1)
sigma.B~dunif(0,1)
tau.Y <- 1/(sigma.Y*sigma.Y)
tau.B <- 1/(sigma.B*sigma.B)
```

and comparing the resulting values of posterior mean and standard deviation. The results in this example are robust to such changes.

Example 1b

In Example 1b, the information about the background parameter β is provided in the form of a probability distribution, produced by a Type B evaluation. In this case, the observational model is only in terms of the the five Gaussian measurements; that is,

$$Y_i|\theta, \beta, \sigma_Y^2 \sim N(\theta + \beta, \sigma_Y^2).$$

There are now three parameters that must be assigned prior distributions. For the background parameter β , the prior density is based on the information given in the Introduction; that is,

$$\beta \sim \text{Uniform}(1.126, 1.329).$$

For θ ,

$$\theta \sim \text{Uniform}(0, c)$$

with a large value for c . For σ ,

$$\sigma_Y \sim \text{Uniform}(0, c).$$

This completes the prior distribution specifications. The WinBUGS code for this example is as follows.

```
Example1b{
theta~dunif(0,100)
beta~dunif(1.126,1.329)
sigma.Y~dunif(0,1)

gamma <- theta+beta
tau.Y <- 1/(sigma.Y*sigma.Y)
for(i in 1:n){
y[i]~dnorm(gamma,tau.Y)}
}
```

With the data given in the Introduction, the program produces a posterior mean of θ of 2.312, and a posterior standard deviation of 0.223. A 95% credible interval for θ is (1.871, 2.751). A sensitivity analysis of the results is again satisfactory.

Example 1c

Next consider Example 1c. The only difference from Example 1b is in the actual data values (which are now close to the background) and so the same model and WinBUGS code can be used here. The posterior mean of θ is now 0.068, the posterior standard deviation is 0.066 and the 95% credible interval is (0.00, 0.187). These results are robust to changes in the value of c with the uniform priors. Changing the form of the prior density for σ_Y from uniform to gamma results in a posterior mean of 0.058, posterior standard deviation of 0.053, and 95% credible interval of (0.0, 0.150). This is a larger change than in the previous examples, and indicates that here, because of the closeness of the data to the background, the data are not quite as informative about the measurand. The size of σ_Y (controlled to some degree by the prior distribution because there are only five observations on which an estimate is based) affects how informative the data are. In a case such as this, the conservative solution is to use the longer credible interval based on the uniform distribution. A better way would be to obtain more measurements. This would then greatly reduce the effect of the prior density of σ_Y on the results. (An interesting fact about the Bayes credible intervals such as these can be found in [5]. They show that in models such as Example 1, the 95% Bayes credible interval based on the uniform prior has frequentist coverage of close to 95%, whereas the interval based on the gamma prior usually has lower frequentist coverage.)

To summarize, Example 1a illustrates the case when measurements from two independent sources are used in a single uncertainty analysis. Example 1b shows how Type B evaluated information about the background can be included in the Bayesian model. Example 1c illustrates the ease with which any constraint can be included in the Bayesian model, as the positive constraint on the value of the measurand is critical here. It also shows how the choice of a noninformative prior distribution can affect the results.

4 Fiducial inference for uncertainty assessment

4.1 Basic method

Under the measurement equation (1) the uncertainty assessment for a measurand θ may be based on the fiducial distribution of θ . The following examples serve to illustrate the recipe for obtaining fiducial distributions for parameters of interest.

Suppose $Y \sim N(\theta, 1)$, where θ is the measurand, the measurement process has a known variance equal to 1, and Y is the random variable representing values that may be observed. One might express the relationship between the measured values and the underlying random experimental error process by the following equation,

$$Y = \theta + E, \quad (16)$$

where E is a random error with $N(0, 1)$ distribution. Each measured value is associated with a particular random experimental error. Suppose a single measurement is made and its value is 10. The associated measurement error is denoted by e . So

$$10 = \theta + e.$$

Hence $\theta = 10 - e$. If the value of e were known, then we would know the measurand exactly, but the value of e is not known. Nevertheless, the fact that we know the distribution from which e was generated helps us determine a set of values of θ that we consider plausible. For instance, how plausible is the value $\theta = 2$ for the measurand? For this to be true we need $e = 8$. A value of 8 is highly unlikely to have come from an $N(0, 1)$ distribution. So we conclude that the value $\theta = 2$ is highly unlikely. How likely is it that θ is between 10 and 12? For θ to be between 10 and 12, e needs to be between 0 and 2 and we can calculate the probability for this to be $\Phi(2) - \Phi(0)$, where $\Phi(z)$ is the value of the cumulative standard normal distribution at z . Thus, probabilities associated with E can be transferred to probabilities for θ . Our knowledge about θ , based on the measured value of 10, can be described by the distribution of the random variable $\tilde{\theta}$ whose distribution is given by that of $10 - E$. That is, $\tilde{\theta} \sim N(10, 1)$. We say that the fiducial distribution of θ (i.e., the distribution of $\tilde{\theta}$) is $N(10, 1)$. The random variable $\tilde{\theta}$ is also called a *fiducial quantity* (FQ) for θ . Such an FQ is related to what is called *generalized pivotal quantity* [29], [30] or *fiducial generalized pivotal quantity* [17], [31] in the literature.

In the above example, suppose we consider making two measurements. Let Y_1 and Y_2 be the random variables denoting the possible values one might obtain for the two measurements. We can write

$$\begin{aligned} Y_1 &= \theta + E_1, \\ Y_2 &= \theta + E_2. \end{aligned} \tag{17}$$

Suppose the actual measurements are 10 and 8. We then have the following equations relating the measured values, the measurand, and the realized values of experimental errors, say e_1 and e_2 ,

$$\begin{aligned} 10 &= \theta + e_1, \\ 8 &= \theta + e_2. \end{aligned}$$

Plausible values for θ are related to plausible values of (e_1, e_2) . What makes this example different from the previous example is that here we *know* $e_1 - e_2$ equals 2. So the universe of possible values for (e_1, e_2) is now limited by this requirement. We know (e_1, e_2) is from a standard bivariate Gaussian distribution but is constrained to lie on the line $e_1 - e_2 = 2$. So the probabilities one would associate with θ are the probabilities one would associate with either $10 - e_1$ or $8 - e_2$ knowing that (e_1, e_2) is a realization from a bivariate standard Gaussian distribution subject to the additional condition that $e_1 - e_2 = 2$.

Hence we define an FQ $\tilde{\theta}$ to have a distribution that is equal to the conditional distribution of $10 - E_1$ given that $E_1 - E_2 = 2$. This is the same distribution as the conditional distribution of $8 - E_2$ given that $E_1 - E_2 = 2$. A simple calculation tells us that the distribution of $\tilde{\theta}$ is $N(\bar{y}, 1/2)$ where $\bar{y} = (y_1 + y_2)/2 = (10 + 8)/2 = 9$.

More generally, if we have n independent measurements from $N(\theta, \sigma^2)$, then we can write

$$\begin{aligned} Y_1 &= \theta + \sigma E_1, \\ Y_2 &= \theta + \sigma E_2, \\ &\dots, \\ Y_n &= \theta + \sigma E_n, \end{aligned} \tag{18}$$

where E_1, \dots, E_n are independent, standard Gaussian random variables. The joint fiducial distribution for (θ, σ) can be obtained as follows. Use the first two (or any two) of the above n structural equations to solve for θ and σ , denoted by $\tilde{\theta}$ and $\tilde{\sigma}$, as functions of y_1, y_2, E_1 , and E_2 . The joint fiducial distribution for (θ, σ) is the joint distribution of $(\tilde{\theta}, \tilde{\sigma})$ conditioned on the E_i imposed by the rest of the $n - 2$ equations. In particular, the fiducial distribution for θ is shown to be

$$\tilde{\theta} = \bar{y} - \frac{s}{\sqrt{n}} T_{n-1}, \tag{19}$$

namely, a shifted and scaled Student- t distribution with $n - 1$ degrees of freedom. Here \bar{y} and s are the realized values of the sample mean \bar{X} and the sample standard deviation S of the n measurements, and T_{n-1} is a Student- t random variable with $n - 1$ degrees of freedom.

There is an alternative simpler method than that just outlined to derive a fiducial distribution for θ in (19), which is illustrated in the subsequent examples.

The above argument is fully generalizable and one can develop fiducial distributions for model parameters in very general problems. The starting point for this process is what we call a *structural equation* [12]. We denote this structural equation by $\mathbf{Y} = G(\boldsymbol{\beta}, \mathbf{E})$. In the single-measurement example, Equation (16) constitutes the structural equation. In the n -measurement example, Equations (18) constitute the structural equations. The structural equations relate the measurements \mathbf{Y} with model parameters $\boldsymbol{\beta}$ and error processes \mathbf{E} whose distributions are fully known. For instance, in the single-measurement example we know the distribution of E completely. For any fixed values of $\boldsymbol{\beta}$, the distribution of \mathbf{E} and the structural equation $G(\cdot)$ determine the distribution of the data \mathbf{Y} . After observing the data \mathbf{Y} we can switch the role of data and parameters. In particular, we fix the value of \mathbf{Y} and use the distribution of \mathbf{E} and the structural equation $G(\cdot)$ to infer a distribution on $\boldsymbol{\beta}$. This is what constitutes the fiducial argument.

4.2 Example 1

Example 1a

To illustrate, consider Example 1a described in the Introduction where the physical quantity θ is to be estimated from measurements that follow the model

$$Y_i = \theta + \beta + \epsilon_i, \quad i = 1, \dots, n, \quad (20)$$

where ϵ_i are independent measurement errors with $\epsilon_i \sim N(0, \sigma_Y^2)$. Also, β represents a background and can be estimated from measurements that follow the model

$$B_i = \beta + \delta_i, \quad i = 1, \dots, n_b, \quad (21)$$

where δ_i are independent measurement errors with $\delta_i \sim N(0, \sigma_B^2)$. It is also assumed that ϵ_i and δ_i are independent. From (20) and (21), it follows that $\bar{Y} - \bar{B}$ has a normal distribution with mean θ and variance $\sigma_Y^2/n + \sigma_B^2/n_b$, where \bar{Y} and \bar{B} are the means of Y_i and B_i , respectively. We can write

$$\bar{Y} - \bar{B} = \theta + \sqrt{\frac{\sigma_Y^2}{n} + \frac{\sigma_B^2}{n_b}} Z, \quad (22)$$

where Z is a standard normal random variable. This is a structural equation for $\bar{Y} - \bar{B}$. Also

$$W_y = \frac{(n-1)S_y^2}{\sigma_Y^2} \sim \chi^2(n-1)$$

and

$$W_b = \frac{(n_b-1)S_b^2}{\sigma_B^2} \sim \chi^2(n_b-1),$$

where $\chi^2(\nu)$ stands for the chi-squared distribution with ν degrees of freedom, S_y^2 and S_b^2 are sample variances of Y_i and B_i , respectively. Thus

$$S_y^2 = \frac{\sigma_Y^2 W_y}{n-1} \quad (23)$$

is a structural equation for S_y^2 and

$$S_b^2 = \frac{\sigma_B^2 W_b}{n_b-1} \quad (24)$$

is a structural equation for S_b^2 . By solving the above three structural equations for θ , σ_Y , and σ_B , we obtain an FQ for θ as

$$\tilde{\theta} = \bar{y} - \bar{b} - \sqrt{\frac{(n-1)s_y^2}{nW_y} + \frac{(n_b-1)s_b^2}{n_bW_b}} Z. \quad (25)$$

A $1 - \alpha$ fiducial interval for θ is given by $(\tilde{\theta}_{\alpha/2}, \tilde{\theta}_{1-\alpha/2})$, where $\tilde{\theta}_\alpha$ is the α quantile of the distribution of $\tilde{\theta}$. These quantiles may be determined analytically in simple situations. However, they are most conveniently estimated using a Monte Carlo approach. This involves generating a large number of realizations from the distribution of $\tilde{\theta}$ and determining the empirical $\alpha/2$ and $1 - \alpha/2$ quantiles. These quantiles are used as the estimates for $\tilde{\theta}_{\alpha/2}$ and $\tilde{\theta}_{1-\alpha/2}$. A single realization of $\tilde{\theta}$ may be generated as follows.

1. Generate a realization of a standard normal random variable Z .
2. Generate realizations of independent χ^2 random variables W_y and W_b with $n - 1$ and $n_b - 1$ degrees of freedom, respectively.
3. Calculate $\tilde{\theta}$ as in (25).

For this example, $n = n_b = 5$, $\bar{y} = 3.537$, $s_y = 0.342$, $\bar{b} = 1.228$, and $s_b = 0.131$. An R program for generating the 500,000 realizations of $\tilde{\theta}$ is listed below.

```
nrun = 500000
Z = rnorm(nrun)
W1 = rchisq(nrun, 4)
Wb = rchisq(nrun, 4)
theta = 3.537 - 1.228 - sqrt(4*0.342^2/(5*W1) +
4*0.131^2/(5*Wb))*Z
```

The mean of the simulated distribution is

```
mean(theta)
## 2.308893
```

and a 95% fiducial interval based on the 0.025 and 0.975 quantiles of the simulated distribution is

```
quantile(theta, c(0.025, 0.975))
##      2.5%      97.5%
## 1.857814 2.760931
```

Example 1b

In this example there are no statistical data on the background. It is assumed that the information regarding β is specified in terms of a probability distribution for β and that β and ϵ_i are independent. Furthermore, it is assumed the probability distribution for β is fully known, that is, does not involve any unknown parameters.

The structural equation for \bar{Y} is given by

$$\bar{Y} = \theta + \beta + \frac{\sigma}{\sqrt{n}} Z. \quad (26)$$

Together with the structural equation for S_y^2 in (23), we obtain an FQ for θ as

$$\tilde{\theta} = \bar{y} - \beta - \frac{s_y}{\sqrt{n}} \frac{Z}{\sqrt{W_y/(n-1)}}.$$

Because $Z/\sqrt{W_y/(n-1)} = T_{n-1}$ is a Student- t random variable with $n-1$ degrees of freedom, we have

$$\tilde{\theta} = \bar{y} - \beta - \frac{s_y}{\sqrt{n}} T_{n-1}. \quad (27)$$

A single realization of $\tilde{\theta}$ may be generated as follows.

1. Generate a realization of T_{n-1} of a Student- t random variable with $n-1$ degrees of freedom.
2. Generate β according its distribution, independently of T_{n-1} .
3. Calculate $\tilde{\theta}$ as in (27).

For this example, β is assumed uniformly distributed over the interval (1.126, 1.329). The 500,000 realizations of $\tilde{\theta}$ are generated by

```
beta = runif(nrun, 1.126, 1.329)
theta = 3.537 - beta - 0.342/sqrt(5)*rt(nrun, 4)
```

The mean of the simulated distribution is

```
mean(theta)
## 2.309454
```

and a 95% fiducial interval based on the 0.025 and 0.975 quantiles of the simulated distribution is

```
quantile(theta, c(0.025, 0.975))
##      2.5%      97.5%
## 1.871685 2.745590
```

The above fiducial interval agrees with the uncertainty interval obtained using the method proposed in Supplement 1 to the ISO *Guide*.

Example 1c

This example is identical to Example 1b except $\bar{y} = 1.196$ and $s_y = 0.106$. The 500,000 realizations of $\tilde{\theta}$ are generated by

```
theta = 1.196 - beta - 0.106/sqrt(5)*rt(nrun, 4)
```

The mean of the realizations is

```
mean(theta)
## -0.03158058
```

which lies outside the parameter space for θ . The number of realizations outside the parameter space can be found by

```
length((1:nrun)[theta < 0])
## 319168
```

The approach for handling parameter constraints is to simply truncate the fiducial distribution to the constrained parameter space. That is, we use

$$\max(\tilde{\theta}, 0)$$

to obtain the realizations of the fiducial distribution for θ . A 95% fiducial interval is calculated as

```
quantile(pmax(theta, 0), c(0.025, 0.975))
##      2.5%      97.5%
## 0.0000000 0.1361553
```

The recipe described above can be generalized to arbitrary statistical models. A prescription for constructing FQs is given in [16]. A simpler recipe for more common problems where sufficient statistics exist was given in a technical report [20] and is further discussed in [29], [30]. It is reproduced here for completeness. The recipe consists of the following steps.

1. Express each sufficient statistic as a function of one or more parameters and random variables whose distributions are completely known, free of any unknown parameters. That is, obtain a structural equation for each sufficient statistic.
2. In each structural equation, express each parameter as a function of the sufficient statistics and random variables whose distributions are completely known.
3. Obtain an FQ for each parameter by replacing the sufficient statistics with their corresponding observed values.

5 Example 2

Example 2, which is taken from Annex H.1 of the ISO *Guide*, focuses on the determination of the length of a new end gauge by comparing it with a nominally identical end gauge that has previously been calibrated. The notation used in the ISO *Guide* is followed as closely as possible, but has been modified where needed to agree with the notation conventions discussed earlier that distinguish measurands from observable measurement results. Table 2 lists the physical quantities needed to determine the length of the new end gauge.

Using the new notation given in Table 2 and based on

- the first line of equation H.2 in the ISO *Guide*,
- the relationships $\alpha = \alpha_s + \delta_\alpha$ and $\theta_s = \bar{\theta} + \Delta - \delta_\theta$ defined in Section H.1.2, and
- inferences drawn from the propagation of uncertainties in Sections H.1.3.2 and H.1.3.4 of the ISO *Guide*,

the measurement equation for λ used in the ISO *Guide* analysis of example H.1 can be expressed as

$$\lambda = \frac{\lambda_s [1 + \alpha_s (\bar{\theta} + \Delta - \delta_\theta)] + \delta_\lambda + \delta_{C_r} + \delta_{C_{nr}}}{1 + (\alpha_s + \delta_\alpha) (\bar{\theta} + \Delta)}. \quad (28)$$

Note that Equation (28) is the exact measurement equation as described in Section H.1.1 and the first line of Equation H.2 of the ISO *Guide*, rather than the approximation that is made on the second line of Equation H.2 and then used throughout the rest of Section H.1.

Table 2. New notation used for the reanalysis of Example H.1 of the ISO *Guide* under each of the three statistical paradigms. The random variable corresponding to δ_λ is denoted \bar{D}_λ , and its observed value \bar{d}_λ .

Quantity	Symbol
Length of unknown end gauge at 20°C	λ
Length of standard end gauge at 20°C	λ_s
Difference between end gauge lengths at lab temperature	δ_λ
Correction to difference between end gauge lengths to compensate for random comparator errors	δ_{C_r}
Correction to difference between end gauge lengths to compensate for systematic comparator errors	$\delta_{C_{nr}}$
Coefficient of thermal expansion of the standard end gauge	α_s
Difference in coefficients of thermal expansion of the standard and unknown end gauges	δ_α
Average deviation of testbed temperature from standard conditions during data collection	$\bar{\theta}$
Cyclic variation of testbed temperature from mean temperature due to thermostatic control	Δ
Difference in temperatures of the standard and unknown end gauges	δ_θ

Table 3. Summary of information from the analysis of Example H.1 in the ISO *Guide* needed for its reanalysis.

Quantity	Value	Standard Uncertainty	Degrees of Freedom	Type	Assumed Distribution
λ_s	50000623 nm	25 nm	18	B	Normal
\bar{d}_λ	215 nm	5.8 nm	24	A	Normal
δ_{C_R}	0 nm	3.9 nm	5	B	Normal
$\delta_{C_{nr}}$	0 nm	6.7 nm	8	B	Normal
α_s	$11.5 \times 10^{-6} \text{ }^\circ\text{C}^{-1}$	$1.2 \times 10^{-6} \text{ }^\circ\text{C}^{-1}$		B	Rectangular
δ_α	$0 \text{ }^\circ\text{C}^{-1}$	$0.58 \times 10^{-6} \text{ }^\circ\text{C}^{-1}$	50	B	Rectangular
θ	$-0.1 \text{ }^\circ\text{C}$	$0.2 \text{ }^\circ\text{C}$		B	Not Specified
Δ	$0 \text{ }^\circ\text{C}$	$0.35 \text{ }^\circ\text{C}$		B	Arcsine
δ_θ	$0 \text{ }^\circ\text{C}$	$0.029 \text{ }^\circ\text{C}$	2	B	Rectangular

Equation (28) is also written in terms of the complete list of physical quantities used to determine the end gauge length, rather than pre-summarizing the effects due to the difference between the lengths of the two gauges and in the temperature of the testbed. It is always good practice to write out the measurement equation for the measurand of interest in terms of the complete list of quantities needed to determine it. This practice is important because it helps minimize the potential failure to identify correlations between different physical quantities, such as θ and θ_s and α and α_s as mentioned in Section H.1.2, whose values ultimately might be based on the same data.

Table 3 summarizes the rest of the information taken from the analysis of Example H.1 in the ISO *Guide* needed for the reanalysis of the example under the different statistical paradigms discussed and compared in the remainder of this section.

The description of the example in the ISO *Guide* indicates that there is only one quantity in this example, δ_λ , whose value has been directly estimated via the analysis of statistical data. The distribution of the sample mean of the measurements has been assumed to be Gaussian (or normal) with an expected value that depends on the length of the unknown end gauge and the other physical quantities described in Tables 2 and 3.

The values and standard uncertainties of all other quantities are assumed to have been evaluated by Type B methods. Because the quantities δ_α and δ_θ follow rectangular rather than Gaussian distributions, however, there are no widely accepted statistical methods to account for the degrees of freedom in these two cases. As a result, we do not use the given degrees of freedom for those quantities.

5.1 Frequentist approach

In the situation of end-gauge calibration in this example, the sensitivity coefficients $c_{\alpha_s} = c_{\theta_s}$ vanish, and the second-order terms are to be incorporated

in (5), although just one of them is noticeably different from zero (ISO *Guide*, p 71).

The ISO *Guide*'s answer, $y = 50000838$ nm, with the corresponding combined standard uncertainty $u_c(y) = 34$ nm can be interpreted as the plug-in estimate of the end gauge length. As was mentioned in Section 2.1, the uncertainties of these estimators are approximated by the marginal quadratic error if the parameters λ_s, θ_s are averaged over their normal distributions, and $\delta_\alpha, \alpha_s,$ and δ_θ are integrated out according to their uniform distributions.

These results are confirmed by Monte Carlo simulations which provide a very close answer. Moreover, the approximation by t -distribution (8) seems to be reasonable. See Figure 1 where the empirical percentiles are plotted against t -distribution quantiles when the degrees of freedom are estimated according to (6). More simulation results are reported in [28].

To construct a bootstrap interval for this example, λ is estimated to be 50000838 nm with a combined standard uncertainty $u_c = 31.7$ nm. From (11), the $100(1 - \alpha)\%$ bootstrap- t confidence interval is $(50000838 - \hat{t}_{1-\alpha/2} \cdot 31.7, 50000838 - \hat{t}_{\alpha/2} \cdot 31.7)$ nm, where \hat{t}_β is the 100β th percentile of W^* of (10). The R code for generating $B = 10000$ realizations of W^* is as follows.

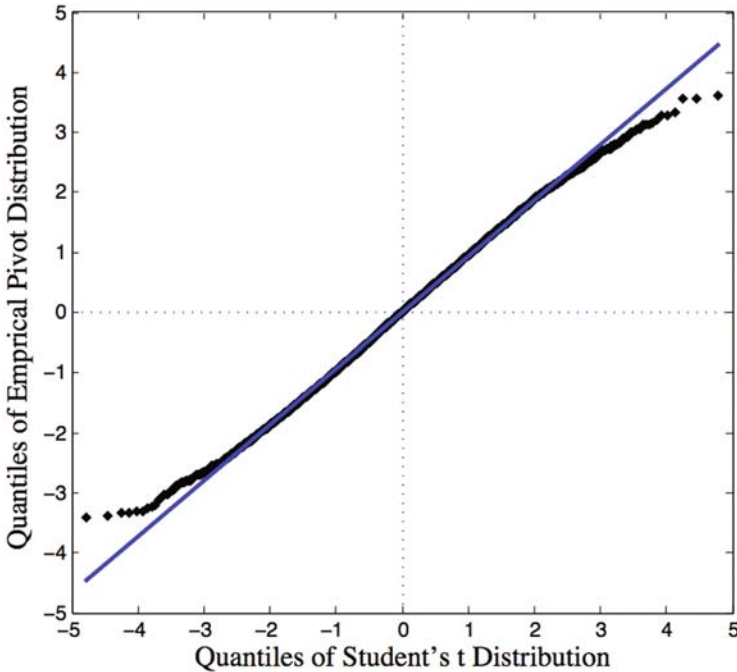


Fig. 1. Plot of empirical percentiles versus percentiles of t -distribution in Example 2.

```

B = 10000
x.star = cbind(
  rnorm(B, mean=50000623, sd=25),
  rnorm(B, mean=215, sd=5.8),
  rnorm(B, mean=0, sd=3.9),
  rnorm(B, mean=0, sd=6.7),
  runif(B, min=0.0000095, max=0.0000135),
  runif(B, min=-0.000001, max=0.000001),
  runif(B, min=-0.45, max=0.25),
  rbeta(B, 0.5, 0.5)-0.5,
  runif(B, min=-0.05, max=0.05))
u.star = cbind(
  25 * sqrt(rchisq(B, df=18)/18),
  5.8 * sqrt(rchisq(B, df=24)/24),
  3.9 * sqrt(rchisq(B, df=5)/5),
  6.7 * sqrt(rchisq(B, df=8)/8),
  0.0000012,
  0.00000058 * sqrt(rchisq(B, df=50)/50),
  0.2,
  0.35,
  0.029 * sqrt(rchisq(B, df=2)/2))
x.name = c("L.s", "D.lambda", "Dc.r", "Dc.s", "A.std", "D.alpha",
           "T.bar", "T.cv", "D.theta")
f = expression((L.s*(1+A.std*(T.bar+T.cv-D.theta))+
               D.lambda+Dc.r+Dc.s)/
               (1+(A.std+D.alpha)*(T.bar+T.cv)))

star = delta(f, x.star, u.star, x.name)
w.star = (star$y - 50000838) / star$uc

```

The R function `delta` takes partial derivatives of the R expression `meq`, the measurement equation, with respect to parameters `namevec`, evaluates them at the input values `x`, and approximates the uncertainty of the computed value of the measurement equation, `eval(meq)`, based on the first-order Taylor approximation to the measurement equation.

```

delta = function(meq,x,u,namevec){
  for(i in 1:ncol(x)) assign(namevec[i], x[,i])
  c = attr(eval(deriv(meq,namevec)), "gradient")
  list(y = eval(meq), uc = sqrt(apply((c*u)^2, 1, sum)))
}

```

The 95% bootstrap-*t* confidence interval based on the 0.025 and 0.975 quantiles of the simulated distribution is


```
50000838 - quantile(w.star, c(0.975, 0.025)) * 31.70511
## 50000777 50000899
```

This interval (50000777, 50000899) nm is almost 10% shorter than ISO *Guide's* answer. This general behavior, the width of a bootstrap interval being shorter than that of an interval derived from the uncertainty analysis based on the first-order Taylor approximation, is further discussed in [7].

5.2 Bayesian approach

The measurement problem is to determine the length of a nominally 50 nm gauge. In the notation introduced here, the measurand is λ . The measurement is performed by comparison with a known standard of the same nominal length. The direct measurement output (the data) of the comparison is the difference in their lengths. This is given as an average \bar{d}_λ of five measurements. Following the process outlined above, an observational model relating the data to the parameters needs to be specified.

The ISO *Guide* can be interpreted as stating that the expected value of the measurement $E(D_\lambda)$ is equal to δ_λ where

$$\delta_\lambda = \lambda (1 + (\delta_\alpha + \alpha_s) (\bar{\theta} + \Delta)) - \lambda_s (1 + ((\bar{\theta} + \Delta) - \delta_\theta) \alpha_s). \quad (29)$$

The expected value of the measurements is a function of the parameter vector $\gamma = (\lambda_s, \bar{\theta}, \Delta, \delta_\alpha, \alpha_s, \delta_\theta)$ and of the measurand λ . The ISO *Guide* gives two additional components of Type B uncertainty associated with the comparator. This implies that there is uncertainty about the expected value of the difference measurement being equal to δ_λ . Similarly to the ISO *Guide*, the two components can be combined additively to obtain an uncertainty of 7.8 nm, with 12 degrees of freedom using the Welch–Satterthwaite formula. The following two-stage statistical model combines all of the available information

$$\bar{D}_\lambda | \delta_{\lambda_r} \sim N(\delta_{\lambda_r}, \sigma_{D_\lambda}^2/5) \quad (30)$$

$$\delta_{\lambda_r} | \delta_\lambda \sim \delta_\lambda + 7.8 * T_{12}.$$

Another piece of data given in the example is the Type A evaluated uncertainty associated with the measured difference s_{d_λ} . This is an estimate of σ_{D_λ} . A well-known result from basic probability theory is that for a sample of size n , from a Gaussian distribution with a known variance σ^2 ,

$$\frac{(n-1)}{\sigma^2} S^2 \sim \chi^2(n-1).$$

As $\chi^2(n-1)$ is also the Gamma $((n-1/2), \frac{1}{2})$ density, this results in

$$S_{d_\lambda}^2 | \sigma_{D_\lambda} \sim \text{Gamma}\left(\frac{(25-1)}{2}, \frac{1}{2\sigma_{D_\lambda}^2}\right).$$

There are eight parameters in the statistical model, including the measurand λ . To find the posterior distribution of λ , it is necessary to first specify the joint prior distribution of the eight parameters. A priori, these random variables can be assumed to be independent and so their joint distribution is the product of their individual prior distributions. For the elements of the parameter vector γ , the Type B evaluated information can be interpreted as informative prior densities as follows.

$$\lambda_s \sim N(50000623, 625) \quad (31)$$

$$\delta_\alpha \sim \text{Uniform}(-1 \times 10^{-6}, 1 \times 10^{-6})$$

$$\bar{\theta} \sim N(-0.1, 0.1681)$$

$$\Delta \sim \text{Beta}(0.5, 0.5) - 0.5$$

$$\alpha_s \sim \text{Uniform}(9.5 \times 10^{-6}, 13.5 \times 10^{-6})$$

$$\delta_\theta \sim \text{Uniform}(-0.05, 0.05).$$

We obtain

$$p(\gamma) = p(\lambda_s) p(\delta_\alpha) p(\bar{\theta}) p(\Delta) p(\alpha_s) p(\delta_\theta).$$

Prior distributions for the measurand λ and σ_{D_λ} are needed to complete the prior specification. In this example, there is no additional information about these two parameters other than that they are not negative. As in Example 1, the parameters are given *reference* priors [4]. For λ , a reference density is approximated as

$$\lambda \sim \text{Uniform}(0, c) \quad (32)$$

with a large value for c . For σ_{D_λ} we can also use

$$\sigma_{D_\lambda} \sim \text{Uniform}(0, c), \quad (33)$$

or the Gamma(c, c). This completes the prior distribution specification.

Note that the two reference prior distributions, which sensitivity analysis shows have minimal impact on the results, are the only distributions not used in some manner by the frequentist or fiducial approaches.

Application of Bayes' theorem results in the joint posterior density of $(\lambda, \gamma, \sigma_{D_\lambda})$ as follows.

$$p(\lambda, \gamma, \sigma_{D_\lambda} | \bar{d}_\lambda, s_{d_\lambda}) = \frac{p(\bar{d}_\lambda | \delta_{\lambda_r}) p(\delta_{\lambda_r} | \delta_\lambda) p(\gamma) p(\lambda) p(\sigma_{D_\lambda})}{\int p(\bar{d}_\lambda | \delta_{\lambda_r}) p(\delta_{\lambda_r} | \delta_\lambda) p(\gamma) p(\lambda) d\gamma d\lambda d\sigma_{D_\lambda}}.$$

The posterior density of λ is then obtained by integration as

$$p(\lambda | \bar{d}_\lambda, s_{d_\lambda}) = \int p(\lambda, \gamma, \sigma_{D_\lambda} | \bar{d}_\lambda, s_{d_\lambda}) d\gamma d\sigma_{D_\lambda}.$$

This posterior distribution summarizes all of the information about λ available after the measurements were obtained. The WinBUGS code for this example is as follows.

```

Example2{
n<-25
df<-(n-1)/2

lambda~dnorm(0,1.0E-18)
delta.a~dunif(-0.000001, 0.000001)
alpha~dunif(0.0000095,0.0000135)
theta~dnorm(-0.1,5.94)
ddelt~dbeta(0.5,0.5)
delta<-ddelt-0.5
delta.t~dunif(-0.05,0.05)
lambda.s~dnorm(50000623, 0.0016)
sigma.D~dunif(0,20)
tau.D<-1/(sigma.D*sigma.D)

delta.l<-lambda*(1+(delta.a+alpha)*(theta+delta))
      -lambda.s*(1+((theta+delta)-delta.t)*alpha)
delta.l.r~dt(delta.l, 0.0164,12)
msg<-5*tau.D
dbar~dnorm(delta.l.r,msg)

pg<-tau.D/2
ssq<-(n-1)*s.y*s.y
ssq~dgamma(df,pg)
}

```

With data input of $\bar{d}_\lambda = 215$ and $s_Y = 13$, this program obtains the posterior mean of λ as 50000837 nm, with posterior standard deviation of 34 nm. The 95% credible interval is (50000768, 50000908) nm. These results are almost identical to the results in the *ISO Guide*.

Note that in the solution given here, the measurement equation in terms of λ , that is, Equation (28), is never used. This avoids the unnecessary and difficult task of determining how the distributions of the various parameters are related. As in Example 1 with the two parameters, the approach given here leads to the correct joint posterior distribution of all eight parameters.

It is interesting to consider an approximate solution for this example based on the Taylor series approximation. In the *ISO Guide* solution, Equation (1) is approximated as

$$\delta_\lambda = \lambda - \lambda_s \left(1 - (\delta_\alpha (\bar{\theta} + \Delta) + \alpha_s \delta_\theta) \right).$$

Define a parameter $\eta = \lambda - \delta_\lambda$. Using the Taylor series approximation, the probability density of η can be approximated by a Gaussian as $\eta \sim$

$N(50000623, 911.47)$, For simplicity also approximate σ_{D_λ} by s_{d_λ} . Then the statistical model becomes

$$\begin{aligned}\bar{D}_\lambda | \delta_{\lambda_r} &\sim N\left(\delta_{\lambda_r}, \frac{(13)^2}{5}\right) \\ \delta_{\lambda_r} | \delta_\lambda &\sim N\left(\lambda - \eta, (7.8)^2\right) \\ \lambda &\sim N(0, c) \\ \eta &\sim N(50000623, 911.47).\end{aligned}$$

For this model, the posterior density of λ can be obtained analytically [24]. We get

$$\lambda \sim N\left(\bar{d} + 50000623, \frac{(13)^2}{5} + (7.8)^2 + 911.47\right)$$

Because $\bar{d}_\lambda = 215$ nm, we obtain the posterior mean of λ as 50000838 nm with posterior standard deviation of 31.3 nm, again results very close to the ISO *Guide* solution.

5.3 Fiducial approach

We use this example to illustrate the fiducial inference approach in a more complex application. The measurement equation is given in (28). Five independent repeated measurements are available for estimating δ_λ . Based on the information provided in the ISO *Guide*, the following assumptions are made.

1. The estimated value of λ_s (i.e., the value given in the calibration certificate), denoted by l_s , is equal to 50000623 nm. The standard uncertainty of the estimate is 25 nm with 18 degrees of freedom. Under the normality assumption, an FQ for λ_s is given by

$$\tilde{\lambda}_s = 50000623 - 25 T_{18}. \quad (34)$$

This is obtained from (19) with $\bar{y} = 50000623$ nm, $u(\bar{y}) = 25$ nm, and 18 degrees of freedom associated with $u(\bar{y})$.

2. Each replicate measurement has a normal distribution with mean δ_λ and standard deviation σ_{δ_λ} . The observed mean of the five repeated measurements, denoted by \bar{d}_λ , is 215 nm. The value of σ_{δ_λ} is estimated from a separate experiment to be 13 nm with 24 degrees of freedom. This gives $u(\bar{d}_\lambda) = 13/\sqrt{5}$. So an FQ for δ_λ is given by

$$\tilde{\delta}_\lambda = 215 - 13 T_{24}/\sqrt{5}. \quad (35)$$

Also, based on the calibration certificate for the comparator device, the estimate of δ_{C_r} is 0 with a standard uncertainty of 3.9 nm (5 degrees of

freedom), and the estimate of $\delta_{C_{nr}}$ is 0 with a standard uncertainty of 6.7 nm (8 degrees of freedom). Furthermore, the comparator errors can be assumed to be independent of the replication errors. Thus, we may write

$$\tilde{\delta}_{C_r} = 3.9 T_5 \tag{36}$$

and

$$\tilde{\delta}_{C_{nr}} = 6.7 T_8. \tag{37}$$

Mutual independence among the Student-*t* random variables is a consequence of the ISO *Guide* assumption about the measurement process.

3. Let $\bar{\theta}$ be the true deviation of the average temperature of the testbed from the nominal value of 20 °C. An estimate of $\bar{\theta}$ is -0.1 °C with a standard deviation equal to 0.2 °C. The ISO *Guide* gives no additional information concerning this standard deviation, therefore we assume infinite degrees of freedom for it and we assume that $\bar{\theta}$ follows a normal distribution. Hence we have

$$\tilde{\bar{\theta}} = -0.1 - 0.2 Z, \tag{38}$$

where Z is a standard normal random variable, independent of all other random variables.

4. An FQ for Δ has a probability density function given by

$$g(x) = \frac{2}{\pi\sqrt{1-4x^2}}, \quad -0.5 \text{ °C} < x < 0.5 \text{ °C}.$$

For simulating realizations from the arcsine distribution above, it is useful to observe that if U_1 is a uniform (0, 1) random variable, then $-\cos(\pi U_1)/2$ has the required arcsine distribution. So an FQ for Δ may be taken to be

$$\tilde{\Delta} = -\cos(\pi U_1)/2. \tag{39}$$

5. An FQ for δ_α is given by

$$\tilde{\delta}_\alpha = U_2, \tag{40}$$

where U_2 is a uniform random variable over the interval $\pm 1 \times 10^{-6}$ °C⁻¹.

6. An FQ for δ_θ is given by

$$\tilde{\delta}_\theta = U_3, \tag{41}$$

where U_3 is a uniform random variable over the interval ± 0.05 °C.

7. An FQ for α_s is given by

$$\tilde{\alpha}_s = 11.5 \times 10^{-6} + U_4, \tag{42}$$

where U_4 is a uniform random variable over the interval $\pm 2 \times 10^{-6}$ °C⁻¹.

Substituting the fiducial quantities in (34)–(42) into (28), we obtain a fiducial quantity for λ . We estimate the distribution of $\tilde{\lambda}$ using 500000 Monte Carlo samples. The mean and the standard deviation of this simulated distribution are 50000838 nm and 35 nm, respectively. A 95% fiducial interval for λ can be obtained by using the interval between the 0.025 and 0.975 quantiles

of the simulated distribution, which is given by (50000768, 50000907) nm. An R program for generating the 500000 realizations of $\tilde{\lambda}$ is listed below.

```
nrun = 500000
lambda.s = 50000623 - 25 * rt(nrun, 18)
delta.lambda = 215 - 13/sqrt(5) * rt(nrun, 24)
delta.cr = 3.9*rt(nrun, 5)
delta.cnr = 6.7*rt(nrun, 8)
theta = rnorm(nrun, -0.1, 0.2)
Delta = (-cos(pi*runif(nrun))/2)
delta.alpha = runif(nrun, -10^(-6), 10^(-6))
delta.theta = runif(nrun, -0.05, 0.05)
alpha.s = runif(nrun, (11.5-2)*10^(-6), (11.5+2)*10^(-6))
lambda = (lambda.s * (1 + alpha.s*(theta + Delta -
  delta.theta)) + delta.lambda + delta.cr + delta.cnr)/
  (1 + (alpha.s + delta.alpha)*(theta + Delta))
```

6 Discussion

Table 4 summarizes the results for Example 1. The frequentist bootstrap, Bayesian, and fiducial solutions for Example 1a and Example 1b are very similar. The bootstrap and the ISO *Guide* solutions produce slightly shorter intervals in both Example 1a and Example 1b. More substantial differences are seen in the solution for Example 1c. Here the Bayes solution based on the uniform prior density produces an interval that is quite a bit longer than most of the other methods; only the conservative Eisenhart interval is longer.

Because the Bayesian and fiducial approaches actually both produce full probability distributions for the measurand θ , their results for Example 1a and Example 1c are further compared in Figure 2, in addition to the comparison of the expanded uncertainty intervals in Table 4. The results for Example 1b are not displayed because they are visually indistinguishable from the results for Example 1a. From the histograms in Figure 2 it is clear that the Bayesian posterior probability distribution for θ and the fiducial distribution for θ are quite similar when the signal is well above the background. When the signal is near the background, however, the two distributions have very different

Table 4. Expanded uncertainty intervals constructed under the three statistical paradigms for Example 1.

	ISO <i>Guide</i>	Eisenhart	Bootstrap	Bayes	Fiducial
Example 1a	(1.89, 2.73)	(1.89, 2.73)	(1.90, 2.73)	(1.83, 2.80)	(1.86, 2.76)
Example 1b	(1.90, 2.72)	(1.78, 2.84)	(1.92, 2.70)	(1.87, 2.75)	(1.87, 2.75)
Example 1c	(0.00, 0.12)	(0.00, 0.20)	(0.00, 0.11)	(0.00, 0.19)	(0.00, 0.14)

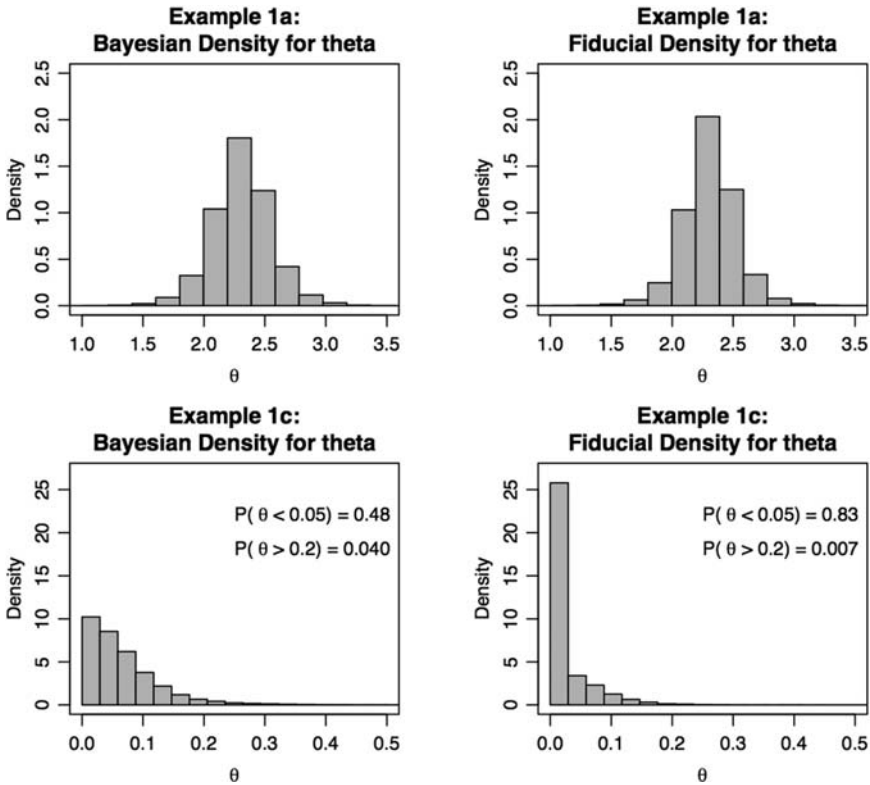


Fig. 2. Comparisons of the simulated Bayesian and fiducial densities for Example 1a and Example 1c.

characteristics due to their different methods of incorporating the physical constraints inherent in the problem.

In the frequentist context, the measurand θ and the input quantities μ_1, \dots, μ_p of the measurement equation (1) are all assumed to be fixed unknown quantities. This approach seems to be quite reasonable if the measurand represents a physical constant for which previous studies do not provide for an appropriate (informative) prior distribution or structural equation. It is favored by statisticians who do not believe that all parameters must be modeled as random variables, although it typically handles uncertainties of Type B as nuisance parameters, that is, by assigning them a probability distribution and integrating over this distribution. In this regard it is similar to the Bayesian approach, where all parameters must have their distributions, but needs fewer distributional assumptions.

The bootstrap is a well-established statistical method that can replace complicated and often inaccurate approximate confidence intervals by computer simulations. There are various bootstrap schemes developed to construct

confidence intervals under different conditions. The parametric bootstrap- t interval, introduced in this chapter, is the natural choice as an improvement to the Student- t interval of the ISO *Guide*. The advantage of bootstrapping is its simplicity: it is straightforward to apply the bootstrap to derive confidence intervals as demonstrated in the examples. Nonparametric bootstrap methods provide even bigger advantages when the datasets are large.

It was shown with the examples that Bayesian uncertainty assessment via the statistical model is conceptually simple, and can be applied to complex measurement problems without any changes to the basic method. Systematic effects, which are not estimable from the measurements (i.e., there are no functions of the observations whose expected values are equal to the systematic effect) and for which information is obtained by Type B evaluations, can easily be included in the Bayesian model. Computation of posterior distributions can be done using MCMC methods, often using existing software. As was seen, there is no need for asymptotic arguments to justify the probability statements, small and large samples share a common, exact probabilistic justification.

There are some drawbacks to the Bayesian methods described here. Most important, to use them one must specify prior distributions for all parameters in the measurement model, including the measurand. Even though in metrology informative prior distributions are often available in the form of the Type B uncertainty evaluations, it is usually the case that one or two of the parameters will need to be assigned vague (noninformative) prior distributions because of lack of prior knowledge. Such distributions are not unique, and as was demonstrated in Example 1c, they can influence the results. It is therefore always necessary to perform sensitivity analysis to determine the size of any such effect, which should be small. This was so for all of our examples except 1c. Large effects from the specification of a noninformative prior require further study of the measurement system. Generally this means that there is not enough information in the data about the measurand and thus the prior distribution has too much influence on the result. In some cases, this problem can be solved by increasing the sample size, or by changing the way in which the data are collected, for example, by improving resolution. In other situations, it may be that the mathematical model being used has too many parameters about which we have no real prior information and so the model must be simplified.

There is a silver lining to the need for prior specification and that is that when substantial prior information about the measurand does exist, it can be introduced simply, and updated efficiently via Bayes' theorem. Furthermore, sensitivity to the prior form, not just for the measurand, but especially for the standard deviation associated with the likelihood function, can be a good indication that there are problems with the measurement system. These can then be corrected.

Fiducial inference provides a framework for associating a distribution with a parameter of interest. Recent research results [17] show that fiducial

inference is a valid statistical method with generally good operating characteristics. The examples used demonstrated that the fiducial approach can easily and naturally incorporate the uncertainty information into the measurement equation, and calculate the final estimate and its combined standard uncertainty for a measurand of interest by propagating the component statistical distributions. There is no need for propagation of uncertainty based on Taylor series expansions or the Welch–Satterthwaite approximation under the fiducial approach.

There is an issue of nonuniqueness in the fiducial distribution due to the choice of a particular form of the structural equation. However, it is important to note that, in many practical applications, the physical process by which the data was generated is known. In this case we can and should choose the structural equation to reflect this process, thus eliminating the problem of nonuniqueness due to the choice of structural equation. In the field of metrology where an unknown measurand is measured using some known processes, one typically knows that random errors influence the measurement in some pre-specified known fashion. The resulting measured values are expressed as an equation combining some unknown measured quantities and errors. This formula can be taken as the structural equation.

7 Chapter summary

In this chapter we have discussed three different approaches for constructing uncertainty intervals that each have a clear probabilistic interpretation. This contrasts with much of the other work in this area that has focused on assessing the statistical properties of procedures currently in popular use across the metrology community. One of the goals in approaching the study of methods for uncertainty assessment from a different vantage point was to try to gain insight into current methods and highlight new options that may also prove useful.

As Lira and Wöger [25] observed, the uncertainty intervals obtained under the different paradigms will often be similar numerically. Even when this is the case, however, their interpretations are quite different from one another.

Frequentist uncertainty intervals make probabilistic statements about the long-term performance of a particular procedure for constructing uncertainty intervals during repeated use under identical conditions. Thus the probability statement is not directly about the value of the measurand, but is about the long-term relationship between the interval construction procedure and the measurand. Once the data have been observed and a frequentist uncertainty interval has been computed, there is no longer anything random about the results. Although it is not known whether the value of the measurand is captured in a particular realized interval, the analyst does know that such intervals will capture the value of the measurand with a specified probability. Unlike a traditional confidence interval based only on statistical data, the

frequentist uncertainty interval is typically constructed so that the desired confidence level is attained on average after integrating over the probability distributions of any quantities that must be evaluated using Type B methods.

Bayesian and fiducial uncertainty intervals, on the other hand, are based on probability distributions that directly describe our knowledge of the value of the measurand. The methods used to obtain these two types of intervals are different, but the results are similar in this aspect of their interpretation. The Bayesian results are obtained by combining probability distributions for each parameter specified prior to analysis of the data with a probability model that describes the variation in the data using Bayes' theorem. The resulting posterior distributions for each parameter reflect the probability of the parameter values given the prior information and the data. The fiducial results are obtained by inverting a probability model for the data given the parameters to find a distribution for the parameter values given the data.

Of course if the numerical results were always quite similar, then each of the different interpretations would be simultaneously applicable (at least approximately) to every uncertainty interval. However, as we have seen, the numerical results can differ significantly from one another in some instances, even though each can be justified probabilistically and they share a common level of significance (generally 95%). Other differences also may be observed. For example, if one of the dominant sources of uncertainty in a particular application has a skewed distribution, the uncertainty intervals obtained using the Bayesian or fiducial approaches will often reflect that asymmetry whereas an approximate confidence interval obtained using the procedures of the *ISO Guide* will produce a symmetric uncertainty interval (and may be longer than necessary on one side). Frequentist results based on other statistical principles may match the Bayesian or fiducial results in some cases, but the different methods will never all agree in general because each paradigm is ultimately based on a different set of unique mathematical assumptions and criteria.

The existence of different paradigms for uncertainty assessment that do not always agree might be seen as an unfortunate complication by some. However, we feel it is better seen as an indication of further opportunity. It is only by continually working together to appreciate the features of different paradigms that we will arrive at methods for uncertainty assessment that meet all of our scientific and economic needs: methods that are practical to implement, make efficient use of resources, are applicable to many types of measurements, both old and new, and are transparent in meaning.

Acknowledgments

The authors thank Antonio Possolo, Charles Hagwood, William Strawderman, David Duewer, Leon Gleser, Pin-Hao Wang, and the anonymous referees for comments that helped to improve this chapter.

References

1. http://www.bipm.org/utils/en/pdf/news_jcgm-wg1.pdf
2. <http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml>
3. T. Bayes. An essay toward solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418, 1764. (facsimile available at <http://www.stat.ucla.edu/history/essay.pdf>)
4. J. M. Bernardo and A. F. M. Smith. *Bayesian Theory*. John Wiley and Sons, Chichester, UK, 1994
5. W. J. Browne and D. Draper. A comparison of Bayesian and likelihood-based methods for fitting multi-level models. *Bayesian Analysis*, 1:473–514, 2006
6. G. Casella and R. Berger. *Statistical Inference*. Duxbury, MA, 2 edition, 2002
7. B. Efron. Six questions raised by the bootstrap. In R. LePage and L. Billard, editors, *Exploring the Limits of Bootstrap*, pages 99–126. Wiley, New York, 1992
8. B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*, volume 57 of *Monographs of Statistics and Applied Probability*. Chapman and Hall, Boca Raton, FL 1993
9. C. Eisenhart. Expression of the uncertainties of final measurements results. NBS special publication, NIST, Gaithersburg, MD, 1983
10. C. Elster, W. Wöger, and M. G. Cox. Draft GUM Supplement 1 and Bayesian analysis. *Metrologia*, 44:L31–L32, 2007
11. R. A. Fisher. Inverse probability. *Proc. Comb. Philos. Soc.*, 26:528–535, 1930
12. D. A. S. Fraser. *The Structure of Inference*. Krieger, New York, 1968
13. A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, Boca Raton, FL, 1995
14. L. J. Gleaser. Assessing uncertainty in measurement. *Statistical Science*, 13: 277–290, 1998
15. W. F. Guthrie. Should (T1-T2) have larger uncertainty than T1? *Proceedings of the 8th International Conference on Temperature: Its Measurements and Control*, 2:887–892, 2002. <http://www.itl.nist.gov/div898/pubs/author/guthrie/guthrie-2002-01.pdf>
16. J. Hannig. On fiducial inference – the good, the bad and the ugly. Technical Report 2006/3, Department of Statistics, Colorado State University, Fort Collins, CO, 2006. http://www.stat.colostate.edu/research/2006_3.pdf
17. J. Hannig, H. K. Iyer, and P. L. Patterson. Fiducial generalized confidence intervals. *Journal of the American Statistical Association*, 101:254–269, 2006
18. T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York, 2001
19. International Organization for Standardization. *Guide to the Expression of Uncertainty in Measurement*, 1993
20. H. K. Iyer and P. L. Patterson. A recipe for constructing generalized pivot quantities and generalized confidence intervals. Technical Report 2002/10, Department of Statistics, Colorado State University, Fort Collins, CO, 2002. http://www.stat.colostate.edu/research/2002_10.pdf
21. JCGM/WG 2. *International Vocabulary of Metrology - Basic and General Concepts and Associated Terms (VIM)*, 3rd edition, 2007
22. R. Kacker and A. Jones. On use of Bayesian statistics to make the guide to the expression of uncertainty in measurement consistent. *Metrologia*, 40:235–248, 2003

23. L. Kirkup and B. Frenkel. *An Introduction to Uncertainty in Measurement Using the GUM*. Cambridge University Press, Cambridge, UK, 2006
24. D. Lindley and A. F. M. Smith. Bayes estimates for the linear model. *JRSS B*, 34:1–41, 1972
25. I. Lira and W. Wöger. Comparison between the conventional and Bayesian approaches to evaluate measurement data. *Metrologia*, 43:S249–S259, 2006
26. D. B. Percival and A. T. Walden. *Wavelet Methods for Time Series Analysis*. Cambridge University Press, New York, 2000
27. R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, 2003. ISBN 3-900051-00-3, <http://www.R-project.org>
28. A. L. Rukhin and N. Sedransk. Statistics in metrology: international key comparisons and interlaboratory studies. *Journal of Data Science*, 7:393–412, 2007
29. C. M. Wang and H. K. Iyer. Propagation of uncertainties in measurements using generalized inference. *Metrologia*, 42:145–153, 2005
30. C. M. Wang and H. K. Iyer. A generalized confidence interval for a measurand in the presence of Type-A and Type-B uncertainties. *Measurement*, 39:856–863, 2006a
31. C. M. Wang and H. K. Iyer. Uncertainty analysis for vector measurands using fiducial inference. *Metrologia*, 43:486–494, 2006b
32. R. Willink. A procedure for the evaluation of measurement uncertainty based on moments. *Metrologia*, 42:329–343, 2005
33. N. F. Zhang. Calculation of the uncertainty of the mean of autocorrelated measurements. *Metrologia*, 43:S276–S281, 2006

Interval Computations and Interval-Related Statistical Techniques: Tools for Estimating Uncertainty of the Results of Data Processing and Indirect Measurements

Vladik Kreinovich

University of Texas, Department of Computer Science, El Paso, TX 79968-0518,
USA vladik@utep.edu

Summary. In many practical situations, we only know the upper bound Δ on the (absolute value of the) measurement error Δx ; that is, we only know that the measurement error is located on the interval $[-\Delta, \Delta]$. The traditional engineering approach to such situations is to assume that Δx is uniformly distributed on $[-\Delta, \Delta]$, and to use the corresponding statistical techniques. In some situations, however, this approach underestimates the error of indirect measurements. It is therefore desirable to directly process this interval uncertainty. Such “interval computations” methods have been developed since the 1950s. In this chapter, we provide a brief overview of related algorithms, results, and remaining open problems.


Key words: Interval computation, interval-related statistics, interval uncertainty, error bounds, indirect measurements

1 Importance of data processing and indirect measurements

In many real-life situations, we are interested in the value of a physical quantity y that is difficult or impossible to measure directly. Examples of such quantities are the distance to a star and the amount of oil in a given well. Because we cannot measure y directly, a natural idea is to measure y *indirectly*. Specifically, we find some easier-to-measure quantities $x_1, \dots, x_i, \dots, x_n$ that are related to y by a known relation $y = f(x_1, \dots, x_i, \dots, x_n)$; this relation may be a simple functional transformation, or complex algorithm (e.g., for the amount of oil, numerical solution to an inverse problem). Then, to estimate y , we first measure the values of the quantities $x_1, \dots, x_i, \dots, x_n$, and then we use the results $\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n$ of these measurements to compute an estimate \tilde{y} for y as $\tilde{y} = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)$; see Table 1 and Figure 1.

Table 1. List of symbols.

y	actual (unknown) value of the desired quantity
x_i	actual (unknown) value of the i th auxiliary quantity
n	number of auxiliary quantities
$f(x_1, \dots, x_n)$	relation between x_1, \dots, x_n and y : $y = f(x_1, \dots, x_n)$
\tilde{x}_i	result of measuring x_i
\tilde{y}	estimate for y : the result of indirect measurement
$\Delta x_i = \tilde{x}_i - x_i$	the i th measurement error
$\Delta y = \tilde{y} - y$	inaccuracy of indirect measurement
Δ_i	upper bound on the absolute value $ \Delta x_i $ of Δx_i
Δ	resulting upper bound on $ \Delta y $
$x_i = [x_i, \bar{x}_i]$	interval of possible values of x_i
$y = [y, \bar{y}]$	interval of possible values of the desired quantity y
\mathbf{Y}	enclosure for y , i.e., an interval such that $y \subseteq \mathbf{Y}$
$\rho(x)$	probability density function (pdf)
$F(x)$	cumulative distribution function (cdf)
δx_i	simulated value of the i th measurement error
$a^{(k)}$	value of a quantity a on the k th iteration

**Fig. 1.** Indirect measurement.

For example, to find the resistance R , we measure current I and voltage V , and then use the known relation $R = V/I$ to estimate resistance as $\tilde{R} = \tilde{V}/\tilde{I}$.

Computing an estimate for y based on the results of direct measurements is called *data processing*; data processing is the main reason why computers were invented in the first place, and data processing is still one of the main uses of computers as number-crunching devices.

Comment. In this chapter, for simplicity, we consider the case when the relation between x_i and y is known exactly; in some practical situations, we only know an approximate relation between x_i and y .

2 Estimating uncertainty for the results of data processing and indirect measurements: An important metrological problem

Measurements are never 100% accurate, so in reality, the actual value x_i of the i th measured quantity can differ from the measurement result \tilde{x}_i . Because of these *measurement errors* $\Delta x_i \stackrel{\text{def}}{=} \tilde{x}_i - x_i$, the result $\tilde{y} = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)$ of data processing is, in general, different from the actual value $y = f(x_1, \dots, x_i, \dots, x_n)$ of the desired quantity y .

It is desirable to describe the error $\Delta y \stackrel{\text{def}}{=} \tilde{y} - y$ of the result of data processing. To do that, we must have some information about the errors of direct measurements.

3 Uncertainty of direct measurements: Brief description, limitations, need for overall error bounds (i.e., interval uncertainty)

Upper bounds on measurement errors. What do we know about the errors Δx_i of direct measurements? First, the manufacturers of a measuring device usually provide us with an upper bound Δ_i for the (absolute value of) possible measurement errors, that is, with the bound Δ_i for which we are guaranteed that $|\Delta x_i| \leq \Delta_i$.

The need for such a bound comes from the very nature of a measurement process. Indeed, if no such bound is provided, this means that the actual value x_i can be arbitrarily different from the “measurement result” \tilde{x}_i as possible. Such a value \tilde{x}_i is not a measurement; it is a wild guess.

Because the (absolute value of the) measurement error $\Delta x_i = \tilde{x}_i - x_i$ is bounded by the given bound Δ_i , we can therefore guarantee that the actual (unknown) value of the desired quantity belongs to the interval $\mathbf{x}_i \stackrel{\text{def}}{=} [\tilde{x}_i - \Delta_i, \tilde{x}_i + \Delta_i]$.

Example. If the measured value of a quantity is $\tilde{x}_i = 1.0$, and the upper bound Δ_i on the measurement error is 0.1, this means that the (unknown) actual value of the measured quantity can be anywhere between $1 - 0.1 = 0.9$ and $1 + 0.1 = 1.1$, that is, that it can take any value from the interval $[0.9, 1.1]$.

Probabilities. In many practical situations, we not only know the *interval* $[-\Delta_i, \Delta_i]$ of possible values of the measurement error; we also know the *probabilities* of different values Δx_i within this interval (see, e.g., Rabinovich [1]).

In most practical applications, it is assumed that the corresponding measurement errors are normally distributed with 0 mean and known standard deviations.

Numerous engineering techniques are known (and widely used) for processing this uncertainty (see, e.g., Rabinovich [1]).

In practice, we can determine the desired probabilities of different values of Δx_i by comparing the results \tilde{x}_i of measuring with this instrument with the results \tilde{x}_i^{st} of measuring the same quantity by a standard (much more accurate) measuring instrument. Because the standard measuring instrument is much more accurate than the one used, the difference between these two measurement results is practically equal to the measurement error; thus, the empirical distribution of this difference $\tilde{\Delta}x_i = \tilde{x}_i - \tilde{x}_i^{\text{st}}$ is close to the desired probability distribution for the measurement error $\Delta x_i = \tilde{x}_i - x_i$.

Sometimes, one does not know probabilities. There are two cases when this determination is not done.

- First is the case of cutting-edge measurements, for example, measurements in fundamental science. When a Hubble telescope detects the light from a distant galaxy, there is no “standard” (much more accurate) telescope floating nearby that we can use to calibrate the Hubble; the Hubble telescope is the best we have.
- The second case is the case of measurements on the shop floor. In this case, in principle, every sensor can be thoroughly calibrated, but sensor calibration is so costly—usually costing ten times more than the sensor itself—that manufacturers rarely do it.

In both cases, we have no information about the probabilities of Δx_i ; the only information we have is the upper bound on the measurement error.

4 Data processing and indirect measurements under interval uncertainty: The main problem of interval computations

In the case when the only information we have is the upper bound on the measurement error, after we performed a measurement and got a measurement result \tilde{x}_i , the only information that we have about the actual value x_i of the measured quantity is that it belongs to the interval $\mathbf{x}_i = [\tilde{x}_i - \Delta_i, \tilde{x}_i + \Delta_i]$. In such situations, the only information that we have about the (unknown) actual value of $y = f(x_1, \dots, x_i, \dots, x_n)$ is that y belongs to the range $\mathbf{y} = [\underline{y}, \bar{y}]$ of the function f over the box $\mathbf{x}_1 \times \dots \times \mathbf{x}_i \times \dots \times \mathbf{x}_n$:

$$\mathbf{y} = [\underline{y}, \bar{y}] = \{f(x_1, \dots, x_i, \dots, x_n) \mid x_1 \in \mathbf{x}_1, \dots, x_i \in \mathbf{x}_i, \dots, x_n \in \mathbf{x}_n\}. \quad (1)$$

The process of computing this interval range based on the input intervals \mathbf{x}_i is called *interval computations* (Fig.2); see, for example, Jaulin et al. [2] and Kearfott and Reinovich [3].




Fig. 2. Interval computations.

5 Uniform distributions: Traditional engineering approach to interval uncertainty

Brief description. In the case of interval uncertainty, we only know the intervals; we do not know the probability distributions on these intervals. A traditional statistical approach to the situation when several probability distributions are possible is to select the “most uncertain” distribution, that is, the distribution that has the largest possible value of the entropy $S \stackrel{\text{def}}{=} -\int \rho(x) \cdot \ln(\rho(x)) dx$ (here $\rho(x)$ denotes the probability density). For details on this maximum entropy approach and its relation to interval uncertainty (and Laplace’s principle of indifference), see, for example, Jaynes et al. [4].

One can easily check that for a single variable x_1 , among all distributions located on a given interval, the entropy is the largest when this distribution is *uniform* on this interval. Indeed, a function $\rho(x) \geq 0$ is a probability density function on the given interval if $\int \rho(x) dx = 1$. Thus, to find the probability density function that maximizes entropy, we must maximize entropy $-\int \rho(x) \cdot \ln(\rho(x)) dx$ under the constraint $\int \rho(x) dx = 1$. According to the Lagrange multiplier method, for some value λ (Lagrange multiplier), the desired constraint optimization problem is equivalent to an unconstrained optimization problem of maximizing the expression $-\int \rho(x) \cdot \ln(\rho(x)) dx + \lambda \cdot (\int \rho(x) dx - 1)$. Differentiating this expression with respect to each of the variables $\rho(x)$ and equating the derivative to 0, we conclude that $-\ln(\rho(x)) - 1 + \lambda = 0$, hence $\rho(x) = \exp(\lambda - 1)$. The probability density has the same value for all x from the given interval, hence we indeed have a uniform distribution.

In the case of several variables, we can similarly conclude that the distribution with the largest value of the entropy is the one which is uniformly distributed in the corresponding box $x_1 \times \dots \times x_i \times \dots \times x_n$, that is, a distribution in which:

- Each variable Δx_i is uniformly distributed on the corresponding interval $[-\Delta_i, \Delta_i]$.

- Variables corresponding to different inputs are statistically independent.

This is indeed one of the main ways how interval uncertainty is treated in engineering practice; if we only know that the value of some variable is in the interval $[\underline{x}_i, \bar{x}_i]$, and we have no information about the probabilities, then we assume that the variable x_i is uniformly distributed on this interval.

Limitations. To explain the limitations of this engineering approach, let us consider the simplest possible algorithm $y = f(x_1, \dots, x_i, \dots, x_n) = x_1 + \dots + x_i + \dots + x_n$. For simplicity, let us assume that the measured values of all n quantities are 0s $\tilde{x}_1 = \dots = \tilde{x}_i = \dots = \tilde{x}_n = 0$, and that all n measurements have the same error bound Δ_x ; $\Delta_1 = \dots = \Delta x_i = \dots = \Delta_n = \Delta_x$.

In this case, $\Delta y = \Delta x_1 + \dots + \Delta x_i + \dots + \Delta x_n$. Each of n component measurement errors can take any value from $-\Delta_x$ to Δ_x , so the largest possible value of Δy is attained when all of the component errors attain the largest possible value $\Delta x_i = \Delta_x$. In this case, the largest possible value Δ of Δy is equal to $\Delta = n \cdot \Delta_x$.

Let us see what the maximum entropy approach will predict in this case. According to this approach, we assume that Δx_i are independent random variables, each of which is uniformly distributed on the interval $[-\Delta_x, \Delta_x]$. According to the central limit theorem (see, e.g., Sheskin [5]), when $n \rightarrow \infty$, the distribution of the sum of n independent identically distributed bounded random variables tends to Gaussian. This means that for large values n , the distribution of Δy is approximately normal.

A normal distribution is uniquely determined by its mean and variance. When we add several independent variables, their means and variances add up. For each uniform distribution Δx_i on the interval $[-\Delta_x, \Delta_x]$ of width $2\Delta_x$, the probability density is equal to $\rho(x) = 1/(2\Delta_x)$, so the mean is 0 and the variance is

$$V = \int_{-\Delta_x}^{\Delta_x} x^2 \cdot \rho(x) dx = \frac{1}{2\Delta_x} \cdot \int_{-\Delta_x}^{\Delta_x} x^2 dx = \frac{1}{2\Delta_x} \cdot \frac{1}{3} \cdot x^3 \Big|_{-\Delta_x}^{\Delta_x} = \frac{1}{3} \cdot \Delta_x^2. \quad (2)$$

Thus, for the sum Δy of n such variables, the mean is 0, and the variance is equal to $(n/3) \cdot \Delta_x^2$. Thus, the standard deviation is equal to $\sigma = \sqrt{V} = \Delta_x \cdot \sqrt{n}/\sqrt{3}$.

It is known that in a normal distribution, with probability close to 1, all the values are located within the $k \cdot \sigma$ vicinity of the mean: for $k = 3$, it is true with probability 99.9%, for $k = 6$, it is true with probability $1 - 10^{-6}\%$, and so on. So, practically with certainty, Δy is located within an interval $k \cdot \sigma$ which grows with n as \sqrt{n} .

For large n , we have $k \cdot \Delta_x \cdot (\sqrt{n}/\sqrt{3}) \ll \Delta_x \cdot n$, so we get a serious underestimation of the resulting measurement error. This example shows that estimates obtained by selecting a single distribution can be very misleading.

6 Techniques for estimating the uncertainty of the results of indirect measurements in situations when the measurement errors of direct measurements are relatively small

Linearization: main idea. When the measurement errors Δx_i are relatively small, we can use linearization.

By definition of the measurement error $\Delta x_i = \tilde{x}_i - x_i$, hence $x_i = \tilde{x}_i - \Delta x_i$. When the measurement errors Δx_i of direct measurements are relatively small, we can expand the expression

$$\begin{aligned} \Delta y &= \tilde{y} - y = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) - f(x_1, \dots, x_n) \\ &= f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) - f(\tilde{x}_1 - \Delta x_1, \dots, \tilde{x}_i - \Delta x_i, \dots, \tilde{x}_n - \Delta x_n) \end{aligned} \quad (3)$$

in Taylor series and only keep linear terms in the resulting expansion. Because

$$\begin{aligned} y &= f(\tilde{x}_1 - \Delta x_1, \dots, \tilde{x}_i - \Delta x_i, \dots, \tilde{x}_n - \Delta x_n) \\ &\approx f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) - \sum_{i=1}^n \frac{\partial f}{\partial x_i} \cdot \Delta x_i, \end{aligned} \quad (4)$$

we conclude that $\Delta y = \tilde{y} - y = \sum_{i=1}^n c_i \cdot \Delta x_i$, where $c_i = \partial f / \partial x_i$.

The dependence of Δy on Δx_i is linear; it is increasing relative to x_i if $c_i \geq 0$ and decreasing if $c_i < 0$. So, to find the largest possible value Δ of Δy , we must take:

- The largest possible value $\Delta x_i = \Delta_i$ when $c_i \geq 0$
- The smallest possible value $\Delta x_i = -\Delta_i$ when $c_i < 0$

In both cases, the corresponding term in the sum has the form $|c_i| \cdot \Delta_i$, so we can conclude that

$$\Delta = \sum_{i=1}^n |c_i| \cdot \Delta_i. \quad (5)$$

Similarly, the smallest possible value of Δy is equal to $-\Delta$. Thus, the range of possible values of y is equal to $[y, \bar{y}] = [\tilde{y} - \Delta, \tilde{y} + \Delta]$. So, to compute Δ , it is sufficient to know the partial derivatives c_i .

Case of analytical formulas. In the simplest case when the algorithm $f(x_1, \dots, x_i, \dots, x_n)$ consists of a simple analytical expression, we can find explicit analytical formulas for the partial derivatives and thus compute the desired bound Δ .

Techniques based on sensitivity analysis (automatic differentiation). In the general case, a natural way to compute partial derivatives comes directly from the definition. By definition, a partial derivative is defined as a limit

$$\frac{\partial f}{\partial x_i} = \lim_{h_i \rightarrow 0} \frac{f(\tilde{x}_1, \dots, \tilde{x}_{i-1}, \tilde{x}_i + h_i, \tilde{x}_{i+1}, \dots, \tilde{x}_n) - f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)}{h_i}. \quad (6)$$

In turn, a limit, by its definition, means that when the values of h_i are small, the corresponding ratio is very close to the partial derivative. Thus, we can estimate the partial derivative as the ratio

$$c_i = \frac{\partial f}{\partial x_i} \approx \frac{f(\tilde{x}_1, \dots, \tilde{x}_{i-1}, \tilde{x}_i + h_i, \tilde{x}_{i+1}, \dots, \tilde{x}_n) - f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)}{h_i} \quad (7)$$

for some small value h_i .

After we have computed n such ratios, we can then compute the desired bound Δ on $|\Delta y|$ as $\Delta = |c_i| \cdot \Delta_i$.

In general, this procedure requires n divisions by h_i and n multiplications by Δ_i . The procedure can be made faster if we select $h_i = \Delta_i$. In this case, we get

$$\Delta = \sum_{i=1}^n |f(\tilde{x}_1, \dots, \tilde{x}_{i-1}, \tilde{x}_i + \Delta_i, \tilde{x}_{i+1}, \dots, \tilde{x}_n) - \tilde{y}|. \quad (8)$$

Advanced Monte Carlo simulation techniques. The above algorithm requires that we call the data processing algorithm $n + 1$ times: first to compute the value $\tilde{y} = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)$, and then n more times to compute the values

$$f(\tilde{x}_1, \dots, \tilde{x}_{i-1}, \tilde{x}_i + \Delta_i, \tilde{x}_{i+1}, \dots, \tilde{x}_n) \quad (9)$$

and thus, the corresponding partial derivatives.

In many practical situations, the data processing algorithms are time-consuming, and we process large amounts of data, with the number n of data points in thousands. In this case, the use of the above linearization algorithm would be thousands of times longer than data processing itself, which itself is already time-consuming. Is it possible to estimate Δ faster?

The answer is “yes;” it is possible to have a Monte Carlo type algorithm that estimates Δ by using only a constant number of calls to the data processing algorithm f ; for details, see, for example, Kreinovich et al. [6] and Kreinovich and Ferson [7].

At first glance, because we know that the measurement errors are located within the intervals $[-\Delta_i, \Delta_i]$, it sounds reasonable to select distributions located on these intervals. However, it can be shown that this does not lead to the desired estimates. It turns out that it is possible to estimate the interval uncertainty if we use a distribution d which is *not* located on the interval $[-\Delta_i, \Delta_i]$, namely, a distribution related to the *basic Cauchy* distribution with the probability density function

$$\rho(x) = \frac{1}{\pi \cdot (x^2 + 1)}.$$

The resulting Cauchy deviate method works in the linearized case, when the function $f(x_1, \dots, x_i, \dots, x_n)$ is reasonably smooth and the box $[\underline{x}_1, \bar{x}_1] \times \dots \times [\underline{x}_i, \bar{x}_i] \times \dots \times [\underline{x}_n, \bar{x}_n]$ is small enough, so that on this box, we can reasonably approximate the function f by its linear terms.

If we multiply a random variable distributed according to the above basic Cauchy distribution d by a value Δ , then we get a Cauchy distribution with a parameter Δ , that is, a distribution described by the following density function,

$$\rho(x) = \frac{\Delta}{\pi \cdot (x^2 + \Delta^2)}.$$

It is known that if $\xi_1, \dots, \xi_i, \dots, \xi_n$ are independent variables distributed according to Cauchy distributions with parameters Δ_i , then, for every n real numbers $c_1, \dots, c_i, \dots, c_n$, the corresponding linear combination $c_1 \cdot \xi_1 + \dots + c_i \cdot \xi_i + \dots + c_n \cdot \xi_n$ is also Cauchy distributed, with the parameter Δ equal to the desired value $\Delta = |c_1| \cdot \Delta_1 + \dots + |c_i| \cdot \Delta_i + |c_n| \cdot \Delta_n$.

Thus, if for some number of iterations N , we simulate $\delta x_i^{(k)}$ ($1 \leq k \leq N$) as Cauchy distributed with parameter Δ_i , then, in the linear approximation, the corresponding differences

$$\delta y^{(k)} \stackrel{\text{def}}{=} f(\tilde{x}_1 + \delta x_1^{(k)}, \dots, \tilde{x}_i + \delta x_i^{(k)}, \dots, \tilde{x}_n + \delta x_n^{(k)}) - \tilde{y} \quad (10)$$

are distributed according to the Cauchy distribution with the parameter Δ . The resulting values $\delta y^{(1)}, \dots, \delta y^{(k)}, \dots, \delta y^{(N)}$ are therefore a sample from the Cauchy distribution with the unknown parameter Δ . Based on this sample, we can estimate the value Δ .

In order to estimate Δ , we can apply the maximum likelihood method which leads to the following equation,

$$\frac{1}{1 + \left(\frac{\delta y^{(1)}}{\Delta}\right)^2} + \dots + \frac{1}{1 + \left(\frac{\delta y^{(k)}}{\Delta}\right)^2} + \dots + \frac{1}{1 + \left(\frac{\delta y^{(N)}}{\Delta}\right)^2} = \frac{N}{2}. \quad (11)$$

The left-hand side of this equation is an increasing function that is equal to 0 (hence smaller than $N/2$) for $\Delta = 0$ and larger than $N/2$ for $\Delta = \max |\delta y^{(k)}|$; therefore the solution to this equation can be found by applying a bisection method to the interval $[0, \max |\delta y^{(k)}|]$.

Simulation of the Cauchy distribution with parameter Δ_i can be based on the functional transformation of uniformly distributed sample values:

$$\delta x_i^{(k)} = \Delta_i \cdot \tan(\pi \cdot (r_i - 0.5)), \quad (12)$$

where r_i is uniformly distributed on the interval $[0, 1]$.

As a result, we arrive at the following algorithm (see, e.g., Kreinovich and Ferson [7] and Trejo et al. [8]).

- Apply f to the midpoints: $\tilde{y} := f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n)$.
- For $k = 1, 2, \dots, N$, repeat the following.
 - Use the standard random number generator to compute n numbers $r_i^{(k)}$, $i = 1, 2, \dots, n$, that are uniformly distributed on the interval $[0, 1]$.
 - Compute Cauchy distributed values $c_i^{(k)} := \tan(\pi \cdot (r_i^{(k)} - 0.5))$.
 - Compute the largest value of $|c_i^{(k)}|$ so that we will be able to normalize the simulated approximation errors and apply f to the values that are within the box of possible values: $K := \max_i |c_i^{(k)}|$.
 - Compute the simulated approximation errors $\delta x_i^{(k)} := \Delta_i \cdot c_i^{(k)} / K$.
 - Compute the simulated “actual values” $x_i^{(k)} := \tilde{x}_i + \delta x_i^{(k)}$.
 - Apply the program f to the simulated measurement results and compute the simulated approximation error for y :

$$\Delta y^{(k)} := K \cdot \left(f \left(x_1^{(k)}, \dots, x_i^{(k)}, \dots, x_n^{(k)} \right) - \tilde{y} \right). \quad (13)$$

- Compute Δ by applying the bisection method to solve the equation

$$\frac{1}{1 + \left(\frac{\Delta y^{(1)}}{\Delta} \right)^2} + \dots + \frac{1}{1 + \left(\frac{\Delta y^{(k)}}{\Delta} \right)^2} + \dots + \frac{1}{1 + \left(\frac{\Delta y^{(N)}}{\Delta} \right)^2} = \frac{N}{2}. \quad (14)$$

In Kreinovich and Ferson [7] and Trejo et al. [8], we found the number of iterations N that would provide the desired (relative) accuracy ε in estimating Δ , that is, the number of iterations that are needed to guarantee that

$$(1 - \varepsilon) \cdot \tilde{\Delta} \leq \Delta \leq (1 + \varepsilon) \cdot \tilde{\Delta} \quad (15)$$

with a given certainty p_0 .

In practice, it is reasonable to get a certainty $p_0 = 95\%$ and accuracy $\varepsilon = 0.2$ (20%).

To get an accuracy ε with 95% certainty, we must pick $N = 8/\varepsilon^2$. In particular, to get a 20% accuracy ($0.2 \cdot \Delta$) with 95% certainty, that is, to guarantee that

$$0.8 \cdot \tilde{\Delta} \leq \Delta \leq 1.2 \cdot \tilde{\Delta} \quad (16)$$

with certainty $\geq 95\%$, we need $N = 8/(0.2)^2 = 200$ runs.

In general, the required number of calls to a model depends only on the desired accuracy ε and not on n , so for large n , these methods are much faster.

Comment. It is important to mention that we assumed that the function f is reasonably linear within the box

$$[\tilde{x}_1 - \Delta_1, \tilde{x}_1 + \Delta_1] \times \dots \times [\tilde{x}_i - \Delta_i, \tilde{x}_i + \Delta_i] \times \dots \times [\tilde{x}_n - \Delta_n, \tilde{x}_n + \Delta_n]. \quad (17)$$

However, the simulated values δ_i may be outside the box. When we get such values, we do not use the function f for them, we use a linearized function that is equal to f within the box, and that is extended linearly for all other values.

7 Techniques for error estimation in the general case of interval uncertainty

Need for interval computations. In many application areas, it is sufficient to have an approximate estimate of y . However, in some applications, it is important to guarantee that the (unknown) actual value y of a certain quantity does not exceed a certain threshold y_0 . The only way to guarantee this is to have an interval $\mathbf{Y} = [\underline{Y}, \bar{Y}]$ which is guaranteed to contain y (i.e., for which $\mathbf{y} \subseteq \mathbf{Y}$) and for which $\bar{Y} \leq y_0$.

For example, in nuclear engineering, we must make sure that the temperatures and the neutron flows do not exceed the critical values; when planning a space flight, we want to guarantee that the spaceship lands on the planet and does not fly past it, and so on.

The interval \mathbf{Y} that is guaranteed to contain the actual range \mathbf{y} is usually called an *enclosure* for this range. So, in such situations, we need to compute either the original range or at least an enclosure for this range. Computing such an enclosure is also one of the main tasks of interval computations.

Traditional numerical methods are often not sufficient. The main limitations of the traditional numerical mathematics approach to error estimation was that often no clear distinction was made between approximate (non-guaranteed) and guaranteed (= interval) error bounds.

For example, for iterative methods, many papers on numerical mathematics consider the rate of convergence as an appropriate measure of approximation error. Clearly, if we know that the error decreases as $O(1/n)$ or as $O(a^{-n})$, we gain some information about the corresponding algorithms, and we also gain knowledge that for large n , the second method is more accurate. However, in real life, we make a fixed number n of iterations. If the only information we have about the approximation error is the above asymptotics, then we still have no idea how close the result of n th iteration is to the actual (desired) value.

It is therefore important to emphasize the need for guaranteed methods, and to develop techniques for producing *guaranteed* estimates. Such guaranteed estimates are what interval computations are about.

Interval computations: A brief history. The notion of interval computations is reasonably recent, it dates back to the 1950s, but the main problem has been known since Archimedes who used guaranteed two-sided bounds to compute π (see, e.g., Archimedes [9]).

Since then, many useful guaranteed bounds have been developed for different numerical methods. There have also been several general descriptions of such bounds, often formulated in terms similar to what we described above. For example, in the early 20th century, the concept of a function having values that are bounded within limits was discussed by W. H. Young in [10]. The concept of operations with a set of multivalued numbers was introduced by R. C. Young, who developed a formal algebra of multivalued numbers [11]. The special case of closed intervals was further developed by P. S. Dwyer in [12].

Interval computations in their current form were independently invented by three researchers in three different parts of the world: by M. Warmus in Poland [13], by T. Sunaga in Japan [14], and by R. Moore in the United States [15].

The active interest in interval computations started with Moore's 1966 monograph [16]. This interest was enhanced by the fact that in addition to estimates for general numerical algorithms, Moore's monograph also described *practical* applications that had already been developed in his earlier papers and technical reports: in particular, interval computations were used to make sure that even when we take all the uncertainties into account, the trajectory of a space flight is guaranteed to reach the moon.

Since then, interval computations have been actively used in many areas of science and engineering; see, for example, interval website [17] and Jaulin et al. [2].

Comment. Early papers on interval computations can be found on the interval computations website [17].

First step: Interval arithmetic. Our goal is to find the range of a given function $f(x_1, \dots, x_i, \dots, x_n)$ on the given intervals $\mathbf{x}_1 = [\underline{x}_1, \bar{x}_1], \dots, \mathbf{x}_i = [\underline{x}_i, \bar{x}_i], \dots, \mathbf{x}_n = [\underline{x}_n, \bar{x}_n]$.

This function $f(x_1, \dots, x_i, \dots, x_n)$ is given as an algorithm. In particular, we may have an explicit analytical expression for f , in which case this algorithm consists of simply computing this expression.

When we talk about algorithms, we usually mean an algorithm (program) written in a high-level programming language like Java or C. Such programming languages allow us to use arithmetic expressions and many other complex constructions. Most of these constructions, however, are not directly hardware-supported inside a computer. Usually, only simple arithmetic operations are implemented: addition, subtraction, multiplication, and $1/x$ (plus branching). Even division a/b is usually not directly supported, it is performed as a sequence of two elementary arithmetic operations:

- First, we compute $1/b$.
- Then, we multiply a by $1/b$.

When we input a general program into a computer, the computer *parses* it, that is, represents it a sequence of elementary arithmetic operations.

Because a computer performs this parsing anyway, we can safely assume that the original algorithm $f(x_1, \dots, x_i, \dots, x_n)$ is already represented as a sequence of such elementary arithmetic operations.

Let us start our analysis of the interval computation techniques with the simplest possible case when the algorithm $f(x_1, \dots, x_i, \dots, x_n)$ simply consists of a single arithmetic operation: addition, subtraction, multiplication, or computing $1/x$.

Let us start by estimating the range of the addition function $f(x_1, x_2) = x_1 + x_2$ on the intervals $[\underline{x}_1, \bar{x}_1]$ and $[\underline{x}_2, \bar{x}_2]$. This function is increasing with respect to both its variables. We already know how to compute the range $[\underline{y}, \bar{y}]$ of a monotonic function. So, the range of addition is equal to $[\underline{x}_1 + \underline{x}_2, \bar{x}_1 + \bar{x}_2]$.

The desired range is usually denoted as $f(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n)$; in particular, for addition, this notation takes the form $\mathbf{x}_1 + \mathbf{x}_2$. Thus, we can define “addition” of two intervals as follows,

$$[\underline{x}_1, \bar{x}_1] + [\underline{x}_2, \bar{x}_2] = [\underline{x}_1 + \underline{x}_2, \bar{x}_1 + \bar{x}_2]. \quad (18)$$

This formula makes perfect intuitive sense: if one town has between 700 and 800 thousand people, and it merges with a nearby town whose population is between 100 and 200 thousand, then:

- The smallest possible value of the total population of the new big town is when both populations are the smallest possible, $700 + 100 = 800$, and
- The largest possible value is when both populations are the largest possible, that is, $800 + 200 = 1000$.

The subtraction function $f(x_1, x_2) = x_1 - x_2$ is increasing with respect to x_1 and decreasing with respect to x_2 , so we have

$$[\underline{x}_1, \bar{x}_1] - [\underline{x}_2, \bar{x}_2] = [\underline{x}_1 - \bar{x}_2, \bar{x}_1 - \underline{x}_2]. \quad (19)$$

These operations are also in full agreement with common sense. For example, if a warehouse originally had between 6.0 and 8.0 tons, and we moved between 1.0 and 2.0 tons to another location, then the smallest amount left is when we start with the smallest possible value 6.0 and move the largest possible value 2.0, resulting in $6.0 - 2.0 = 4.0$. The largest amount left is when we start with the largest possible value 8.0 and move the smallest possible value 1.0, resulting in $8.0 - 1.0 = 7.0$.

For multiplication $f(x_1, x_2) = x_1 \cdot x_2$, the direction of monotonicity depends on the actual values of x_1 and x_2 : for example, when $x_2 > 0$, the product increases with x_1 ; otherwise it decreases with x_1 . So, unless we know the signs of the product beforehand, we cannot tell whether the maximum is attained at $x_1 = \underline{x}_1$ or at $x_1 = \bar{x}_1$. However, we know that it is always attained at one of these endpoints. So, to find the range of the product, it is sufficient to try all $2 \cdot 2 = 4$ combinations of these endpoints

$$\begin{aligned}
 & [\underline{x}_1, \bar{x}_1] \cdot [\underline{x}_2, \bar{x}_2] \\
 = & [\min(\underline{x}_1 \cdot \underline{x}_2, \underline{x}_1 \cdot \bar{x}_2, \bar{x}_1 \cdot \underline{x}_2, \bar{x}_1 \cdot \bar{x}_2), \max(\underline{x}_1 \cdot \underline{x}_2, \underline{x}_1 \cdot \bar{x}_2, \bar{x}_1 \cdot \underline{x}_2, \bar{x}_1 \cdot \bar{x}_2)].
 \end{aligned}
 \tag{20}$$

Finally, the function $f(x_1) = 1/x_1$ is decreasing wherever it is defined (when $x_1 \neq 0$), so if $0 \notin [\underline{x}_1, \bar{x}_1]$; then

$$\frac{1}{[\underline{x}_1, \bar{x}_1]} = \left[\frac{1}{\bar{x}_1}, \frac{1}{\underline{x}_1} \right].
 \tag{21}$$

The formulas for addition, subtraction, multiplication, and reciprocal of intervals are called formulas of *interval arithmetic*.

Comment. Alternative faster-to-compute formulas that lead to slightly wider intervals are described, for example, in Cerimele and Venturini Zilli [18].

Straightforward (“naive”) interval computations. Historically the first method for computing the enclosure for the general case is the method which is sometimes called “straightforward” interval computations. In this method, we repeat the computations forming the program f step by step, replacing each operation with real numbers by the corresponding operation of interval arithmetic. It is known that, as a result, we get an enclosure $\mathbf{Y} \supseteq \mathbf{y}$ for the desired range.

In some cases, this enclosure is exact. In more complex cases (see example below), the enclosure has excess width.

Example. Let us illustrate the above idea on the example of estimating the range of the function $f(x_1) = x_1 - x_1^2$ on the interval $x_1 \in [0, 0.8]$.

We start with parsing the expression for the function (i.e., describing how a computer will compute this expression); it will implement the following sequence of elementary operations,

$$r_1 := x_1 \cdot x_1; \quad r_2 := x_1 - r_1.
 \tag{22}$$

According to straightforward interval computations, we perform the same operations, but with *intervals* instead of *numbers*:

$$\mathbf{r}_1 := [0, 0.8] \cdot [0, 0.8] = [0, 0.64]; \quad \mathbf{r}_2 := [0, 0.8] - [0, 0.64] = [-0.64, 0.8].
 \tag{23}$$

For this function, the actual range is $f(\mathbf{x}_1) = [0, 0.25]$; see Figure 3.

Interval computations go beyond straightforward technique. People who are vaguely familiar with interval computations sometimes erroneously assume that the above straightforward (“naive”) techniques are all there is in interval computations. In conference presentations (and even in published papers), one often encounters a statement, “I tried interval computations, and it did




Fig. 3. Range of the function $f(x_1) = x_1 - x_1^2$ on the interval $[0, 0.8]$.

not work.” What this statement usually means is that they tried the above straightforward approach and—not surprisingly—it did not work well.

In reality, interval computation is *not a single algorithm*, it is a *problem* for which many different techniques exist. Let us now describe some of such techniques.

Comment. For each of the known techniques, there are cases when we get an excess width. The reason is that the problem of computing the exact range is NP-hard even for polynomial functions $f(x_1, \dots, x_i, \dots, x_n)$, actually, even for quadratic functions f (see, e.g., Kreinovich et al. [19]).

Centered form. One such technique is the centered form (see, e.g., Jaulin et al. [2]). This technique is based on the same Taylor series expansion ideas as linearization. We start by representing each interval $\mathbf{x}_i = [\underline{x}_i, \bar{x}_i]$ in the form $[\tilde{x}_i - \Delta_i, \tilde{x}_i + \Delta_i]$, where $\tilde{x}_i = (\underline{x}_i + \bar{x}_i)/2$ is the midpoint of the interval \mathbf{x}_i and $\Delta_i = (\bar{x}_i - \underline{x}_i)/2$ is the half-width of this interval.

After that, we use the Taylor expansion. In linearization, we simply ignored quadratic and higher-order terms. Here, instead, we use the Taylor form with a remainder term. Specifically, the centered form is based on the formula

$$f(x_1, \dots, x_i, \dots, x_n) = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\eta_1, \dots, \eta_i, \dots, \eta_n) \cdot (x_i - \tilde{x}_i), \quad (24)$$

where each η_i is some value from the interval \mathbf{x}_i .

Because $\eta_i \in \mathbf{x}_i$, the value of the i th derivative belongs to the interval range of this derivative on these intervals. We also know that $x_i - \tilde{x}_i \in [-\Delta_i, \Delta_i]$. Thus, we can conclude that

$$f(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n) \subseteq f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n) \cdot [-\Delta_i, \Delta_i]. \quad (25)$$

To compute the ranges of the partial derivatives, we can use straightforward interval computations.

Example. Let us illustrate this method on the above example of estimating the range of the function $f(x_1) = x_1 - x_1^2$ over the interval $[0, 0.8]$. For this interval, the midpoint is $\tilde{x}_1 = 0.4$; at this midpoint, $f(\tilde{x}_1) = 0.24$. The half-width is $\Delta_1 = 0.4$. The only partial derivative here is $\partial f / \partial x_1 = 1 - 2x_1$, its range on $[0, 0.8]$ is equal to $1 - 2 \cdot [0, 0.8] = [-0.6, 1]$. Thus, we get the following enclosure for the desired range \mathbf{y} ,

$$\mathbf{y} \subseteq \mathbf{Y} = 0.24 + [-0.6, 1] \cdot [-0.4, 0.4] = 0.24 + [-0.4, 0.4] = [-0.16, 0.64]. \quad (26)$$

This enclosure is narrower than the “naive” estimate $[-0.64, 0.8]$, but it still contains excess width.

How can we get better estimates? In the centered form, we, in effect, ignored quadratic and higher-order terms, that is, terms of the type $(\partial^2 f / \partial x_i \partial x_j) \cdot \Delta x_i \cdot \Delta x_j$. When the estimate is not accurate enough, it means that this ignored term is too large. There are two ways to reduce the size of the ignored term:

- We can try to decrease this quadratic term.
- We can try to explicitly include higher-order terms in the Taylor expansion formula, so that the remainder term will be proportional to say Δx_i^3 and thus, be much smaller.

Let us describe these two ideas in detail.

First idea: Bisection. Let us first describe the situation in which we try to minimize the second-order remainder term. In the above expression for this term, we cannot change the second derivative. The only thing we can decrease is the difference $\Delta x_i = x_i - \tilde{x}_i$ between the actual value and the midpoint. This value is bounded by the half-width Δ_i of the box. So, to decrease this value, we can subdivide the original box into several narrower subboxes. Usually, we divide into two subboxes, so this subdivision is called *bisection*.

The range over the whole box is equal to the union of the ranges over all the subboxes. The widths of each subbox are smaller, so we get smaller Δx_i and it is hoped, more accurate estimates for ranges over each subbox. Then, we take the union of the ranges over subboxes.

Example. Let us illustrate this idea on the above $x_1 - x_1^2$ example. In this example, we divide the original interval $[0, 0.8]$ into two subintervals $[0, 0.4]$ and $[0.4, 0.8]$. For both intervals, $\Delta_1 = 0.2$.

In the first subinterval, the midpoint is $\tilde{x}_1 = 0.2$, so $f(\tilde{x}_1) = 0.2 - 0.04 = 0.16$. The range of the derivative is equal to $1 - 2 \cdot [0, 0.4] = 1 - [0, 0.8] = [0.2, 1]$, hence we get an enclosure $0.16 + [0.2, 1] \cdot [-0.2, 0.2] = [-0.04, 0.36]$.

For the second interval, $\tilde{x}_1 = 0.6$, $f(0.6) = 0.24$, the range of the derivative is $1 - 2 \cdot [0.4, 0.8] = [-0.6, 0.2]$, hence we get an enclosure

$$0.24 + [-0.6, 0.2] \cdot [-0.2, 0.2] = [0.12, 0.36]. \quad (27)$$

The union of these two enclosures is the interval $[-0.04, 0.36]$. This enclosure is much more accurate than before.

Further bisection leads to even more accurate estimates: the smaller the subintervals, the more accurate is the enclosure.

Bisection: General comment. The more subboxes we consider, the smaller Δx_i and thus, the more accurate are the corresponding enclosures. However, once we have more boxes, we need to spend more time processing these boxes. Thus, we have a trade-off between computation time and accuracy: the more computation time we allow, the more accurate estimates we will be able to compute.

Additional idea: Monotonicity checking. If the function $f(x_1, \dots, x_i, \dots, x_n)$ is monotonic over the original box $\mathbf{x}_1 \times \dots \times \mathbf{x}_i \times \dots \times \mathbf{x}_n$, then we can easily compute its exact range. Because we used the centered form for the original box, this probably means that on that box, the function is not monotonic: for example, with respect to x_1 , it may be increasing at some points in this box, and decreasing at other points.

However, as we divide the original box into smaller subboxes, it is quite possible that at least some of these subboxes will be outside the areas where the derivatives are 0 and thus, the function $f(x_1, \dots, x_i, \dots, x_n)$ will be monotonic. So, after we subdivide the box into subboxes, we should first check monotonicity on each of these subboxes, and if the function is monotonic, we can easily compute its range.

In calculus terms; a function is increasing with respect to x_i if its partial derivative $d_i \stackrel{\text{def}}{=} \partial f / \partial x_i$ is nonnegative everywhere on this subbox. Thus, to check monotonicity, we should find the range $[\underline{d}_i, \bar{d}_i]$ of this derivative (we need to do it anyway to compute the centered form expression):

- If $\underline{d}_i \geq 0$, this means that the derivative is everywhere nonnegative and thus, the function f is increasing in x_i
- If $\bar{d}_i \leq 0$, this means that the derivative is everywhere nonpositive and thus, the function f is decreasing in x_i .

If $\underline{d}_i < 0 < \bar{d}_i$, then we have to use the centered form.

If the function is monotonic (e.g., increasing) only with respect to some of the variables x_i , then

- To compute \bar{y} , it is sufficient to consider only the value $x_i = \bar{x}_i$.
- To compute \underline{y} , it is sufficient to consider only the value $x_i = \underline{x}_i$.

For such subboxes, we reduce the original problem to two problems with fewer variables, problems which are thus easier to solve.

Example. For the example $f(x_1) = x_1 - x_1^2$, the partial derivative is equal to $1 - 2 \cdot x_1$.

On the first subbox $[0, 0.4]$, the range of this derivative is $1 - 2 \cdot [0, 0.4] = [0.2, 1]$. Thus, the derivative is always nonnegative, the function is increasing on this subbox, and its range on this subbox is equal to $[f(0), f(0.4)] = [0, 0.16]$.

On the second subbox $[0.4, 0.8]$, the range of the derivative is $1 - 2 \cdot [0.4, 0.8] = [-0.6, 0.2]$. Here, we do not have guaranteed monotonicity, so we can use the centered form to get the enclosure $[0.12, 0.36]$ for the range.

The union of these two enclosures is the interval $[0, 0.36]$, which is slightly more accurate than before. Further bisection leads to even more accurate estimates.

Comment. We got the exact range because of the simplicity of our example, in which the extreme point 0.5 of the function $f(x_1) = x_1 - x_1^2$ is exactly in the middle of the interval $[0, 1]$. Thus, when we divided the box in two, both subboxes have the monotonicity property. In the general case, the extremal point will be inside one of the subboxes, so we will have excess width.

General Taylor techniques. As we have mentioned, another way to get more accurate estimates is to use so-called *Taylor techniques*, that is, to explicitly consider second-order and higher-order terms in the Taylor expansion (see, e.g., Berz and Hoffstätter [20], Neumaier [21], and references therein).

Let us illustrate the main ideas of Taylor analysis on the case when we allow second-order terms. In this case, the formula with a remainder takes the form

$$f(x_1, \dots, x_i, \dots, x_n) = f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\tilde{x}_1, \dots, \tilde{x}_n) \cdot (x_i - \tilde{x}_i) + \frac{1}{2} \cdot \sum_{i=1}^n \sum_{j=1}^m \frac{\partial^2 f}{\partial x_i \partial x_j}(\eta_1, \dots, \eta_m) \cdot (x_i - \tilde{x}_i) \cdot (x_j - \tilde{x}_j). \quad (28)$$

Thus, we get the enclosure

$$f(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n) \subseteq \mathbf{Y} \stackrel{\text{def}}{=} f(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\tilde{x}_1, \dots, \tilde{x}_i, \dots, \tilde{x}_n) \cdot [-\Delta_i, \Delta_i] + \frac{1}{2} \cdot \sum_{i=1}^n \sum_{j=1}^m \frac{\partial^2 f}{\partial x_i \partial x_j}(\mathbf{x}_1, \dots, \mathbf{x}_n) \cdot [-\Delta_i, \Delta_i] \cdot [-\Delta_j, \Delta_j]. \quad (29)$$

Example. Let us illustrate this idea on the above example of $f(x_1) = x_1 - x_1^2$. Here, $\Delta_1 = 0.4$, $\tilde{x}_1 = 0.4$, so $f(\tilde{x}_1) = 0.24$ and $(\partial f / \partial x_1)(\tilde{x}_1) = 1 - 2 \cdot 0.4 = 0.2$. The second derivative is equal to -2 , so the Taylor estimate takes the form

$$\mathbf{Y} = 0.24 + 0.2 \cdot [-0.4, 0.4] - [-0.4, 0.4]^2.$$

Strictly speaking, if we interpret Δx_1^2 as $\Delta x_1 \cdot \Delta x_1$ and use the formulas of interval multiplication, we get the interval

$$[-0.4, 0.4]^2 = [-0.4, 0.4] \cdot [-0.4, 0.4] = [-0.16, 0.16]$$

and thus, the enclosure

$$\mathbf{Y} = 0.24 + [-0.08, 0.08] - [-0.16, 0.16] = [0.16, 0.32] - [-0.16, 0.16] = [0, 0.48]$$

for the desired range. However, we can view x^2 as a special function, for which the range over $[-0.4, 0.4]$ is known to be $[0, 0.16]$. In this case, the above enclosure takes the form

$$\mathbf{Y} = 0.24 + [-0.08, 0.08] - [0, 0.16] = [0.16, 0.32] - [0, 0.16] = [0, 0.32]$$

which is much closer to the actual range $[0, 0.25]$.

Taylor methods: General comment. The more terms we consider in the Taylor expansion, the smaller the remainder term and thus, the more accurate the corresponding enclosures. However, once we have more terms, we need to spend more time computing these terms. Thus, for Taylor methods, we also have a trade-off between computation time and accuracy: the more computation time we allow, the more accurate estimates we will be able to compute.

An alternative version of affine and Taylor arithmetic. The main idea of Taylor methods is to approximate the given function $f(x_1, \dots, x_i, \dots, x_n)$ by a polynomial of a small order plus an interval remainder term.

In these terms, straightforward interval computations can be viewed as 0th-order Taylor methods in which all we have is the corresponding interval (or, equivalently, the constant term plus the remainder interval). To compute this interval, we repeated the computation of f step by step, replacing operations with numbers by operations with intervals.

We can do the same for higher-order Taylor expansions as well. Let us illustrate how this can be done for the first-order Taylor terms. We start with the expressions $x_i = \tilde{x}_i - \Delta x_i$. Then, at each step, we keep a term of the type $a = \tilde{a} + \Delta x_i + \mathbf{a}$. (To be more precise, keep the coefficients \tilde{a} and a_i and the interval \mathbf{a} .)

Addition and subtraction of such terms are straightforward:

$$\left(\tilde{a} + \sum_{i=1}^n a_i \cdot \Delta x_i + \mathbf{a}\right) + \left(\tilde{b} + \sum_{i=1}^n b_i \cdot \Delta x_i + \mathbf{b}\right) = (\tilde{a} + \tilde{b}) + \sum_{i=1}^n (a_i + b_i) \cdot \Delta x_i + (\mathbf{a} + \mathbf{b}); \quad (30)$$

$$\left(\tilde{a} + \sum_{i=1}^n a_i \cdot \Delta x_i + \mathbf{a}\right) - \left(\tilde{b} + \sum_{i=1}^n b_i \cdot \Delta x_i + \mathbf{b}\right) = (\tilde{a} - \tilde{b}) + \sum_{i=1}^n (a_i - b_i) \cdot \Delta x_i + (\mathbf{a} - \mathbf{b}). \quad (31)$$

For multiplication, we add terms proportional to $\Delta x_i \cdot \Delta x_j$ to the interval part

$$\begin{aligned}
& (\tilde{\mathbf{a}} + \sum_{i=1}^n a_i \cdot \Delta x_i + \mathbf{a}) \cdot (\tilde{\mathbf{b}} + \sum_{i=1}^n b_i \cdot \Delta x_i + \mathbf{b}) = (\tilde{\mathbf{a}} \cdot \tilde{\mathbf{b}}) + \sum_{i=1}^n (\tilde{\mathbf{a}} \cdot b_i + \tilde{\mathbf{b}} \cdot a_i) \cdot \Delta x_i \\
& + (\tilde{\mathbf{a}} \cdot \mathbf{b} + \tilde{\mathbf{b}} \cdot \mathbf{a} + \sum_{i=1}^n a_i \cdot b_i \cdot [0, \Delta_i^2] + \sum_{i=1}^n \sum_{j \neq i} a_i \cdot b_j \cdot [-\Delta_i, \Delta_i] \cdot [\Delta_j \cdot \Delta_j]).
\end{aligned} \tag{32}$$

At the end, we get an expression of the above type for the desired quantity y : $y = \tilde{y} + \Delta x_i + \mathbf{y}$. We already know how to compute the range of a linear function, so we get the following enclosure for the final range, $\mathbf{Y} = \tilde{y} + [-\Delta, \Delta] + \mathbf{y}$, where $\Delta = |y_i| \cdot \Delta_i$.

Example. For $f(x_1) = x_1 - x_1^2$, we first compute $x_2 = x_1^2$ and then $y = x_1 - x_2$. We start with the interval $\mathbf{x}_1 = \tilde{x}_1 - \Delta x_1 = 0.4 + (-1) \cdot \Delta_1 + [0, 0]$.

On the next step, we compute the square of this expression. This square is equal to $0.16 + (-0.8) \cdot \Delta x_1 + \Delta x_1^2$. Because $\Delta x_1 \in [-0.4, 0.4]$, we conclude that $\Delta x_1^2 \in [0, 0.16]$ and thus, that $x_2 = 0.16 + (-0.8) \cdot \Delta x_1 + [0, 0.16]$.

For $y = x_1 - x_2$, we now have

$$\begin{aligned}
y &= (0.4 - 0.16) + ((-1) - (-0.8)) \cdot \Delta x_1 + ([0, 0] - [0, 0.16]) \\
&= 0.24 + (-0.2) \cdot \Delta x_1 + [-0.16, 0].
\end{aligned} \tag{33}$$

Because $\Delta x_1 \in [-0.4, 0.4]$, we get the enclosure

$$\mathbf{Y} = 0.24 + (-0.2) \cdot [-0.4, 0.4] + [-0.16, 0] = [0, 0.32]. \tag{34}$$

Comment. We have described several methods and several ideas. On our simple example, some ideas work better and some lead to wider enclosures. The fact that a method works better on the simple example does not mean that it always works better; it depends on the function. In large-scale practical examples, it is useful to *combine* all these methods and ideas: for example, bisect and use centered form and monotonicity on subboxes (see, e.g., Jaulin et al. [2]).

The interval method—one of the above or their combination—has to be carefully chosen to match the function at hand. There exist several semi-empirical heuristics on which method to choose (see, e.g., Jaulin et al. [2]).

8 Situations when, in addition to the upper bounds on the measurement error, we also have partial information about the probabilities of different error values

Practical problem. In interval computations, we assume that the uncertainty in x_i can be described by the interval of possible values. In real life, in addition to the intervals, we often have some information about the probabilities of different values within this interval. What can we then do?

Which is the best way to describe the corresponding probabilistic uncertainty?

One of the main objectives of data processing is to make decisions. A standard way of making a decision is to select the action a for which the expected utility (gain) is the largest possible. This is where probabilities are used: in computing, for every possible action a , the corresponding expected utility. To be more precise, we usually know, for each action a and for each actual value of the (unknown) quantity x , the corresponding value of the utility $u_a(x)$. We must use the probability distribution for x to compute the expected value $e[u_a(x)]$ of this utility.

In view of this application, the most useful characteristics of a probability distribution would be the ones that would enable us to compute the expected value $e[u_a(x)]$ of different functions $u_a(x)$.

Which representations are the most useful for this intended usage? General idea. Which characteristics of a probability distribution are the most useful for computing mathematical expectations of different functions $u_a(x)$? The answer to this question depends on the type of the function, that is, on how the utility value u depends on the value x of the analyzed parameter.

Smooth utility functions naturally lead to moments. One natural case is when the utility function $u_a(x)$ is smooth. We have already mentioned, in Section 1, that we usually know a (reasonably narrow) interval of possible values of x . So, to compute the expected value of $u_a(x)$, all we need to know is how the function $u_a(x)$ behaves on this narrow interval. Because the function is smooth, we can expand it into Taylor series. Because the interval is narrow, we can safely consider only linear and quadratic terms in this expansion and ignore higher-order terms: $u_a(x) \approx c_0 + c_1 \cdot (x - x_0) + c_2 \cdot (x - x_0)^2$, where x_0 is a point inside the interval. Thus, we can approximate the expectation of this function by the expectation of the corresponding quadratic expression: $e[u_a(x)] \approx e[c_0 + c_1 \cdot (x - x_0) + c_2 \cdot (x - x_0)^2]$, that is, by the following expression, $e[u_a(x)] \approx c_0 + c_1 \cdot e[x - x_0] + c_2 \cdot e[(x - x_0)^2]$. So, to compute the expectations of such utility functions, it is sufficient to know the first and second moments of the probability distribution.

In particular, if we use, as the point x_0 , the average $e[x]$, the second moment turns into the variance of the original probability distribution. So, instead of the first and the second moments, we can use the mean E and the variance V .

In decision making, nonsmooth utility functions are common. In decision making, not all dependencies are smooth. There is often a threshold x_0 after which, say, a concentration of a certain chemical becomes dangerous.

This threshold sometimes comes from the detailed chemical and/or physical analysis. In this case, when we increase the value of this parameter, we see the drastic increase in effect and hence, the drastic change in utility value. Sometimes, this threshold simply comes from regulations. In this case, when

we increase the value of this parameter past the threshold, there is no drastic increase in effects, but there is a drastic decrease of utility due to the necessity to pay fines, change technology, and so on. In both cases, we have a utility function that experiences an abrupt decrease at a certain threshold value x_0 .

Nonsmooth utility functions naturally lead to cdfs. We want to be able to compute the expected value $e[u_a(x)]$ of a function $u_a(x)$ that changes smoothly until a certain value x_0 , then drops its value and continues smoothly for $x > x_0$. We usually know the (reasonably narrow) interval that contains all possible values of x . Because the interval is narrow and the dependence before and after the threshold is smooth, the resulting change in $u_a(x)$ before x_0 and after x_0 is much smaller than the change at x_0 . Thus, with a reasonable accuracy, we can ignore the small changes before and after x_0 , and assume that the function $u_a(x)$ is equal to a constant u^+ for $x < x_0$, and to some other constant $u^- < u^+$ for $x > x_0$.

The simplest case is when $u^+ = 1$ and $u^- = 0$. In this case, the desired expected value $e[u_a^{(0)}(x)]$ coincides with the probability that $x < x_0$, that is, with the corresponding value $F(x_0)$ of the cumulative distribution function (cdf). A generic function $u_a(x)$ of this type, with arbitrary values u^- and u^+ , can be easily reduced to this simplest case, because, as one can easily check, $u_a(x) = u^- + (u^+ - u^-) \cdot u^{(0)}(x)$ and hence, $e[u_a(x)] = u^- + (u^+ - u^-) \cdot F(x_0)$.

Thus, to be able to compute easily the expected values of all possible nonsmooth utility functions, it is sufficient to know the values of the cdf $F(x_0)$ for all possible x_0 .

How to represent partial information about probabilities: General idea. In many cases, we have complete information about the probability distributions that describe the uncertainty of each of n inputs.

However, a practically interesting case is how to deal with situations when we only have partial information about the probability distributions. How can we represent this partial information?

Case of cdf. If we use cdf $F(x)$ to represent a distribution, then full information corresponds to the case when we know the exact value of $F(x)$ for every x . Partial information means:

- Either that we only know approximate values of $F(x)$ for all x , that is, that for every x , we only know the interval that contains $F(x)$; in this case, we get a *p-box* (Person [22]).
- Or that we only know the values of $F(x)$ for some x , that is, that we only know the values $F(x_1), \dots, F(x_i), \dots, F(x_n)$ for finitely many values $x = x_1, \dots, x_i, \dots, x_n$; in this case, we have a *histogram*.

It is also possible that we know only approximate values of $F(x)$ for some x ; in this case, we have an *interval-valued histogram*.

Case of moments. If we use moments to represent a distribution, then partial information means that we either know the exact values of finitely many moments, or that we know intervals of possible values of several moments.

Resulting algorithms. This discussion leads to a natural classification of possible algorithms.

- If we have complete information about the distributions of x_i , then, to get validated estimates on uncertainty of y , we have to use Monte Carlo-type techniques (see, e.g., Lodwick and Jamison [23]).
- If we have p-boxes, we can use methods from Ferson [22] and Ferson et al. [24].
- If we have histograms, we can use methods from Berleant and Zhang [25].
- If we have moments, then we can use methods from Granvilliers et al. [26] and Kreinovich [27].

Case study: First moments. In some practical situations, in addition to the lower and upper bounds on each random variable x_i , we know the bounds $\mathbf{E}_i = [\underline{E}_i, \overline{E}_i]$ on its mean E_i . Indeed, in measurement practice (see, e.g., Rabinovich [1]), the overall measurement error Δx is usually represented as a sum of two components:

- A *systematic* error component $\Delta_s x$ which is defined as the expected value $e[\Delta x]$
- A *random* error component $\Delta_r x$ which is defined as the difference between the overall measurement error and the systematic error component:

$$\Delta_r x \stackrel{\text{def}}{=} \Delta x - \Delta_s x$$

In addition to the bound Δ on the overall measurement error, the manufacturers of the measuring instrument often provide an upper bound Δ_s on the systematic error component: $|\Delta_s x| \leq \Delta_s$.

This additional information is provided because, with this additional information, we not only get a bound on the accuracy of a single measurement, but we also get an idea of what accuracy we can attain if we use repeated measurements to increase the measurement accuracy. Indeed, the very idea that repeated measurements can improve the measurement accuracy is natural: we measure the same quantity by using the same measurement instrument several (N) times, and then take, for example, an arithmetic average

$$\bar{x} = \frac{\tilde{x}^{(1)} + \dots + \tilde{x}^{(k)} + \dots + \tilde{x}^{(N)}}{N}$$

of the corresponding measurement results $\tilde{x}^{(1)} = x + \Delta x^{(1)}, \dots, \tilde{x}^{(k)} = x + \Delta x^{(k)}, \dots, \tilde{x}^{(N)} = x + \Delta x^{(N)}$.

- If systematic error is the only error component, then all the measurements lead to exactly the same value $\tilde{x}^{(1)} = \dots = \tilde{x}^{(k)} = \dots = \tilde{x}^{(N)}$, and averaging does not change the value, hence does not improve the accuracy.

- On the other hand, if we know that the systematic error component is zero (i.e., $e[\Delta x] = 0$ and $e[\tilde{x}] = x$), then, as $N \rightarrow \infty$, the arithmetic average tends to the actual value x . In this case, by repeating the measurements sufficiently many times, we can determine the actual value of x with an arbitrary given accuracy.

In general, by repeating measurements sufficiently many times, we can arbitrarily decrease the random error component and thus attain accuracy as close to Δ_s as we want.

When this additional information is given, then, after we perform a measurement and get a measurement result \tilde{x} , then not only do we get the information that the actual value x of the measured quantity belongs to the interval $\mathbf{x} = [\tilde{x} - \Delta, \tilde{x} + \Delta]$, but we can also conclude that the expected value of $x = \tilde{x} - \Delta x$ (which is equal to $e[x] = \tilde{x} - e[\Delta x] = \tilde{x} - \Delta_s x$) belongs to the interval $\mathbf{E} = [\tilde{x} - \Delta_s, \tilde{x} + \Delta_s]$.

If we have this information for every x_i , then, in addition to the interval \mathbf{y} of possible values of y , we would also like to know the interval of possible values of $e[y]$. This additional interval will, it is hoped, provide us with the information on how repeated measurements can improve the accuracy of this indirect measurement. Thus, we arrive at the following problem.

Precise formulation of the problem. Given an algorithm computing a function $f(x_1, \dots, x_i, \dots, x_n)$ from R^n to R , and values $\underline{x}_1, \bar{x}_1, \dots, \underline{x}_i, \bar{x}_i, \dots, \underline{x}_n, \bar{x}_n, \underline{E}_1, \bar{E}_1, \dots, \underline{E}_i, \bar{E}_i, \dots, \underline{E}_n, \bar{E}_n$, we want to find

$$\begin{aligned} \underline{E} &\stackrel{\text{def}}{=} \min\{e[f(x_1, \dots, x_i, \dots, x_n)] \mid \text{all distributions of } (x_1, \dots, x_i, \dots, x_n) \\ &\text{for which } x_1 \in [\underline{x}_1, \bar{x}_1], \dots, x_i \in [\underline{x}_i, \bar{x}_i], \dots, x_n \in [\underline{x}_n, \bar{x}_n], \\ &e[x_1] \in [\underline{E}_1, \bar{E}_1], \dots, e[x_i] \in [\underline{E}_i, \bar{E}_i], \dots, e[x_n] \in [\underline{E}_n, \bar{E}_n]\}; \end{aligned} \quad (35)$$

and \bar{E} which is the maximum of $e[f(x_1, \dots, x_n)]$ for all such distributions.

In addition to considering all possible distributions, we can also consider the case when all the variables x_i are independent.

Algorithms for solving the problem: Case of exactly known moments. The main idea behind straightforward interval computations can be applied here as well. Namely, first, we find out how to solve this problem for the case when $n = 2$ and $f(x_1, x_2)$ is one of the standard arithmetic operations. Then, once we have an arbitrary algorithm $f(x_1, \dots, x_n)$, we parse it and replace each elementary operation on real numbers with the corresponding operation on quadruples $(\underline{x}, \underline{E}, \bar{E}, \bar{x})$.

To implement this idea, we must therefore know how to solve the above problem for elementary operations.

For *addition*, the answer is simple. Because $e[x_1 + x_2] = e[x_1] + e[x_2]$, if $y = x_1 + x_2$, there is only one possible value for $E = e[y]$: the value $E = E_1 + E_2$.

This value does not depend on whether we have correlation, or whether we have any information about the correlation. Thus, $\mathbf{E} = \mathbf{E}_1 + \mathbf{E}_2$.

Similarly, the answer is simple for *subtraction*: if $y = x_1 - x_2$, there is only one possible value for $E = e[y]$: the value $E = E_1 - E_2$. Thus, $\mathbf{E} = \mathbf{E}_1 - \mathbf{E}_2$.

For *multiplication*, if the variables x_1 and x_2 are independent, then $e[x_1 \cdot x_2] = e[x_1] \cdot e[x_2]$. Hence, if $y = x_1 \cdot x_2$ and x_1 and x_2 are independent, there is only one possible value for $E = e[y]$: the value $E = E_1 \cdot E_2$; hence $\mathbf{E} = \mathbf{E}_1 \cdot \mathbf{E}_2$.

The first nontrivial case is the case of multiplication in the presence of possible correlation. When we know the exact values of E_1 and E_2 , the solution to the above problem is as follows (see, e.g., Granvilliers et al. [26] and Kreinovich [27]): For multiplication $y = x_1 \cdot x_2$, when we have no information about the correlation,

$$\begin{aligned} \underline{E} &= \max(p_1 + p_2 - 1, 0) \cdot \bar{x}_1 \cdot \bar{x}_2 + \min(p_1, 1 - p_2) \cdot \bar{x}_1 \cdot \underline{x}_2 + \\ &\quad \min(1 - p_1, p_2) \cdot \underline{x}_1 \cdot \bar{x}_2 + \max(1 - p_1 - p_2, 0) \cdot \underline{x}_1 \cdot \underline{x}_2; \end{aligned} \tag{36}$$

$$\begin{aligned} \bar{E} &= \min(p_1, p_2) \cdot \bar{x}_1 \cdot \bar{x}_2 + \max(p_1 - p_2, 0) \cdot \bar{x}_1 \cdot \underline{x}_2 + \\ &\quad \max(p_2 - p_1, 0) \cdot \underline{x}_1 \cdot \bar{x}_2 + \min(1 - p_1, 1 - p_2) \cdot \underline{x}_1 \cdot \underline{x}_2, \end{aligned} \tag{37}$$

where $p_i \stackrel{\text{def}}{=} (E_i - \underline{x}_i)/(\bar{x}_i - \underline{x}_i)$.

For the *reciprocal* $y = 1/x_1$, the finite range is possible only when $0 \notin \mathbf{x}_1$. Without losing generality, we can consider the case when $0 < \underline{x}_1$. In this case, the range of possible values of E is $\mathbf{E} = [1/E_1, p_1/\bar{x}_1 + (1 - p_1)/\underline{x}_1]$.

Similar formulas can be produced for max and min, and also for the cases when there is a strong correlation between x_i : namely, when x_1 is (nonstrictly) increasing or decreasing in x_2 .

Algorithms for solving the problem: General case. For multiplication (under no assumption about correlation), if we only know the intervals of possible values of E_i , then to find \underline{E} it is sufficient to consider the following combinations of p_1 and p_2 .

- $p_1 = \underline{p}_1$ and $p_2 = \underline{p}_2$; $p_1 = \underline{p}_1$ and $p_2 = \bar{p}_2$; $p_1 = \bar{p}_1$ and $p_2 = \underline{p}_2$; $p_1 = \bar{p}_1$ and $p_2 = \bar{p}_2$.
- $p_1 = \max(\underline{p}_1, 1 - \bar{p}_2)$ and $p_2 = 1 - p_1$ (if $1 \in \mathbf{p}_1 + \mathbf{p}_2$).
- $p_1 = \min(\bar{p}_1, 1 - \underline{p}_2)$ and $p_2 = 1 - p_1$ (if $1 \in \mathbf{p}_1 + \mathbf{p}_2$).

The smallest value of \underline{E} for all these cases is the desired lower bound \underline{E} .

To find \bar{E} , it is sufficient to consider the following combinations of p_1 and p_2 .

- $p_1 = \underline{p}_1$ and $p_2 = \underline{p}_2$; $p_1 = \underline{p}_1$ and $p_2 = \bar{p}_2$; $p_1 = \bar{p}_1$ and $p_2 = \underline{p}_2$; $p_1 = \bar{p}_1$ and $p_2 = \bar{p}_2$.
- $p_1 = p_2 = \max(\underline{p}_1, \underline{p}_2)$ (if $\mathbf{p}_1 \cap \mathbf{p}_2 \neq \emptyset$).
- $p_1 = p_2 = \min(\bar{p}_1, \bar{p}_2)$ (if $\mathbf{p}_1 \cap \mathbf{p}_2 \neq \emptyset$).

The largest value of \bar{E} for all these cases is the desired upper bound \bar{E} .

Important open problems. What if, in addition to intervals and first moments, we also know second moments? This problem is important for design of computer chips.

What if, in addition to moments, we also know p-boxes?

Additional problem: How to estimate bounds on the moments. If we knew the exact values Δx_i of the measurement errors, we could estimate the moments by using the standard formulas: the mean as $e(\Delta x) = 1/n \cdot (\Delta x_1 + \dots + \Delta x_i + \dots + \Delta x_n)$, the variance as $\text{var}(\Delta x) = (1/(n-1)) \cdot \sum_{i=1}^n (\Delta x_i - E(\Delta x))^2$, and the covariance as

$$\text{Cov}(\Delta x, \Delta y) = \frac{1}{n-1} \cdot \sum_{i=1}^n (\Delta x_i - E(\Delta x)) \cdot (\Delta y_i - E(\Delta y)). \quad (38)$$

In practice, we do not know the actual value of $\Delta x_i = \tilde{x}_i - x_i$; we only know an approximate value $\tilde{\Delta}x_i = \tilde{x}_i - \tilde{x}_i^{\text{st}}$, where \tilde{x}_i^{st} is the result of measuring the same quantity by a standard (much more accurate) measuring instrument.

For the standard measuring instrument, we often only know the upper bound Δ_i^{st} on its measurement error: $|\tilde{x}_i^{\text{st}} - x_i| \leq \Delta_i^{\text{st}}$. In this case, we only know that $\Delta x_i \in [\tilde{\Delta}x_i - \Delta_i^{\text{st}}, \tilde{\Delta}x_i + \Delta_i^{\text{st}}]$. So, to find guaranteed bounds for each of the above statistical characteristics $c(\Delta x_1, \dots, \Delta x_i, \dots, \Delta x_n)$, we must find the *range* of possible values of the corresponding characteristics when Δx_i belongs to the corresponding interval $[\underline{\Delta}x_i, \overline{\Delta}x_i]$.

For some characteristics, computing the corresponding range is easy. For example, the mean $E(\Delta x)$ is a monotonic function of all its variables, so its range can be computed as $e(\Delta) = [\underline{E}, \overline{E}]$, where $\underline{E} = (1/n) \cdot (\underline{\Delta}x_1 + \dots + \underline{\Delta}x_i + \dots + \underline{\Delta}x_n)$ and $\overline{E} = (1/n) \cdot (\overline{\Delta}x_1 + \dots + \overline{\Delta}x_i + \dots + \overline{\Delta}x_n)$.

For other statistics such as variance $\text{var}(\Delta x)$ or covariance $\text{Cov}(\Delta x, \Delta y)$, the problem is, in general, NP-hard; (Ferson et al. [28]). In such cases, in general, we have to use approximate techniques. There are, however, practical meaningful situations in which it is possible to efficiently compute the exact range of the variance and of other characteristics; the corresponding algorithms are summarized in Ferson et al. [29], and Kreinovich et al. [30] and [31].

Comment. Similar algorithms can be used in the general situation of statistical processing under interval uncertainty. Interval uncertainty can come from measurement errors, but there are also other sources of interval uncertainty:

- A source of interval uncertainty is the existence of detection limits for different sensors: if a sensor, for example, did not detect any ozone, this means that the ozone concentration is below its detection limit DL , that is, in the interval $[0, DL]$.

- Yet another source of interval uncertainty is discretized data: if we experiment on the fish and watch it daily, and a fish is alive on Day 5 but dead on Day 6, then all we know about its lifetime is that it is in the interval [5, 6].
- Expert estimates often come as intervals.
- The need to keep privacy in statistical (e.g., medical) databases also often leads to the fact that instead of recording, for example, exact age, what we only record is the interval [40, 50].

In all these situations, the algorithms from Ferson et al. [29] and Kreinovich et al. [30] and [31] can be used.

9 Final Remarks

The traditional statistical approach to processing measurement errors Δx_i is based on the assumption that we have *full* information about the probability distributions for these errors. Typically, it is assumed that these errors are independent and normally distributed, with known means and standard deviations.

In practice, however, we often only have *partial* information about the corresponding probability distributions. For example, sometimes, we only know the upper bound Δ_i on the (absolute values of the) measurement errors; that is, we only know that Δx_i belongs to the interval $[-\Delta_i, \Delta_i]$. In this case, a usual engineering approach is to select, from several possible distributions, the most “reasonable one,” for example, the uniform distribution on $[-\Delta_i, \Delta_i]$. We show that this selection sometimes drastically underestimates the error of indirect measurements. To get more adequate estimates, we must use *robust* statistical techniques, that is, techniques that take into account all the probability distributions which are consistent with our knowledge. For the case of interval uncertainty, such techniques are called *interval computations*. In this chapter, we overviewed interval computation techniques and more general techniques of robust statistics.

Acknowledgments

This work was supported in part by NSF grants HRD-0734825, EAR-0225670, and EIA-0080940, by Texas Department of Transportation grant No. 0-5453, by the Japan Advanced Institute of Science and Technology (JAIST) International Joint Research Grant 2006-08, and by the Max Planck Institut für Mathematik.

The author is extremely thankful to the anonymous referees for valuable suggestions.

References

1. Rabinovich, S.G.: Measurement Errors and Uncertainty. Theory and Practice, Springer Verlag, Berlin (2005)
2. Jaulin, L., et al.: Applied Interval Analysis. Springer Verlag, London (2001)
3. Kearfott, R.B., Kreinovich, V. (eds.): Applications of Interval Computations. Kluwer, Dordrecht (1996)
4. Jaynes, E.T., Bretthorst, G.L.: Probability Theory: The Logic of science, Cambridge University Press, Cambridge, UK (2003)
5. Sheskin, D.J.: Handbook of Parametric and Nonparametric Statistical Procedures, Chapman & Hall/CRC, Boca Raton, FL (2004)
6. Kreinovich, V., et al.: Monte-Carlo-type techniques for processing interval uncertainty, and their potential engineering applications. *Reliable Computing* **13**(1), 25–69 (2007)
7. Kreinovich, V., Ferson, S.: A new Cauchy-based black-box technique for uncertainty in risk analysis. *Reliability Engineering and Systems Safety* **85**(1–3), 267–279 (2004)
8. Trejo, R., Kreinovich, V.: Error estimations for indirect measurements. In: Rajasekaran, S. et al. (eds.), *Handbook on Randomized Computing*, pp. 673–729, Kluwer, Dordrecht (2001)
9. Archimedes, On the measurement of the circle, In: Heath, T.L. (ed.), *The Works of Archimedes*. Cambridge University Press, Cambridge (1897); Dover (1953)
10. Young, W.H.: Sull due funzioni a piu valori costituite dai limiti d'una funzione di variable reale a destra ed a sinistra di ciascun punto. *Rendiconti Accademia di Lincei, Classes di Scienza Fiziche*, **17**(5), 582–587 (1908)
11. Young, R.C.: The algebra of multi-valued quantities. *Mathematische Annalen* **104**, 260–290 (1931)
12. Dwyer, P.S.: *Linear Computations*, Wiley, New York (1951)
13. Warmus, M.: Calculus of approximations. *Bulletin de l'Academie Polonaise de Sciences* **4**(5), 253–257 (1956)
14. Sunaga, T.: Theory of interval algebra and its application to numerical analysis, *RAAG Memoirs, Ggujutsu Bunken Fukuy-kai, Tokyo* **2**, 29–46 (547–564) (1958)
15. Moore R.E.: Automatic error analysis in digital computation, Space Div. Report LMSD84821, Lockheed Missiles and Space Co. (1959)
16. Moore, R.E.: *Interval Analysis*. Prentice-Hall, Englewood Cliffs, NJ (1966)
17. Interval computations website <http://www.cs.utep.edu/interval-comp>
18. Cerimele, M.M., Venturini Zilli, M.: Effective numerical approximation by intervals. *Freiburger Intervall-Berichte*, 85/4, 1–24 (1985)
19. Kreinovich, V., et al.: Computational Complexity and Feasibility of Data Processing and Interval Computations, Kluwer, Dordrecht (1997)
20. Berz, M., Hoffstätter, G.: Computation and application of Taylor polynomials with interval remainder bounds. *Reliable Computing* **4**(1), 83–97 (1998)
21. Neumaier, A.: Taylor forms. *Reliable Computing* **9**, 43–79 (2002)
22. Ferson, S.: *Risk Assessment with Uncertainty Numbers: RiskCalc*, CRC Press, Boca Raton, FL (2002)
23. Lodwick, W.A., Jamison, K.D.: Estimating and validating the cumulative distribution of a function of random variables. *Reliable Computing* **9**(2), 127–141 (2003)
24. Ferson, S., et al.: Constructing Probability Boxes and Dempster-Shafer Structures. Sandia National Laboratories, Report SAND2002-4015 (2003)

25. Berleant, D., Zhang, J.: Representation and problem solving with the distribution envelope determination (DEnv) method. *Reliability Engineering and System Safety* **85**(1–3), 153–168 (2004)
26. Granvilliers, L., Kreinovich, V., Mueller, N.: Novel approaches to numerical software with result verification. In: Alt, R. et al. (eds.), *Numerical Software with Result Verification*, Springer Lectures Notes in Computer Science **2291**, 274–305 Springer, New York (2004)
27. Kreinovich, V.: Probabilities, intervals, what next? *Journal of Global Optimization* **29**(3) 265–280 (2004)
28. Ferson, S., et al.: Exact bounds on finite populations of interval data. *Reliable Computing* **11** (3), 207–233 (2005)
29. Ferson, S., et al.: *Experimental Uncertainty Estimation and Statistics for Data Having Interval Uncertainty*. Sandia National Laboratories, Report SAND2007-0939, May 2007; available as <http://www.ramas.com/intstats.pdf> and as part of the book DVD
30. Kreinovich, V., et al.: Towards combining probabilistic and interval uncertainty in engineering calculations. *Reliable Computing* **12**(6), 471–501 (2006)
31. Kreinovich, V., et al.: Interval versions of statistical techniques. *Journal of Computational and Applied Mathematics* **199**(2), 418–423 (2007)

Parameter Estimation Based on Least Squares Methods

Alistair B. Forbes

National Physical Laboratory, Teddington, Middlesex, TW11 0LW, UK
alistair.forbes@npl.co.uk

Summary. This chapter describes how standard linear and nonlinear least squares methods can be applied to a large range of regression problems. In particular, it is shown that for many problems for which there are correlated effects it is possible to develop algorithms that use structure associated with the variance matrices to solve the problems efficiently. It is also shown how least squares methods can be adapted to cope with outliers.

Key words: Least squares, parameter estimation, uncertainty evaluation, numerical algorithms

1 Introduction

Least squares parameter estimation methods are used widely throughout science and engineering and have particular relevance to metrology. There are perhaps three main reasons for this: (i) least squares fitting problems lead to straightforward computational problems for which reliable algorithms have been developed and used successfully, (ii) least squares approximation has a natural geometrical interpretation, and (iii) least squares estimates can be justified from a probabilistic point of view. In different circumstances, any one of these reasons could dictate a least squares approach. In recent years, metrology has developed a more probabilistic approach to data analysis and in many circumstances a least squares approach is appropriate for the statistical model associated with the data. In this chapter, we consider all three aspects with an emphasis on applying standard least squares algorithms to as broad a range of problems as possible. In Section 2, we set the context of model fitting in metrology, and in Sections 3 and 4 we overview the main algorithmic and statistical aspects associated with linear and nonlinear least squares methods. In Sections 5 to 7 we describe how various classes of regression problems can be reformulated as standard least squares methods. In Section 8 we show how least squares methods can be used for robust estimation, that is, model

fitting to data that could include outliers. Our concluding remarks are given in Section 9.

2 Model fitting in metrology

Many experimental data analysis problems [3]¹ involve characterising how the response ϕ of a system depends on variables $\mathbf{x} = (x_1, \dots, x_{p-1})^T$. The response ϕ can be modelled as a function $\phi = \phi(\mathbf{x}, \mathbf{a})$ where $\mathbf{a} = (a_1, \dots, a_n)^T$ are model parameters controlling the shape of the response function ϕ . For example, ϕ may be a polynomial function of a single variable x and \mathbf{a} represents the vector of polynomial coefficients. In many situations, it is the parameters \mathbf{a} that are the direct objects of interest because they describe features of the system under study. The values of the model parameters \mathbf{a} are often determined from experimental data $\mathbf{y} = (y_1, \dots, y_m)^T$ in which the responses y_i corresponding to variable values $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p-1})^T$ are measured. From this point of view, model fitting is a type of inverse problem [44] in as much as we are trying to determine the parameters \mathbf{a} indirectly from the measured responses. The measurements will reflect not only the system response but also random effects associated with the measurement process. These random effects can perhaps be modelled as

$$y_i = \phi_i(\mathbf{a}) + e_i,$$

where $\phi_i(\mathbf{a}) = \phi(\mathbf{x}_i, \mathbf{a})$ is the model response, y_i the recorded measurement of ϕ_i , and e_i models the perturbatory effect of the random influence factors. The exact contribution of these random effects is generally unknown but they can be characterised from a probabilistic point of view. For example, we could model these effects by stating that e_i is a sample from a probability distribution whose mean μ_i and variance $V_i = \sigma_i^2$ are known. Usually $\mu_i = 0$. The standard deviations σ_i tell us how much to weight each measurement y_i in determining an estimate $\hat{\mathbf{a}}$ of \mathbf{a} from the measurement data $\{(\mathbf{x}_i, y_i)\}$.

3 Linear least squares problems (LS)

Given $m \times n$ matrix C , m -vector \mathbf{y} , the linear least squares problem is to determine \mathbf{a}_{LS} that solves

$$\min_{\mathbf{a}} F_{LS}(\mathbf{a}) = \frac{1}{2}(\mathbf{y} - C\mathbf{a})^T(\mathbf{y} - C\mathbf{a}).$$

(The fraction $\frac{1}{2}$ is used so that related expressions are simpler.) The matrix C often arises from a model $\phi(\mathbf{x}, \mathbf{a})$ described as a linear combination of basis functions

¹ Citations in this chapter are given in alphabetical order.

$$\phi(\mathbf{x}, \mathbf{a}) = \sum_{j=1}^n a_j \phi_j(\mathbf{x}).$$

For datapoints (\mathbf{x}_i, y_i) , $i = 1, \dots, m$, if $C_{ij} = \phi_j(\mathbf{x}_i)$, then the model responses are given by $C\mathbf{a}$. At the solution, the partial derivatives of F_{LS} with respect to the parameters are zero, leading to the system of linear equations,

$$C^T C \mathbf{a} = C^T \mathbf{y}, \tag{1}$$

known as the *normal equations*. If C is full rank, so that $C^T C$ is invertible, the solution parameters are given by

$$\mathbf{a}_{LS} = C^\dagger \mathbf{y}, \quad C^\dagger = (C^T C)^{-1} C^T. \tag{2}$$

We note that (2) defines \mathbf{a}_{LS} as a linear transformation of the data vector \mathbf{y} .

The linear least squares estimate has the following geometrical interpretation. The columns \mathbf{c}_j , $j = 1, \dots, n$, and \mathbf{y} are vectors (or points) in \mathbb{R}^m , (i.e., m -dimensional Euclidean vector space). Linear combinations

$$C\mathbf{a} = \sum_{j=1}^n a_j \mathbf{c}_j = a_1 \mathbf{c}_1 + \dots + a_n \mathbf{c}_n$$

of the vectors \mathbf{c}_j define points in the n -dimensional linear subspace \mathcal{C} (a hyperplane) defined by these column vectors. The linear least squares solution defines the point $\hat{\mathbf{y}} = C\mathbf{a}$ on the linear subspace \mathcal{C} closest to \mathbf{y} . The vector $\mathbf{y} - \hat{\mathbf{y}}$ must be orthogonal (perpendicular) to the plane and in particular perpendicular to the vectors \mathbf{c}_j : $\mathbf{c}_j^T (\mathbf{y} - C\mathbf{a}) = 0$, $j = 1, \dots, n$. Writing these equations in matrix terms,

$$C^T (\mathbf{y} - C^T C \mathbf{a}) = \mathbf{0},$$

from which we derive the normal equations (1). The solution necessarily defines a minimum of the convex function $F_{LS}(\mathbf{a})$.

3.1 Orthogonal factorisation method to determine parameter estimates

If C has orthogonal factorisation [3, 28]

$$C = QR = [Q_1 \ Q_2] \begin{bmatrix} R_1 \\ 0 \end{bmatrix} = Q_1 R_1, \tag{3}$$

then, using the fact that for any vector \mathbf{v} , $\|Q\mathbf{v}\| = \|\mathbf{v}\|$, we have

$$\|\mathbf{y} - C\mathbf{a}\| = \|Q^T \mathbf{y} - Q^T C\mathbf{a}\| = \left\| \begin{bmatrix} \mathbf{t}_1 \\ \mathbf{t}_2 \end{bmatrix} - \begin{bmatrix} R_1 \\ 0 \end{bmatrix} \mathbf{a} \right\|,$$

where \mathbf{t}_1 is the first n and \mathbf{t}_2 the last $m - n$ elements of $\mathbf{t} = Q^T \mathbf{y}$; that is, $\mathbf{t}_1 = Q_1^T \mathbf{y}$ and $\mathbf{t}_2 = Q_2^T \mathbf{y}$. From this it is seen that $\|\mathbf{y} - C\mathbf{a}\|$ is minimised if \mathbf{a}_{LS} solves the upper triangular system

$$R_1 \mathbf{a} = \mathbf{t}_1.$$

The QR factorisation has the following geometrical interpretation. The column vectors \mathbf{c}_j of C define an n -dimensional subspace \mathcal{C} of \mathbb{R}^m . The orthogonal matrix Q defines an axis system for \mathbb{R}^m such that the n columns of Q_1 define an axis system for \mathcal{C} and the $m - n$ columns of Q_2 define an axis system for the space of vectors \mathcal{C}^\perp orthogonal to \mathcal{C} . The columns \mathbf{q}_j , $j = 1, \dots, n$, of Q_1 are constructed with \mathbf{q}_1 aligned with \mathbf{c}_1 so that there is an r_{11} such that $\mathbf{c}_1 = r_{11}\mathbf{q}_1$. The vector \mathbf{q}_2 is chosen to lie in the plane defined by \mathbf{c}_1 and \mathbf{c}_2 , and so there are scalars r_{12} and r_{22} such that $\mathbf{c}_2 = r_{12}\mathbf{q}_1 + r_{22}\mathbf{q}_2$, and so on. This gives the factorisation

$$[\mathbf{c}_1 \ \mathbf{c}_2 \ \cdots \ \mathbf{c}_n] = [\mathbf{q}_1 \ \mathbf{q}_2 \ \cdots \ \mathbf{q}_n] \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1n} \\ 0 & r_{22} & \cdots & r_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & r_{nn} \end{bmatrix};$$

that is, in matrix notation $C = Q_1 R_1$.

The coefficients \mathbf{t}_1 define the point $\hat{\mathbf{y}}$ in \mathcal{C} closest to \mathbf{y} as a linear combination of the columns of Q_1 ; solving $R_1 \mathbf{a} = \mathbf{t}_1$ redefines $\hat{\mathbf{y}}$ as a linear combination of the columns \mathbf{c}_j of C .

3.2 Minimum variance property of the least squares estimate

As well as a geometrical interpretation, the least squares estimate has the following statistical properties. Suppose \mathbf{y} represents the measurement of $C\mathbf{a}$ subject to random effects so that

$$\mathbf{y} = C\mathbf{a} + \mathbf{e},$$

where \mathbf{e} models the perturbatory effects associated with the measurement system. We regard \mathbf{e} and hence \mathbf{y} as realisations of vectors of random variables \mathbf{E} and \mathbf{Y} , respectively. If the expectation $E(\mathbf{Y})$ and variance $V(\mathbf{Y})$ are given by $C\mathbf{a}$ and $\sigma^2 I$, respectively, then \mathbf{a}_{LS} is a realisation of a vector of random variables \mathbf{A}_{LS} with

$$E(\mathbf{A}_{\text{LS}}) = \mathbf{a}, \quad V(\mathbf{A}_{\text{LS}}) = V(\mathbf{A}_{\text{LS}}) = \sigma^2 (C^T C)^{-1}. \quad (4)$$

The first property is summarised by saying that \mathbf{a}_{LS} is an unbiased estimate of \mathbf{a} , generally a desirable property. Algebraically, it corresponds to the fact that $C^\dagger C = I$, so that

$$E(\mathbf{A}_{\text{LS}}) = E(C^\dagger \mathbf{Y}) = C^\dagger C \mathbf{a} = \mathbf{a}.$$

For this reason C^\dagger is sometimes referred to as the (left) pseudo-inverse of C . Given an $m \times n$ matrix A , the constraint that $A^T C = I$ can be implemented using the orthogonal factorisation (3) of C . We can express A as

$$A = [Q_1 \ Q_2] \begin{bmatrix} D_1 \\ D_2 \end{bmatrix} = Q_1 D_1 + Q_2 D_2,$$

where D_1 is $n \times n$ and D_2 is $(m - n) \times n$. The requirement that $A^T C = I$ can be written as

$$I = (D_1^T Q_1^T + D_2^T Q_2^T) Q_1 R_1 = D_1^T R_1 \Rightarrow D_1 = R_1^{-T}.$$

Hence, any A of the form $A = Q_1 R_1^{-T} + Q_2 D_2$ will determine an unbiased linear estimator $\hat{\mathbf{a}} = A^T \mathbf{y}$ of \mathbf{a} . The matrix $C^\dagger = R_1^{-1} Q_1^T$ is defined by $D_2 = \mathbf{0}$ and so in some sense is the simplest such estimator.

Given A such that $A = Q_1 R_1^{-T} + Q_2 D_2$, the variance of $\mathbf{A} = A^T \mathbf{Y}$ is

$$\sigma^2 A^T A = \sigma^2 (R_1^{-1} R_1^{-T} + D_2^T D_2),$$

and is minimised if $D_2 = \mathbf{0}$. This shows that if $V(\mathbf{Y}) = \sigma^2 I$, the least squares estimate (2) is the minimum variance, unbiased linear estimate of \mathbf{a} and is associated with uncertainty matrix $V_{\mathbf{a}_{LS}}$ given in (4).

3.3 Linear Gauss–Markov problem (GM)

Given C and \mathbf{y} as above and $m \times m$ a strictly positive definite matrix $V_{\mathbf{y}}$, the linear Gauss–Markov problem is to determine \mathbf{a}_{GM} that solves

$$\min_{\mathbf{a}} F_{GM}(\mathbf{a}) = \frac{1}{2} (\mathbf{y} - C\mathbf{a})^T V_{\mathbf{y}}^{-1} (\mathbf{y} - C\mathbf{a}). \tag{5}$$

If $V_{\mathbf{y}}$ has Cholesky factorisation $V_{\mathbf{y}} = LL^T$ [28] with L lower triangular, and

$$\tilde{C} = L^{-1}C, \quad \tilde{\mathbf{y}} = L^{-1}\mathbf{y},$$

then (5) is equivalent to the linear least squares problem

$$\min_{\mathbf{a}} \frac{1}{2} (\tilde{\mathbf{y}} - \tilde{C}\mathbf{a})^T (\tilde{\mathbf{y}} - \tilde{C}\mathbf{a}), \tag{6}$$

and therefore has solution

$$\mathbf{a}_{GM} = \left(\tilde{C}^T \tilde{C} \right)^{-1} \tilde{C}^T \tilde{\mathbf{y}} = \left(C^T V_{\mathbf{y}}^{-1} C \right)^{-1} C^T V_{\mathbf{y}}^{-1} \mathbf{y}. \tag{7}$$

From this last equation, \mathbf{a}_{GM} is such that

$$(\mathbf{y} - C\mathbf{a}_{GM})^T V_{\mathbf{y}}^{-1} C = \mathbf{0};$$

that is, the vector $\mathbf{y} - C\mathbf{a}_{GM}$ is orthogonal to the columns of C in the metric defined by $V_{\mathbf{y}}^{-1}$. If $\mathbf{y} \in \mathbf{Y}$, that is, \mathbf{y} is a realisation of random variables \mathbf{Y} , and $V(\mathbf{Y}) = V_{\mathbf{y}}$, then $\tilde{\mathbf{y}} \in \tilde{\mathbf{Y}}$ with $V(\tilde{\mathbf{Y}}) = I$, so that \mathbf{a}_{GM} is the minimum variance, unbiased linear estimate of \mathbf{a} associated with (6) and therefore (5). The Gauss–Markov estimate \mathbf{a}_{GM} can be determined using a QR factorisation approach to solve (6). The uncertainty matrix $V_{\mathbf{a}_{GM}}$ associated with \mathbf{a}_{GM} is

$$V_{\mathbf{a}_{GM}} = \left(\tilde{C}^T \tilde{C} \right)^{-1} = \left(C^T V_{\mathbf{y}}^{-1} C \right)^{-1}. \tag{8}$$

If $V_{\mathbf{y}} = \sigma^2 I$, then $\mathbf{a}_{GM} = \mathbf{a}_{LS}$ and $V_{\mathbf{a}_{GM}} = V_{\mathbf{a}_{LS}}$.

3.4 Generalised QR factorisation approach to the Gauss–Markov problem

The generalised QR decomposition can be employed to solve (5), avoiding the calculation of the inverse of a matrix in forming \tilde{C} in (7). In fact the generalised QR approach can be applied even if the uncertainty matrix $V_{\mathbf{y}}$ is rank deficient [6, 12, 29, 43]. The generalised QR approach starts with a factorisation of $V_{\mathbf{y}}$ of the form $V_{\mathbf{y}} = BB^T$, where B is $m \times p$. Often $p = m$ but the approach applies in the more general case. The matrix B could be the Cholesky factor of $V_{\mathbf{y}}$ but an uncertainty matrix usually can be expressed naturally in such a factored form; see, for example, Section 5. The estimate \mathbf{a}_{GM} can be found by solving

$$\min_{\mathbf{a}, \mathbf{d}} \frac{1}{2} \mathbf{d}^T \mathbf{d} \quad \text{subject to constraints} \quad \mathbf{y} = C\mathbf{a} + B\mathbf{d}. \quad (9)$$

Note that if B is invertible,

$$\mathbf{d} = B^{-1}(\mathbf{y} - C\mathbf{a}), \quad \mathbf{d}^T \mathbf{d} = (\mathbf{y} - C\mathbf{a})^T V_{\mathbf{y}}^{-1} (\mathbf{y} - C\mathbf{a}).$$

We factorise $C = QR$ and $Q^T B = TU$ where R and T are upper-triangular and Q and U are orthogonal. Multiplying the constraints by Q^T , we have

$$\begin{bmatrix} \tilde{\mathbf{y}}_1 \\ \tilde{\mathbf{y}}_2 \end{bmatrix} = \begin{bmatrix} R_1 \\ \mathbf{0} \end{bmatrix} \mathbf{a} + \begin{bmatrix} T_{11} & T_{12} \\ & T_{22} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{d}}_1 \\ \tilde{\mathbf{d}}_2 \end{bmatrix}, \quad (10)$$

where $\tilde{\mathbf{y}} = Q^T \mathbf{y}$ and $\tilde{\mathbf{d}} = U\mathbf{d}$. From the second set of equations, $\tilde{\mathbf{d}}_2$ must satisfy $\tilde{\mathbf{y}}_2 = T_{22} \tilde{\mathbf{d}}_2$. Given any $\tilde{\mathbf{d}}_1$, the first set of equations is satisfied if $R_1 \mathbf{a} = \tilde{\mathbf{y}}_1 - T_{11} \tilde{\mathbf{d}}_1 - T_{12} \tilde{\mathbf{d}}_2$. We choose $\tilde{\mathbf{d}}_1 = \mathbf{0}$ in order to minimise

$$\mathbf{d}^T \mathbf{d} = \tilde{\mathbf{d}}^T \tilde{\mathbf{d}} = \tilde{\mathbf{d}}_1^T \tilde{\mathbf{d}}_1 + \tilde{\mathbf{d}}_2^T \tilde{\mathbf{d}}_2,$$

so that $\hat{\mathbf{a}}_{\text{GM}}$ solves $R_1 \mathbf{a} = \tilde{\mathbf{y}}_1 - T_{12} \tilde{\mathbf{d}}_2$. The uncertainty matrix $V_{\mathbf{a}_{\text{GM}}}$ associated with $\hat{\mathbf{a}}_{\text{GM}}$ is given by

$$V_{\mathbf{a}_{\text{GM}}} = K K^T, \quad \text{where } K \text{ solves } R_1 K = T_{11}.$$

3.5 Linear least squares, maximum likelihood estimation, and the posterior distribution $p(\mathbf{a}|\mathbf{y})$

The uncertainty matrices $V_{\mathbf{a}}$ and $V_{\mathbf{a}_{\text{GM}}}$ defined in (4) and (8), respectively, are derived from the law of propagation of uncertainty [GUM95, 13]. If we make the further assumption that the observed data \mathbf{y} are a sample from a multivariate normal distribution,

$$\mathbf{y} \in N(C\mathbf{a}, V_{\mathbf{y}}), \quad (\text{i.e., } \mathbf{y} \in \mathbf{Y}, \quad \mathbf{Y} \sim N(C\mathbf{a}, V_{\mathbf{y}})),$$

then the Gauss–Markov solution \mathbf{a}_{GM} is a sample from a multivariate normal distribution:

$$\mathbf{a}_{\text{GM}} \in N(\mathbf{a}, V_{\mathbf{a}_{\text{GM}}}).$$

The Gauss–Markov solution is also the maximum likelihood estimate. The probability $p(\mathbf{y}|\mathbf{a})$ of observing \mathbf{y} , given \mathbf{a} , is such that

$$p(\mathbf{y}|\mathbf{a}) \propto \exp \left\{ -\frac{1}{2}(\mathbf{y} - C\mathbf{a})^T V_{\mathbf{y}}^{-1}(\mathbf{y} - C\mathbf{a}) \right\},$$

and is maximised by \mathbf{a}_{GM} .

In a Bayesian context [9, 26, 38], we regard \mathbf{a} as a vector of parameters, information about which is described in terms of probability distributions. If there is no substantial prior information about \mathbf{a} so that the prior distribution $p(\mathbf{a})$ can be taken to be a constant (i.e., $p(\mathbf{a}) \equiv 1$), the posterior probability distribution $p(\mathbf{a}|\mathbf{y})$ is proportional to the likelihood, regarded as a function of \mathbf{a} :

$$p(\mathbf{a}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{a}).$$

The vector $\mathbf{y} - C\mathbf{a}$ can be written as the sum of two mutually $V_{\mathbf{y}}^{-1}$ -orthogonal vectors

$$\mathbf{y} - C\mathbf{a} = [\mathbf{y} - C\mathbf{a}_{\text{GM}}] + [C(\mathbf{a}_{\text{GM}} - \mathbf{a})],$$

so that

$$\begin{aligned} (\mathbf{y} - C\mathbf{a})^T V_{\mathbf{y}}^{-1}(\mathbf{y} - C\mathbf{a}) &= (\mathbf{y} - C\mathbf{a}_{\text{GM}})^T V_{\mathbf{y}}^{-1}(\mathbf{y} - C\mathbf{a}_{\text{GM}}) \\ &\quad + (\mathbf{a}_{\text{GM}} - \mathbf{a})^T C^T V_{\mathbf{y}}^{-1} C(\mathbf{a}_{\text{GM}} - \mathbf{a}). \end{aligned}$$

The first term on the right does not depend on \mathbf{a} and so

$$p(\mathbf{a}|\mathbf{y}) \propto \exp \left\{ -\frac{1}{2}(\mathbf{a} - \mathbf{a}_{\text{GM}})^T C^T V_{\mathbf{y}}^{-1} C(\mathbf{a} - \mathbf{a}_{\text{GM}}) \right\}. \quad (11)$$

Comparing the right-hand side with the multivariate normal distribution, we see that

$$\mathbf{a}|\mathbf{y} \sim N(\mathbf{a}_{\text{GM}}, V_{\mathbf{a}_{\text{GM}}}),$$

where $V_{\mathbf{a}_{\text{GM}}}$ is defined in (8). These calculations show that the distribution $p(\mathbf{a}_{\text{GM}}|\mathbf{a})$ associated with observing a GM estimate \mathbf{a}_{GM} , given \mathbf{a} , is

$$\mathbf{a}_{\text{GM}}|\mathbf{a} \sim N(\mathbf{a}, V_{\mathbf{a}_{\text{GM}}}).$$

On the other hand, the distribution $p(\mathbf{a}|\mathbf{a}_{\text{GM}})$ for \mathbf{a} having observed a GM estimate \mathbf{a}_{GM} is

$$\mathbf{a}|\mathbf{a}_{\text{GM}} \sim N(\mathbf{a}_{\text{GM}}, V_{\mathbf{a}_{\text{GM}}}).$$

The symmetry in these two statements reflects the fact that \mathbf{a} and \mathbf{a}_{GM} appear symmetrically in (11). Furthermore, $p(\mathbf{a}|\mathbf{a}_{\text{GM}}) = p(\mathbf{a}|\mathbf{y})$, so that from this

point of view, the GM estimate does not lose any of the information that can be derived from the data \mathbf{y} .

Algorithms for solving linear least squares systems are described in detail in [6, 28, 37, 46]. There are linear least squares solvers in the NAG library, LINPACK, MINPACK, LAPACK, and Matlab, for example, [14, 25, 39, 42, 43]. LAPACK (and the NAG library) also has software for computing the generalised QR factorisation.

4 Nonlinear least squares (NLS)

The nonlinear least squares problem is: given m functions $f_i(\mathbf{a})$ of parameters $\mathbf{a} = (a_1, \dots, a_n)$, $m \geq n$, determine \mathbf{a}_{NLS} that solves

$$\min_{\mathbf{a}} F_{\text{NLS}}(\mathbf{a}) = \frac{1}{2} \sum_{i=1}^m f_i^2(\mathbf{a}) = \frac{1}{2} \mathbf{f}^T(\mathbf{a}) \mathbf{f}(\mathbf{a}). \quad (12)$$

Necessary conditions for a solution are that

$$\frac{\partial F_{\text{NLS}}}{\partial a_j} = \sum_{i=1}^m f_i \frac{\partial f_i}{\partial a_j} = 0, \quad j = 1, \dots, n.$$

Defining the *Jacobian matrix* $J = J(\mathbf{a})$ by

$$J_{ij} = \frac{\partial f_i}{\partial a_j}(\mathbf{a}),$$

this condition can be written as $J^T(\mathbf{a}) \mathbf{f}(\mathbf{a}) = \mathbf{0}$.

4.1 The Gauss–Newton algorithm for nonlinear least squares

The Gauss–Newton algorithm is a modification of Newton’s algorithm for minimising a function. Let J be the Jacobian matrix associated with $\mathbf{f}(\mathbf{a})$. Then the gradient \mathbf{g} of $F_{\text{NLS}}(\mathbf{a})$ is given by $\mathbf{g} = J^T \mathbf{f}$ and the Hessian matrix H of second partial derivatives is given by $H = J^T J + G$, where

$$G_{jk} = \sum_{i=1}^m f_i \frac{\partial^2 f_i}{\partial a_j \partial a_k}. \quad (13)$$

For Newton’s algorithm, given estimates \mathbf{a} of the solution, updated estimates are given by $\mathbf{a} := \mathbf{a} + \mathbf{p}$ where \mathbf{p} solves $H\mathbf{p} = -\mathbf{g}$. Usually a line search is incorporated so that the update step is of the form $\mathbf{a} := \mathbf{a} + t\mathbf{p}$ where $t > 0$ is chosen to ensure sufficient progress to the minimum is made. The Gauss–Newton variant follows the same approach, only that in determining the update step, H is approximated by $J^T J$; that is, the term G is ignored

and \mathbf{p} is found by solving $J^T J \mathbf{p} = -J^T \mathbf{f}$. This corresponds to the linear least squares problem $J \mathbf{p} = -\mathbf{f}$ and can be solved using an orthogonal factorisation approach, for example; see Section 3.1. The Gauss–Newton algorithm in general converges linearly at a rate that depends on the condition of the approximation problem, the size of the residuals \mathbf{f} near the solution, and the curvature (i.e., the degree of nonlinearity). If the problem is well-conditioned, the residuals are small, and the summand functions f_i are nearly linear, then $J^T J$ is a good approximation to the Hessian matrix H and convergence is fast.

If the model is

$$\mathbf{y} = \phi(\mathbf{a}) + \mathbf{e}, \quad \phi(\mathbf{a}) = (\phi_1(\mathbf{a}), \dots, \phi_m(\mathbf{a}))^T, \quad \mathbf{f}(\mathbf{a}) = \mathbf{y} - \phi(\mathbf{a}),$$

the mapping $\mathbf{a} \mapsto \phi(\mathbf{a})$ defines an n -dimensional surface in \mathbb{R}^m , and we look for \mathbf{a}_{NLS} that defines the point on the surface closest to \mathbf{y} . At the solution, the vector $\mathbf{f} = \mathbf{y} - \phi(\mathbf{a}_{\text{NLS}})$ is orthogonal to the surface at $\phi(\mathbf{a}_{\text{NLS}})$. The tangent plane at \mathbf{a}_{NLS} is

$$\phi(\mathbf{a}_{\text{NLS}} + \Delta) \approx \phi(\mathbf{a}_{\text{NLS}}) + J \Delta,$$

and so \mathbf{f} must be orthogonal to the columns of J , or in matrix terms $J^T \mathbf{f} = \mathbf{0}$, the optimality conditions.

The Gauss–Newton algorithm has the following geometrical interpretation. If the current estimate of the parameters is \mathbf{a} , the Jacobian matrix J evaluated at \mathbf{a} is used to construct the linear n -space \mathcal{J} defined by the columns of J . The step \mathbf{p} defines the point $J\mathbf{p}$ on \mathcal{J} closest to $\mathbf{y} - \phi(\mathbf{a})$.

4.2 Approximate uncertainty matrix associated with \mathbf{a}_{NLS}

Suppose the observation model is

$$\mathbf{y} = \phi(\mathbf{a}) + \mathbf{e}, \quad \mathbf{y} \in \mathbf{Y}, \quad \mathbf{E}(\mathbf{Y}) = \phi(\mathbf{a}), \quad \mathbf{V}(\mathbf{Y}) = \sigma^2 \mathbf{I},$$

and $\mathbf{f}(\mathbf{a})$ is defined as $\mathbf{f}(\mathbf{a}) = \mathbf{y} - \phi(\mathbf{a})$. In the linear case, the solution parameters \mathbf{a}_{LS} are given by a linear function of \mathbf{y} and it is possible to propagate $V_{\mathbf{y}}$ through to obtain the uncertainty matrix $V_{\mathbf{a}_{\text{LS}}}$ associated with \mathbf{a}_{LS} . In the nonlinear case, the estimate \mathbf{a}_{NLS} is defined as a nonlinear function of \mathbf{y} and there is no simple formula to allow us to calculate the variance matrix associated with the estimate. However, we can use linearisation about the solution \mathbf{a}_{NLS} to determine an approximate estimate. We regard the solution $\mathbf{a}_{\text{NLS}} \in \mathbf{A}_{\text{NLS}}$ as a realisation of a vector of random variables. If J is the Jacobian at the solution \mathbf{a}_{NLS} then

$$\mathbf{E}(\mathbf{A}_{\text{NLS}}) \approx \mathbf{a}, \quad \mathbf{V}(\mathbf{A}_{\text{NLS}}) \approx V_{\mathbf{a}_{\text{NLS}}} = \sigma^2 (J^T J)^{-1}. \quad (14)$$

The accuracy of the approximation depends mainly on the degree of nonlinearity. For highly nonlinear problems, these estimates can be very misleading. If J has QR factorisation $J = Q_1 R_1$, where R_1 is an $n \times n$ upper-triangular matrix, then $V_{\mathbf{a}_{\text{NLS}}} = \sigma^2 (R_1^T R_1)^{-1}$.

4.3 Nonlinear Gauss–Markov problem (NGM)

Given $\mathbf{f}(\mathbf{a})$ as before and an $m \times m$ strictly positive definite matrix $V_{\mathbf{y}}$, the nonlinear Gauss–Markov problem is to determine the solution \mathbf{a}_{NGM} of

$$\min_{\mathbf{a}} F_{\text{NGM}}(\mathbf{a}) = \frac{1}{2} \mathbf{f}^T(\mathbf{a}) V_{\mathbf{y}}^{-1} \mathbf{f}(\mathbf{a}). \quad (15)$$

As in the linear case, we can use the Cholesky decomposition $V_{\mathbf{y}} = LL^T$ to convert this problem to a standard nonlinear least squares problem applied to

$$\min_{\mathbf{a}} \frac{1}{2} \tilde{\mathbf{f}}^T(\mathbf{a}) \tilde{\mathbf{f}}(\mathbf{a}), \quad \tilde{\mathbf{f}} = L^{-1} \mathbf{f}. \quad (16)$$

We note that if J is the Jacobian matrix associated with \mathbf{f} , then $\tilde{J} = L^{-1}J$ is that associated with $\tilde{\mathbf{f}}$. As for the case of linear least squares, if $V_{\mathbf{y}}$, and hence L , is poorly conditioned, the formation and use of L^{-1} could lead to numerical instability. The Gauss–Newton algorithm can be adapted so that at each iteration the Gauss–Newton step is found by solving

$$\min_{\mathbf{a}, \mathbf{d}} \mathbf{d}^T \mathbf{d} \quad \text{subject to constraints} \quad \mathbf{f}(\mathbf{a}) = -J\mathbf{p} + L\mathbf{d},$$

using, for example, the generalised QR decomposition (Section 3.4). More generally, if $V_{\mathbf{y}}$ is given in factored form as $V_{\mathbf{y}} = BB^T$, then B can replace L in the above. There is no requirement for B to be a square matrix.

4.4 Approximate uncertainty matrix associated with \mathbf{a}_{NGM}

Suppose the observation model is

$$\mathbf{y} = \phi(\mathbf{a}) + \mathbf{e}, \quad \mathbf{y} \in \mathbf{Y}, \quad E(\mathbf{Y}) = \phi(\mathbf{a}), \quad V(\mathbf{Y}) = V_{\mathbf{y}},$$

and $\mathbf{f}(\mathbf{a})$ is defined as $\mathbf{f}(\mathbf{a}) = \mathbf{y} - \phi(\mathbf{a})$. We regard the solution $\mathbf{a}_{\text{NGM}} \in \mathbf{A}_{\text{NGM}}$ as a realisation of a vector of random variables. If J is the Jacobian for \mathbf{f} at the solution \mathbf{a}_{NGM} then,

$$E(\mathbf{A}_{\text{NGM}}) \approx \mathbf{a}, \quad V(\mathbf{A}_{\text{NGM}}) \approx V_{\mathbf{a}_{\text{NGM}}} = \left(J^T V_{\mathbf{y}}^{-1} J \right)^{-1}. \quad (17)$$

If $V_{\mathbf{y}}$ has Cholesky factorisation $V_{\mathbf{y}} = LL^T$ and $\tilde{J} = L^{-1}J$ has QR factorisation $\tilde{J} = Q_1 R_1$, where R_1 is an $n \times n$ upper-triangular matrix, then $V_{\mathbf{a}_{\text{NGM}}} = (R_1^T R_1)^{-1}$.

4.5 Nonlinear least squares, maximum likelihood estimation, and the posterior distribution $p(\mathbf{a}|\mathbf{y})$

If the measurement model is $\mathbf{y} \in N(\phi(\mathbf{a}), V_{\mathbf{y}})$, then \mathbf{a}_{NGM} is also the maximum likelihood estimate. From the definition of the multivariate normal distribution, the probability $p(\mathbf{y}|\mathbf{a})$ of observing \mathbf{y} , given parameter values \mathbf{a} , is such that

$$p(\mathbf{y}|\mathbf{a}) \propto \exp \left\{ -\frac{1}{2}(\mathbf{y} - \phi(\mathbf{a}))^T V_{\mathbf{y}}^{-1}(\mathbf{y} - \phi(\mathbf{a})) \right\},$$

and is maximised by \mathbf{a}_{NGM} .

In a Bayesian context, if there is no substantial prior information about \mathbf{a} so that the prior distribution $p(\mathbf{a})$ can be taken to be a constant, the posterior probability distribution $p(\mathbf{a}|\mathbf{y})$ is proportional to the likelihood. For the case of $V_{\mathbf{y}} = \sigma^2 I$, we have

$$p(\mathbf{a}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{a}) \propto \exp \left\{ -\frac{1}{2\sigma^2}(\mathbf{y} - \phi(\mathbf{a}))^T(\mathbf{y} - \phi(\mathbf{a})) \right\}.$$

The term on the right represents a multivariate normal distribution with respect to \mathbf{y} , up to a normalising constant. If $\phi(\mathbf{a})$ is a linear function of \mathbf{a} then it also represents a multivariate normal distribution with respect to \mathbf{a} . For nonlinear $\phi(\mathbf{a})$, the distribution can be quite different from a multinormal. (If m is much greater than n , then asymptotic results show that it is likely to be close to a multinormal distribution.) The shape of $p(\mathbf{a}|\mathbf{y})$ is determined by the term

$$d^2(\mathbf{y}, \phi(\mathbf{a})) = (\mathbf{y} - \phi(\mathbf{a}))^T(\mathbf{y} - \phi(\mathbf{a})),$$

where we can regard $d(\mathbf{y}, \phi(\mathbf{a}))$ as the distance from \mathbf{y} to the point on the surface $\phi(\mathbf{a})$. If we approximate $d^2(\mathbf{y}, \phi(\mathbf{a}))$ by a quadratic about the nonlinear least squares estimates \mathbf{a}_{NLS} , we find

$$p(\mathbf{a}|\mathbf{y}) \approx K \exp \left\{ -\frac{1}{2\sigma^2}(\mathbf{a} - \mathbf{a}_{\text{NLS}})^T H(\mathbf{a} - \mathbf{a}_{\text{NLS}}) \right\},$$

where $H = J^T J + G$ is the Hessian matrix associated with F_{NLS} in (12); that is, $p(\mathbf{a}|\mathbf{y})$ is approximated by the normal distribution $N(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}}^Q)$, where

$$V_{\mathbf{a}}^Q = \sigma^2 [J^T J + G]^{-1}.$$

If H is approximated by $J^T J$, then $p(\mathbf{a}|\mathbf{y})$ is approximated by $N(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}_{\text{NLS}}})$, where $V_{\mathbf{a}_{\text{NLS}}} = \sigma^2 (J^T J)^{-1}$, as in (14).

The approximation based on H is derived from a quadratic approximation to the surface $\phi(\mathbf{a})$ at \mathbf{a}_{NLS} which involves the matrix G of second partial derivatives whereas the approximation based on $J^T J$ is derived only from a linear approximation. Furthermore, the distribution $N(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}_{\text{NLS}}})$ depends only on the estimate \mathbf{a}_{NLS} whereas $N(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}}^Q)$ depends on \mathbf{y} through its contribution to G .

These calculations can be compared with (14) which states that the distribution $p(\mathbf{a}_{\text{NLS}}|\mathbf{a})$ for \mathbf{a}_{NLS} , given \mathbf{a} , is approximated by $N(\mathbf{a}, V_{\mathbf{a}})$, $V_{\mathbf{a}} = \sigma^2 (J^T J)^{-1}$, where J is the Jacobian matrix evaluated at \mathbf{a} . This corresponds to a forward evaluation of uncertainty: given \mathbf{a} , this distribution gives the likely distribution of the estimates \mathbf{a}_{NLS} , given the likely variation associated with the data vector \mathbf{y} . Because \mathbf{a} is generally unknown, \mathbf{a} is replaced

by its estimate \mathbf{a}_{NLS} and $V_{\mathbf{a}}$ by $V_{\mathbf{a}_{\text{NLS}}}$. The distribution $p(\mathbf{a}|\mathbf{y})$ corresponds to an inverse uncertainty evaluation: given that we have observed \mathbf{y} , $p(\mathbf{a}|\mathbf{y})$ describes the distribution of parameter values \mathbf{a} that could have given rise to \mathbf{y} . To a linear approximation, both $p(\mathbf{a}|\mathbf{y})$ and $p(\mathbf{a}_{\text{NLS}}|\mathbf{a})$ (evaluated with \mathbf{a} set to \mathbf{a}_{NLS}) are given by $N(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}_{\text{NLS}}})$. However, the more nonlinear the model, the less good are these linear approximations and the more disparate these two distributions can become. In particular, $p(\mathbf{a}|\mathbf{y})$ changes shape with \mathbf{y} whereas the shape of $p(\mathbf{a}_{\text{NLS}}|\mathbf{a})$ depends only on \mathbf{a} : two data vectors with the same least squares solution will lead to the same uncertainty matrix for \mathbf{a}_{NLS} . The distribution $p(\mathbf{a}_{\text{NLS}}, V_{\mathbf{a}}^{\text{Q}})$ will generally be a better approximation of $p(\mathbf{a}|\mathbf{y})$ because $V_{\mathbf{a}}^{\text{Q}}$ does have some dependence on \mathbf{y} . For mildly nonlinear problems and accurate data (say, a relative accuracy of better than 0.1%) then the Gaussian approximation should be accurate enough for most purposes.

The analysis for general $V_{\mathbf{y}}$ is entirely similar.

5 Exploiting structure in the uncertainty matrix

In many applications the uncertainty matrix $V_{\mathbf{y}}$ associated with the data vector \mathbf{y} is a full matrix. If the number m of observations is large, then solving the least squares problem can be computationally expensive because the number of operations required to perform a Cholesky factorisation of $V_{\mathbf{y}}$, for example, is $O(m^3)$ (i.e., scales with m^3). For this reason correlated effects are commonly ignored so that $V_{\mathbf{y}}$ is replaced by a diagonal matrix in order to reduce the computational requirement to $O(m)$. However, in many circumstances, a full uncertainty matrix $V_{\mathbf{y}}$ can be specified more compactly in a factored form that allows a more efficient organisation of the calculations.

5.1 Structure due to common random effects, linear case

Suppose the random effects associated with measurements are modelled as

$$\mathbf{y} = C\mathbf{a} + \mathbf{e} + K\mathbf{e}_0.$$

Here, e_i represents the random effect particular to the i th measured value, $\mathbf{e}_0 = (e_{1,0}, \dots, e_{k,0})^{\text{T}}$ those common to all the measurements, and K is an $m \times k$ matrix storing the sensitivities K_{ij} of y_i with respect to $e_{j,0}$. If \mathbf{e} and \mathbf{e}_0 are associated with uncertainty matrices $\sigma^2 I$ and U_0 , respectively, then the uncertainty matrix $V_{\mathbf{y}}$ associated with the data vector \mathbf{y} is given by

$$V_{\mathbf{y}} = \sigma^2 I + KU_0K^{\text{T}}.$$

The matrix $V_{\mathbf{y}}$ (and its Cholesky factor) is a full matrix by virtue of the common effects \mathbf{e}_0 . Estimate \mathbf{a}_{GM} of \mathbf{a} can be found by solving

$$\min_{\mathbf{a}} \frac{1}{2}(\mathbf{y} - C\mathbf{a})V_{\mathbf{y}}^{-1}(\mathbf{y} - C\mathbf{a}), \quad (18)$$

using the techniques described in Section 3.3. However, if $D = \sigma I$, U_0 has Cholesky factorisation $U_0 = L_0 L_0^T$ and $B_0 = K L_0$, then $V_{\mathbf{y}}$ can be factored as

$$V_{\mathbf{y}} = B B^T, \quad B = \begin{bmatrix} D & B_0 \end{bmatrix}, \quad (19)$$

and (18) has the same solution as

$$\min_{\mathbf{a}, \mathbf{d}, \mathbf{d}_0} \frac{1}{2} \left\{ \mathbf{d}^T \mathbf{d} + \mathbf{d}_0^T \mathbf{d}_0 \right\} \quad \text{subject to} \quad \mathbf{y} = C \mathbf{a} + D \mathbf{d} + B_0 \mathbf{d}_0. \quad (20)$$

The optimisation problem (20) can be written as

$$\min_{\mathbf{a}, \mathbf{d}, \mathbf{d}_0} \frac{1}{2} \left\{ \mathbf{d}^T \mathbf{d} + \mathbf{d}_0^T \mathbf{d}_0 \right\}, \quad \mathbf{d} = D^{-1} (\mathbf{y} - C \mathbf{a} - B_0 \mathbf{d}_0),$$

so that if

$$\tilde{C} = \begin{bmatrix} D^{-1} C & D^{-1} B_0 \\ \mathbf{0} & I \end{bmatrix}, \quad \tilde{\mathbf{y}} = \begin{bmatrix} D^{-1} \mathbf{y} \\ \mathbf{0} \end{bmatrix}, \quad \tilde{\mathbf{a}} = \begin{bmatrix} \mathbf{a} \\ \mathbf{d}_0 \end{bmatrix},$$

then (20), and hence (18), is equivalent to the standard linear least squares problem

$$\min_{\tilde{\mathbf{a}}} \frac{1}{2} (\tilde{\mathbf{y}} - \tilde{C} \tilde{\mathbf{a}})^T (\tilde{\mathbf{y}} - \tilde{C} \tilde{\mathbf{a}}).$$

By introducing the parameters \mathbf{d}_0 explicitly into the optimisation to explain the correlating effects, a simpler and more efficient solution method using the QR factorisation can be implemented. In the example above, $D = \sigma I$, and the approach can be extended to any D that is well-conditioned and for which it is computationally efficient to compute $D^{-1} C$, and so on.

Numerical example: Incorporating offsets

Suppose the response y of a system is a quadratic function $a_1 + a_2 x + a_3 x^2$ of x and that the measurements of y are made using an instrument with two accuracy settings, depending on the magnitude of x . The observation equations are of the form

$$y_i = a_1 + a_2 x_i + a_3 x_i^2 + e_{k,i} + e_{k,0}, \quad (21)$$

$e_{k,i} \in N(0, \sigma_k^2)$ and $e_{k,0} \in N(0, \sigma_{k,0}^2)$, where $k = 1, 2$, indicates the setting of the instrument. The term $e_{k,0}$ represents a fixed offset (systematic effect) associated with the k th setting. The uncertainty matrix $V_{\mathbf{y}}$ associated with \mathbf{y} can be factored as

$$V_{\mathbf{y}} = B B^T, \quad B = \begin{bmatrix} D & B_0 \end{bmatrix}, \quad D = \begin{bmatrix} \sigma_1 I & \\ & \sigma_2 I \end{bmatrix}, \quad B_0 = \begin{bmatrix} \sigma_{1,0} \mathbf{1} & \\ & \sigma_{2,0} \mathbf{1} \end{bmatrix},$$

where $\mathbf{1}$ represents a column vector of 1s.




Fig. 1. Least squares estimate associated with the model (21). The solid curve is generated by \mathbf{a} , the points marked by dots represent the data $\{(x_i, y_i)\}$, the dotted curve is generated by GM estimate \mathbf{a}_{GM} , and the dashed curve is generated by \mathbf{a}_{WLS} , and ignores the fixed offsets. The points marked by circles represent the adjusted data $\{(x, \hat{y}_i)\}$, where $\hat{y}_i = y_i - \hat{d}_{k,0}$.

Figure 1 graphs the GM estimate to simulated data generated with $\sigma_1 = 0.1$, $\sigma_2 = \sigma_{1,0} = 0.2$ and $\sigma_{2,0} = 0.4$, along with the weighted least squares fit \mathbf{a}_{WLS} determined by the solution of

$$\min_{\mathbf{a}} \frac{1}{2}(\mathbf{y} - C\mathbf{a})^T D^{-1}(\mathbf{y} - C\mathbf{a}),$$

that is, ignoring the contribution of the uncertainties associated with the fixed offsets \mathbf{d}_0 . An advantage of introducing \mathbf{d}_0 explicitly into the optimisation is that it is possible to provide estimates $\hat{\mathbf{d}}_0$ of \mathbf{d}_0 and use them to provide adjusted values $\hat{y}_i = y_i - \hat{d}_{k,0}$ of the measured response values. The standard uncertainties associated with estimates \mathbf{a}_{GM} are $\mathbf{u}(\mathbf{a}_{\text{GM}}) = (0.20, 0.20, 0.12)^T$, whereas those associated with \mathbf{a}_{WLS} , ignoring the contribution arising from the offsets, are $\mathbf{u}(\mathbf{a}_{\text{WLS}}) = (0.08, 0.20, 0.10)^T$, underestimating the uncertainty associated with a_1 , in particular. Solving the Gauss–Markov problem with the true variance matrix $V_{\mathbf{y}}$ allows the uncertainties associated with the systematic effects to be propagated through correctly to those associated with the solution parameters. Using the natural factorisation of $V_{\mathbf{y}}$ allows these calculations to be performed with no significant computational resource penalty.

5.2 Structure due to common random effects, nonlinear case

The same approach for the linear case, described above, also applies in the nonlinear case. If $V_{\mathbf{y}}$ can be factored as in (19), then setting

$$\tilde{\mathbf{a}} = \begin{bmatrix} \mathbf{a} \\ \mathbf{d}_0 \end{bmatrix}, \quad \tilde{\mathbf{f}}(\tilde{\mathbf{a}}) = \begin{bmatrix} D^{-1}(\mathbf{f}(\mathbf{a}) - B_0 \mathbf{d}_0) \\ \mathbf{d}_0 \end{bmatrix},$$

the nonlinear Gauss–Markov problem (15) is equivalent to

$$\min_{\tilde{\mathbf{a}}} \frac{1}{2} \tilde{\mathbf{f}}^T(\tilde{\mathbf{a}}) \tilde{\mathbf{f}}(\tilde{\mathbf{a}}),$$

a standard nonlinear least squares problem involving an augmented set of parameters $\tilde{\mathbf{a}}$.

There are a number of nonlinear least squares solvers in the MINPACK and NAG libraries, for example [25, 42]. Nonlinear least squares algorithms are described in [15, 27]; see also [41].

6 Generalised distance regression (GDR)

Linear and nonlinear least squares problems are appropriate if only one measured variable is subject to significant random effects. Typically, these problems are associated with models of the form

$$y_i = \phi(\mathbf{x}_i, \mathbf{a}) + e_i,$$

where it is assumed that $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p-1})^T$ are known accurately but the measurements of $\phi(\mathbf{x}_i, \mathbf{a})$ are subject to random effects represented by e_i . In generalised distance regression, we allow for the fact that the values of the variables are also subject to significant random effects, a situation not uncommon in metrology, but retaining the assumption that the random effects associated with (\mathbf{x}_i, y_i) are independent from those associated with (\mathbf{x}_q, y_q) , $q \neq i$. We can write this more general model as

$$\begin{bmatrix} \mathbf{x}_i \\ y_i \end{bmatrix} = \begin{bmatrix} \mathbf{u}_i \\ \phi(\mathbf{u}_i, \mathbf{a}) \end{bmatrix} + \mathbf{e}_i, \quad \mathbf{e}_i \in \mathbf{E}_i, \quad \mathbf{E}(\mathbf{E}_i) = \mathbf{0}, \quad \mathbf{V}(\mathbf{E}_i) = \mathbf{V}_i. \quad (22)$$

Here \mathbf{e}_i represents the perturbatory effects associated with the measured values \mathbf{x}_i and y_i . The left-hand side of (22) is a datapoint in \mathbb{R}^p and the mapping

$$\mathbf{u} \mapsto \begin{bmatrix} \mathbf{u} \\ \phi(\mathbf{u}, \mathbf{a}) \end{bmatrix}$$

describes a $(p - 1)$ -dimensional surface in \mathbb{R}^p . With this model in mind, we define the explicit generalised distance problem: given a parametrically defined surface $\phi(\mathbf{u}, \mathbf{a})$, where $\phi : \mathbb{R}^{p-1} \times \mathbb{R}^n \rightarrow \mathbb{R}^p$, data $\{\mathbf{x}_i \in \mathbb{R}^p\}$, and $p \times p$ symmetric, strictly positive definite matrices V_i , $i = 1, \dots, m$, determine \mathbf{a}_{GDR} and $\{\mathbf{u}_i^*\}_{i=1}^m$ that solve

$$\min_{\mathbf{a}, \{\mathbf{u}_i\}} F_{\text{GDR}}(\mathbf{a}) = \frac{1}{2} \sum_{i=1}^m (\mathbf{x}_i - \phi(\mathbf{u}_i, \mathbf{a}))^T V_i^{-1} (\mathbf{x}_i - \phi(\mathbf{u}_i, \mathbf{a})). \quad (23)$$

If each $V_i = I$, the identity matrix, the GDR problem is sometimes referred to as the orthogonal distance regression problem (ODR). Setting

$$\mathbf{y} = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_m \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_m \\ \mathbf{a} \end{bmatrix}, \quad \boldsymbol{\eta}(\mathbf{b}) = \begin{bmatrix} \phi(\mathbf{u}_1, \mathbf{a}) \\ \vdots \\ \phi(\mathbf{u}_m, \mathbf{a}) \end{bmatrix}, \quad V\mathbf{y} = \begin{bmatrix} V_1 & & \\ & \ddots & \\ & & V_m \end{bmatrix},$$

then (23) can be formulated as

$$\min_{\mathbf{b}} \frac{1}{2} (\mathbf{y} - \boldsymbol{\eta}(\mathbf{b}))^T V \mathbf{y}^{-1} (\mathbf{y} - \boldsymbol{\eta}(\mathbf{b})). \quad (24)$$

6.1 Algorithms for generalised distance regression

Although the GDR problem can be solved using the techniques for general nonlinear Gauss–Markov problems, the structure associated with the problem allows for a much more efficient solution approach. The structure can be exploited in two ways, as described below.

Structured matrix approach

If V_i has Cholesky factorisation $V_i = L_i L_i^T$, and $\mathbf{f}_i(\mathbf{u}_i, \mathbf{a}) = L_i^{-1}(\mathbf{x}_i - \phi(\mathbf{u}_i, \mathbf{a}))$, then (23) is equivalent to

$$\min_{\mathbf{a}, \{\mathbf{u}_i\}} \frac{1}{2} \sum_{i=1}^m \mathbf{f}_i^T(\mathbf{u}_i, \mathbf{a}) \mathbf{f}_i(\mathbf{u}_i, \mathbf{a}), \quad (25)$$

a standard nonlinear least squares problem. However, the fact that $p - 1$ parameters \mathbf{u}_i only appear in p equations means that the associated Jacobian matrix of partial derivatives has a block-angular structure

$$J = \begin{bmatrix} K_1 & & J_1 \\ & K_2 & J_2 \\ & & \ddots & \vdots \\ & & & K_m & J_m \end{bmatrix}, \quad (26)$$

where K_i is the matrix of derivatives of the i th set of observation equations with respect to \mathbf{u}_i , and the border blocks J_i store their derivatives with respect to \mathbf{a} . The upper-triangular factor R of the Jacobian matrix also has a block-angular structure:

$$R = \begin{bmatrix} R_1 & & B_1 \\ & R_2 & B_2 \\ & & \ddots & \vdots \\ & & & R_m & B_m \\ & & & & R_0 \end{bmatrix}, \quad (27)$$

where $\{R_i\}_{i=1}^m$, are $(p-1) \times (p-1)$ upper-triangular, $\{B_i\}_{i=1}^m$ are $(p-1) \times n$ border blocks and R_0 is the $n \times n$ upper-triangular factor corresponding to the parameters \mathbf{a} . Starting with $R_0 = \mathbf{0}$, $\mathbf{t}_0 = \mathbf{0}$, the upper-triangularisation proceeds by processing p rows at a time. At the i th stage, a $(p+n) \times (p+n)$ orthogonal matrix Q_i is found such that

$$Q_i^T \begin{bmatrix} K_i & J_i \\ & R_0 \end{bmatrix} = \begin{bmatrix} R_i & B_i \\ & R_0 \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad Q_i^T \begin{bmatrix} -\mathbf{f}_i \\ \mathbf{t}_0 \end{bmatrix} = \begin{bmatrix} \mathbf{t}_i \\ \mathbf{t}_0 \\ \mathbf{s} \end{bmatrix}.$$

The Gauss–Newton step is also determined efficiently from triangular systems involving R_0 and R_i :

$$\mathbf{p} = \begin{bmatrix} \mathbf{p}_1 \\ \vdots \\ \mathbf{p}_m \\ \mathbf{p}_0 \end{bmatrix}, \quad R_0 \mathbf{p}_0 = \mathbf{t}_0, \quad R_i \mathbf{p}_i = \mathbf{t}_i - B_i \mathbf{p}_0, \quad i = 1, \dots, m.$$

The use of structure exploiting algorithms for model fitting in metrology is discussed in [10, 11, 16, 18].

Separation of variables approach

The separation of variables approach converts (25) into a standard nonlinear least squares problem in the parameters \mathbf{a} . Denote by $\mathbf{u}_i^* = \mathbf{u}_i^*(\mathbf{a})$ the solution of the footpoint problem,

$$\min_{\mathbf{u}} \frac{1}{2} D(\mathbf{u}) = \frac{1}{2} (\mathbf{x}_i - \phi(\mathbf{u}, \mathbf{a}))^T V_i^{-1} (\mathbf{x}_i - \phi(\mathbf{u}, \mathbf{a})), \quad (28)$$

a special form of a nonlinear Gauss–Markov problem. Let \mathbf{n}_i be any vector orthogonal to the surface at $\mathbf{x}_i^* = \phi(\mathbf{u}_i^*, \mathbf{a})$. The conditions for \mathbf{u}_i^* to be a solution of (28) is that the vector $V_i^{-1} (\mathbf{x}_i - \mathbf{x}_i^*)$ is a scalar multiple of \mathbf{n}_i . From this, it is straightforward to show that if we define the *generalised distance* $d_i = d_i(\mathbf{a})$ by

$$d_i = \frac{1}{s_i} \mathbf{n}_i^T (\mathbf{x}_i - \mathbf{x}_i^*), \quad s_i = (\mathbf{n}_i^T V_i \mathbf{n}_i)^{1/2}, \quad (29)$$

then $d_i^2 = D(\mathbf{u}_i^*)$, and

$$\frac{\partial d_i}{\partial a_j} = -\frac{1}{s_i} \mathbf{n}_i^T \frac{\partial \phi}{\partial a_j}. \quad (30)$$

In this way, the GDR problem can be posed as a standard nonlinear least squares problem $\min_{\mathbf{a}} 1/2 \sum_i d_i^2(\mathbf{a})$, where each function and its gradient are calculated as in (29) and (30), with all quantities evaluated at the solution \mathbf{u}_i^* of the appropriate footpoint problem.

6.2 Sequential quadratic programming for the footpoint parameters

Both d_i and its derivatives are defined in terms of V_i , through s_i , rather than V_i^{-1} . If V_i can be factored as $V_i = B_i B_i^T$, then the footpoint problem can be posed as

$$\min_{\mathbf{u}, \mathbf{d}} \frac{1}{2} \mathbf{d}^T \mathbf{d} \quad \text{subject to} \quad \mathbf{x}_i = \phi(\mathbf{u}, \mathbf{a}) + B_i \mathbf{d}, \quad (31)$$

again avoiding the formation of V_i^{-1} . In fact, there is no requirement in implementing the separation of variables approach that V_i is full rank, only that $\mathbf{n}_i^T V_i \mathbf{n}_i$ is nonzero, where \mathbf{n}_i is normal to the surface. This means that the separation of variables approach applies to a greater range of problems than the structure-exploiting approach which requires that V_i is full rank.

A sequential quadratic programming approach can be applied to solve (31), leading to a quadratically converging algorithm [20]. We first describe an algorithm to minimise a positive definite quadratic function subject to linear equality constraints. Let A be an $n \times n$ positive definite, symmetric matrix, C a $n \times p$ matrix, $p < n$, and \mathbf{b} and \mathbf{d} n - and p -vectors, respectively. The quadratic programming problem is

$$\min_{\mathbf{v}} \frac{1}{2} \mathbf{v}^T A \mathbf{v} + \mathbf{b}^T \mathbf{v} \quad \text{subject to} \quad C^T \mathbf{v} = \mathbf{d}. \quad (32)$$

In a Lagrangian formulation the solution \mathbf{y} is associated with a stationary point of

$$\mathcal{L}(\mathbf{v}, \boldsymbol{\lambda}) = \frac{1}{2} \mathbf{v}^T A \mathbf{v} + \mathbf{b}^T \mathbf{v} - (C^T \mathbf{v} - \mathbf{d})^T \boldsymbol{\lambda},$$

involving \mathbf{v} and Lagrange multipliers $\boldsymbol{\lambda}$. Let

$$C = [Q_1 \ Q_2] \begin{bmatrix} R_1 \\ \mathbf{0} \end{bmatrix} = Q_1 R_1,$$

be the QR factorisation of C^T , where R_1 is $p \times p$ upper triangular and $Q = [Q_1 \ Q_2]$ is an $n \times n$ orthogonal matrix. We look for a solution of (32) of the form $\mathbf{v} = Q_1 \mathbf{v}_1 + Q_2 \mathbf{v}_2$. From the constraint equation we have

$$\mathbf{d} = C^T (Q_1 \mathbf{v}_1 + Q_2 \mathbf{v}_2) = R_1^T Q_1^T Q_1 \mathbf{v}_1 + R_1^T Q_1^T Q_2 \mathbf{v}_2 = R_1^T \mathbf{v}_1,$$

because $Q_1^T Q_1 = I$ and $Q_1^T Q_2 = \mathbf{0}$. This shows that \mathbf{v}_1 must satisfy $R_1^T \mathbf{v}_1 = \mathbf{d}$. These constraints fix \mathbf{v}_1 and we must choose \mathbf{v}_2 to minimise the quadratic expression which amounts to minimising

$$\frac{1}{2} \mathbf{v}_2^T Q_2^T A Q_2 \mathbf{v}_2 + \mathbf{v}_2^T Q_2^T (\mathbf{b} + A Q_1 \mathbf{v}_1)$$

with respect to \mathbf{v}_2 . The conditions for a minimum dictate that \mathbf{v}_2 solves the system

$$Q_2^T A Q_2 \mathbf{v}_2 = -Q_2^T (\mathbf{b} + A Q_1 \mathbf{v}_1).$$

If required, the Lagrange multipliers $\boldsymbol{\lambda}$ can be determined as the solution of $R_1 \boldsymbol{\lambda} = Q_1^T (\mathbf{b} + A \mathbf{v})$.

The footprint problem can be solved through a sequence of quadratic programming problems. Here we describe the algorithm for the case corresponding to (22). Suppose we wish to solve

$$\min_{\mathbf{u}, \mathbf{d}} \frac{1}{2} \mathbf{d}^T \mathbf{d} \quad \text{subject to} \quad \begin{bmatrix} \mathbf{x} \\ y \end{bmatrix} = \begin{bmatrix} \mathbf{u} \\ \phi(\mathbf{u}) \end{bmatrix} + \begin{bmatrix} B_1 \\ \mathbf{b}_2^T \end{bmatrix} \mathbf{d}. \quad (33)$$

The first set of constraint equations can be written as

$$\mathbf{x} = D \begin{bmatrix} \mathbf{u} \\ \mathbf{d} \end{bmatrix}, \quad D = [I \ B_1].$$

If D^T has QR factorisation

$$D^T = [Q_1 \ Q_2] \begin{bmatrix} R_1 \\ \mathbf{0} \end{bmatrix}, \quad Q_2 = \begin{bmatrix} Q_{12} \\ Q_{22} \end{bmatrix},$$

then, for any \mathbf{v} , $\mathbf{u} = \mathbf{x} + Q_{12} \mathbf{v}$ and $\mathbf{d} = Q_{22} \mathbf{v}$ satisfy the constraints. With this factorisation, (33) can be reformulated as

$$\min_{\mathbf{v}} \frac{1}{2} \mathbf{v}^T Q_{22}^T Q_{22} \mathbf{v} \quad \text{subject to} \quad y = \phi(\mathbf{x} + Q_{12} \mathbf{v}) + \mathbf{b}_2^T Q_{22} \mathbf{v},$$

involving only one, nonlinear, constraint. The associated Lagrangian is

$$\mathcal{L}(\mathbf{v}, \lambda) = \frac{1}{2} \mathbf{v}^T Q_{22}^T Q_{22} \mathbf{v} - \lambda \left(y - \phi(\mathbf{x} + Q_{12} \mathbf{v}) - \mathbf{b}_2^T Q_{22} \mathbf{v} \right),$$

and involves only one multiplier λ . The following algorithm updates estimates of \mathbf{v} and λ .

1. Evaluate the surface function ϕ , and its first and second partial derivatives with respect to \mathbf{u} :

$$\mathbf{u} = \mathbf{x} + Q_{12} \mathbf{v}, \quad \phi = \phi(\mathbf{u}), \quad \mathbf{g}_\phi = \nabla \phi, \quad H_\phi = \nabla^2 \phi.$$

2. Set

$$\mathbf{d} = Q_{22} \mathbf{v}, \quad \mathbf{g} = Q_{22}^T \mathbf{d}.$$

3. Evaluate constraint function c and its first and second derivatives with respect to \mathbf{v} :

$$c = y - \phi - \mathbf{b}_2^T \mathbf{d}, \quad \mathbf{h} = -Q_{12}^T \mathbf{g}_\phi - Q_{22}^T \mathbf{b}_2, \quad K = -Q_{12}^T H_\phi Q_{12}.$$

4. Set $A = Q_{22}^T Q_{22} - \lambda K$ and solve, for \mathbf{p} and updated λ , the quadratic programming problem

$$\min_{\mathbf{p}} \frac{1}{2} \mathbf{p}^T A \mathbf{p} + \mathbf{g}^T \mathbf{p} \quad \text{subject to} \quad \mathbf{h}^T \mathbf{p} = -c.$$

- Update $\mathbf{v} := \mathbf{v} + t\mathbf{p}$ for a suitable step length t . (Generally t is chosen to be an approximate minimum of a merit function that involves the function being minimised and a measure of how well the constraints are satisfied; near the solution we expect t to be close to 1.)

6.3 Example application: Response and evaluation calibration curves

We suppose that an instrument’s response y depends approximately linearly (or at least monotonically) on a variable x and that for a sequence of calibrated values $x_i, i = 1, \dots, m$, of x , measurements of the responses y_i are made. Given a model of the form

$$y_i = \phi(u_i, \mathbf{a}) + e_{y,i}, \quad x_i = u_i + e_{x,i}, \quad e_{x,i} \in N(0, \rho^2), \quad e_{y,i} \in N(0, \sigma^2),$$

the *response calibration curve* is found by solving the generalised distance regression problem

$$\min_{\mathbf{a}, \mathbf{u}, \mathbf{d}_x, \mathbf{d}_y} \frac{1}{2} \sum_{i=1}^m \{d_{x,i}^2 + d_{y,i}^2\} \quad \text{subject to} \quad \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} u_i \\ \phi(u_i, \mathbf{a}) \end{bmatrix} + \begin{bmatrix} \rho & 0 \\ 0 & \sigma \end{bmatrix} \begin{bmatrix} d_{x,i} \\ d_{y,i} \end{bmatrix},$$

$i = 1, \dots, m$. If $\rho = 0$, as in the case where the uncertainty in the calibrated values of x is much smaller than those associated with the response measurements, this problem reduces to a standard least squares problem

$$\min_{\mathbf{a}} \frac{1}{2} \sum_{i=1}^m (y_i - \phi(x_i, \mathbf{a}))^2.$$

Given a calibrated value of x , the response curve $\phi(x, \mathbf{a})$ predicts the response of the system. However, in using the instrument, we are interested in estimating the value of the stimulus variable x , given a measurement of the response y . If the instrument is calibrated in terms of the response curve ϕ , then every time we measure with the instrument, recording an uncalibrated response y , we have to use iterative techniques to solve $\phi(x) = y$ in order to output an estimate of the calibrated value x . A more attractive proposition is to model the instrument behaviour as $x = \psi(y, \mathbf{a})$ (so that $\phi = \psi^{-1}$) and the *evaluation calibration curve* is found by solving

$$\min_{\mathbf{a}, \mathbf{v}, \mathbf{d}_x, \mathbf{d}_y} \frac{1}{2} \sum_{i=1}^m \{d_{x,i}^2 + d_{y,i}^2\} \quad \text{subject to} \quad \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} \psi(v_i, \mathbf{a}) \\ v_i \end{bmatrix} + \begin{bmatrix} \rho & 0 \\ 0 & \sigma \end{bmatrix} \begin{bmatrix} d_{x,i} \\ d_{y,i} \end{bmatrix}.$$

Using a separation of variables approach, the case of $\rho = 0$ (or near zero) introduces no complications (nor numerical stability concerns). The output x can be determined from a direct evaluation of $\psi(y, \mathbf{a})$, given a measured response y .

Regarding the response and evaluation calibration curves as parametric curves,

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} u \\ \phi(u, \mathbf{a}) \end{bmatrix}, \quad \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \psi(u, \mathbf{a}) \\ u \end{bmatrix},$$

respectively, both problems can be solved as generalised distance regression problems using the same software.

6.4 Example: GDR line

In the case $\phi(x, \mathbf{a}) = a_1 + a_2x$, each V_i is a 2×2 uncertainty matrix. If

$$\begin{bmatrix} p_i \\ q_i \end{bmatrix} = V_i \begin{bmatrix} -a_2 \\ 1 \end{bmatrix},$$

then

$$u_i = \frac{-q_i x_i + p_i(y_i - a_1)}{-q_i + p_i a_2},$$

defines the point on the line closest to the datapoint (x_i, y_i) as measured using the metric V_i^{-1} . Setting

$$f_i = 1/(-bp_i + q_i)^{1/2}, \quad g_i = u_i f_i, \quad h_i = (-a_2(x_i - u_i) + y_i - a_1 - a_2 u_i) f_i,$$

the update step \mathbf{p} for estimate \mathbf{a} is found from the solution of the linear least squares problem

$$J\mathbf{p} = \mathbf{h}, \quad J = [\mathbf{f} \ \mathbf{g}].$$

The case of orthogonal distance regression is considered in [2, 7, 28, 30, 31, 34, 35, 45], for example. The software package ODRPACK [8] provides a fairly comprehensive facility. Generalised distance regression is considered in [1, 4, 11, 12, 16–18, 20–22, 24, 32, 40]. The component XGENLINE for polynomial generalised distance regression is available for downloading from <http://www.npl.co.uk/ssfm>.

7 Generalised Gauss–Markov regression (GGM)

Generalised Gauss–Markov regression combines generalised distance regression with nondiagonal uncertainty matrices [12]. We consider the case of a parametrically defined surface $\phi(\mathbf{u}, \mathbf{a})$, $\phi : \mathbb{R}^{p-1} \times \mathbb{R}^n \rightarrow \mathbb{R}^p$, and data points $\{\mathbf{x}_i\}_{i=1}^m$ nominally lying on such a surface subject to random effects. The generalised Gauss–Markov problem is formulated as in (24) but for general uncertainty matrix $V_{\mathbf{y}}$ allowing for correlation between effects associated with different datapoints.

The generalised Gauss–Markov problem is a type of nonlinear Gauss–Markov problem and can be solved using nonlinear least squares algorithms

(Section 4.3) using the Cholesky factorisation of $V_{\mathbf{y}}$ possibly in conjunction with a generalised QR factorisation. Because the number of computational steps for solving the generalised Gauss–Markov problem is generally of the order of m^3 , for large datasets this approach is computationally expensive. See below, however.

7.1 Structured generalised Gauss–Markov problems

As with the nonlinear Gauss–Markov problem, the uncertainty matrix $V_{\mathbf{y}}$ often has a structure that allows the generalised nonlinear Gauss–Markov problem to be solved more efficiently. Suppose the measurement model is

$$\mathbf{x}_i = \phi(\mathbf{u}_i, \mathbf{a}) + \mathbf{e}_i + K_i \mathbf{e}_0, \quad \mathbf{e}_i \in N(\mathbf{0}, V_i), \quad \mathbf{e}_0 \in N(\mathbf{0}, V_0),$$

where $\phi : \mathbb{R}^{p-1} \times \mathbb{R}^n \rightarrow \mathbb{R}^p$ is a parametric surface, \mathbf{e}_i represents random effects specific to the i th datapoint \mathbf{x}_i , and \mathbf{e}_0 represents random effects common to all the measurements. The matrix K_i represents the sensitivity of the i th set of measurements to these effects. If V_i has factorisation $V_i = B_i B_i^T$ and V_0 has factorisation $V_0 = B_0 B_0^T$, then the uncertainty matrix $V_{\mathbf{y}}$ associated with measurements $\mathbf{y} = (\mathbf{x}_1^T, \dots, \mathbf{x}_m^T)^T$ is given by

$$V_{\mathbf{y}} = B B^T, \quad B = \begin{bmatrix} B_1 & & B_{0,1} \\ & \ddots & \vdots \\ & & B_m \ B_{0,m} \end{bmatrix}, \quad B_{0,i} = K_i B_0,$$

and (15) can be written as

$$\min_{\mathbf{a}, \{\mathbf{u}_i\}, \{\mathbf{d}_i\}, \mathbf{d}_0} \frac{1}{2} \sum_{i=0}^m \mathbf{d}_i^T \mathbf{d}_i \tag{34}$$

subject to

$$\mathbf{x}_i = \phi(\mathbf{u}_i, \mathbf{a}) + B_i \mathbf{d}_i + B_{0,i} \mathbf{d}_0, \quad i = 1, \dots, m.$$

Holding \mathbf{a} and \mathbf{d}_0 fixed, it is seen that the optimal \mathbf{u}_i^* must solve the footpoint problem (31) but for the surface

$$\tilde{\phi}_i(\mathbf{u}_i, \tilde{\mathbf{a}}) = \phi(\mathbf{u}_i, \mathbf{a}) + B_{0,i} \mathbf{d}_0, \quad \tilde{\mathbf{a}} = \begin{bmatrix} \mathbf{a} \\ \mathbf{d}_0 \end{bmatrix}.$$

Following the same approach as described in Section 6, we define the generalised distance $d_i(\tilde{\mathbf{a}})$ as a function of $\tilde{\mathbf{a}}$ evaluated at the solution of the i th footpoint. Then (34) is equivalent to

$$\min_{\tilde{\mathbf{a}}} \left\{ \mathbf{d}_0^T \mathbf{d}_0 + \frac{1}{2} \sum_{i=1}^m d_i^2(\tilde{\mathbf{a}}) \right\}, \tag{35}$$

and can be solved using standard nonlinear least squares algorithms. This results in an algorithm that requires a number of steps linear in the number m of data points [20].

8 Robust least squares (RLS)

Robust approximation is a term used for data approximation methods that cope well with outlying or rogue datapoints. Standard least squares methods are known to provide estimates that can be significantly influenced by outlying data. We know that least squares methods correspond to maximum likelihood estimation for Gaussian sampling distributions. The Gaussian density function accords vanishingly small probabilities to values more than a few standard deviations from the mean. For this reason, a least squares model fit is forced to move the model value closer to an outlying datapoint, often at the expense of dragging the fit away from the rest of the data. As an alternative to least squares approximation, the L_1 approximation criterion is sometimes used: instead of minimising a sum of squares, the sum of absolute values is minimised:

$$\min_{\mathbf{a}} F_1(\mathbf{a}) = \sum_{i=1}^m |f_i(\mathbf{a})|.$$

L_1 approximation corresponds to maximum likelihood estimation of \mathbf{a} with respect to the double exponential distribution. Figure 2 graphs the PDFs associated with a Gaussian and double exponential distribution, both with standard deviation 1, and Figure 3 graphs the negative logarithm of the PDFs. The double exponential distribution has more probability mass in the tails but has a narrow peaked behaviour, very dissimilar to a Gaussian. The t -distribution $t_\nu(\mu, \sigma^2)$ can be regarded as a generalisation of a Gaussian distribution $N(\mu, \sigma^2)$ with an extra shape parameter $\nu > 0$ specifying the degrees of freedom. As $\nu \rightarrow \infty$, the distribution approaches the corresponding Gaussian but for low degrees of freedom ($\nu < 20$, say) the probability mass far from the mean is significantly greater than that for the corresponding Gaussian distribution; see Figures 2 and 3. The graphs suggest that the maximum likelihood estimation based on a t -distribution will have the good properties of least squares (making efficient use of all the data) but be tolerant of outlying data. We show below that it is possible to use least squares algorithms for this type of approximation, hence the term *robust least squares*.

The PDF for $t_\nu(\mu, \sigma^2)$ is such that

$$p_\nu(x|\mu, \sigma^2) \propto \left[1 + \frac{1}{\nu} \left(\frac{x - \mu}{\sigma} \right)^2 \right]^{-(\nu+1)/2}.$$

For a model of the form

$$y_i = \phi(\mathbf{x}_i, \mathbf{a}) + e_i, \quad e_i \in t_\nu(0, \sigma^2),$$

maximum likelihood estimates of \mathbf{a} are found [19, 23] by minimising

$$F_\nu(\mathbf{a}) = \sum_{i=1}^m (\nu + 1) \log \left[1 + \frac{1}{\nu} \left(\frac{y_i - \phi(\mathbf{x}_i, \mathbf{a})}{\sigma} \right)^2 \right]. \quad (36)$$




Fig. 2. Probability density functions for a Gaussian (solid), double exponential (dot-dash), and t -distribution with eight degrees of freedom (dotted); each have standard deviation 1.




Fig. 3. Negative logarithms of the PDFs graphed in Figure 2.

We can write $F(\mathbf{a})$ in (36) as a sum of squares of the form

$$F_\nu(\mathbf{a}) = \sum_{i=1}^m w_i^2(\mathbf{a}) f_i^2(\mathbf{a}), \quad w_i(\mathbf{a}) = w_\nu(f_i(\mathbf{a})), \quad f_i(\mathbf{a}) = \frac{y_i - \phi(\mathbf{x}_i, \mathbf{a})}{\sigma},$$

where the weighting function $w_\nu(x)$ is defined by

$$w_\nu^2(x) = (\nu + 1) \frac{1}{x^2} \log \left(1 + \frac{x^2}{\nu} \right), \quad \nu > 0.$$

Because $\log(1 + \epsilon) = \epsilon - \epsilon^2/2 + \epsilon^3/3 - \epsilon^4/4 + \dots$ $|\epsilon| \ll 1$, $w_\nu(x)$ is well defined at $x = 0$ with

$$w_\nu^2(x) = \frac{\nu + 1}{\nu} \left[1 - \frac{\zeta}{2} + \frac{\zeta^2}{3} - \frac{\zeta^3}{4} + \dots \right], \quad \zeta = x^2/\nu,$$

and $w_\nu^2(0) = 1 + 1/\nu$. A fourth-order approximation (i.e., quadratic in ζ above) should be accurate enough if $|x| < 10^{-3}$ and the machine precision is of the order of 10^{-16} . As ν increases, $w_\nu(x)$ approaches the constant function $w_\infty(x) = 1$.

8.1 Empirical implementation of RLS

The implementation of RLS considered above used the ‘heavy tail’ property of the t -distribution. The weighting function has the effect of giving less weight to data that are considered to be far from the corresponding model predictions. In an empirical implementation, we look for simple weighting functions with the same qualitative behaviour. One such example [36] is

$$w_\beta(x) = \frac{1}{(1 + \beta_1^2 x^2)^{\beta_2/2}}, \quad 0 \leq \beta_2 \leq 1, \quad \beta = (\beta_1, \beta_2)^T. \quad (37)$$

If $\beta_1 > 0$, then as $|x| \rightarrow \infty$, $w(x)|x|$ behaves as $\beta_1^{-\beta_2}|x|^{1-\beta_2}$. The parameter β_1 controls when the weighting function starts to take effect. If the expected standard deviation of the residuals not subject to outliers is σ then a value of $\beta_1 \approx 0.3/\sigma$ is appropriate. For such a value, the weighting function is close to 1 for residuals that can be explained by the expected noise in the data.

Figure 4 shows the standard LS fit and empirical RLS fit with $\beta_1 = 0.3/\sigma$ and $\beta_2 = 1$ for data by a cubic polynomial. The LS fit is dragged away from the majority of the data by the right-hand outlier. The RLS fit, on the other hand, effectively ignores all the outliers.

8.2 One-sided RLS

For some problems it may be known *a priori* that outliers have a particular sign. For example, in dimensional metrology, dust particles sometimes give rise to spurious measurements. Because the dust particles always have a positive diameter, their effect is one-sided. Some engineering surfaces have oil-bearing cracks which again will give rise to a one-sided effect. This behaviour can be incorporated empirically by defining

$$w_\beta(x) = \begin{cases} 1/[(1 + \beta_1^2 x^2)^{\beta_2/2}], & x > 0, \\ 1, & \text{otherwise.} \end{cases}$$

Figure 5 shows the standard LS fit and one-sided RLS fit of a circle to data with four outlying datapoints. The outliers have a significant influence on the LS fit but very little on the RLS fit.




Fig. 4. LS and empirical RLS fits of a cubic polynomial to data with outliers.

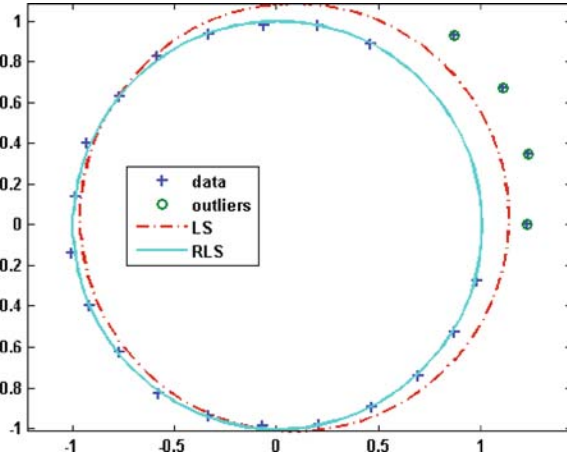


Fig. 5. LS and empirical RLS fits of circle to data with outliers.

8.3 RLS and the Huber M-estimator

One reasonably well-known robust estimation method is the Huber M-estimator [33]. Estimates of the parameters are found by minimising a function of the form

$$F_\gamma(\mathbf{a}) = \sum_{i=1}^m \rho_\gamma(f_i(\mathbf{a})),$$

where ρ_γ is defined as

$$\rho_\gamma(x) = \begin{cases} x^2/(2\gamma), & |x| \leq \gamma, \\ |x| - \gamma/2, & |x| > \gamma. \end{cases}$$

The objective function F_γ can also be written as a nonlinearly weighted sum of squares. Writing $\rho_\gamma(x) = w_\gamma^2(x)x^2$, so that

$$w_\gamma(x) = \begin{cases} 1/(2\gamma)^{1/2}, & |x| \leq \gamma, \\ (|x| - \gamma/2)^{1/2}/|x|, & |x| > \gamma, \end{cases}$$

then

$$F_\gamma(\mathbf{a}) = \sum_{i=1}^m \tilde{f}_{\gamma,i}^2, \quad \tilde{f}_{\gamma,i} = w_\gamma(f_i)f_i.$$

8.4 Algorithms for robust least squares

Given a nonlinear weighting function $w(x)$, let $\tau(x) = w(x)x$ and set

$$\tilde{f}_i(\mathbf{a}) = \tau(f_i(\mathbf{a})), \quad \tilde{F}(\mathbf{a}) = \frac{1}{2} \sum_{i=1}^m \tilde{f}_i^2(\mathbf{a}).$$

Even if f_i is linear in the parameters \mathbf{a} the introduction of the nonlinear τ function makes the minimisation of \tilde{F} a nonlinear least squares problem. To employ a Newton-type algorithm to minimise $\tilde{F}(\mathbf{a})$, we need to calculate

$$\mathbf{g} = \tilde{J}^T \tilde{\mathbf{f}}, \quad \tilde{J}_{ij} = \dot{\tau}_i \frac{\partial f_i}{\partial a_j}, \quad \dot{\tau}_i = \frac{d\tau}{dx}(f_i),$$

and

$$\tilde{H} = \tilde{J}^T \tilde{J} + \tilde{G}, \quad \tilde{G}_{jk} = \sum_i \tilde{f}_i \frac{\partial^2 \tilde{f}_i}{\partial a_j \partial a_k}.$$

We note that

$$\frac{\partial^2 \tilde{f}_i}{\partial a_j \partial a_k} = \ddot{\tau}_i \frac{\partial f_i}{\partial a_j} \frac{\partial f_i}{\partial a_k} + \dot{\tau}_i \frac{\partial^2 f_i}{\partial a_j \partial a_k}, \quad \ddot{\tau}_i = \frac{d^2 \tau}{dx^2}(f_i).$$

The first term on the right is the contribution due to the curvature in τ , the second due to that in F . Even if the second term is small, the first term is likely to be significant. This means that in practice the Gauss–Newton algorithm implemented for RLS will have significantly slower convergence than a Newton algorithm. However, if f is linear with $\mathbf{f} = \mathbf{y} - C\mathbf{a}$, the second term is zero and a Newton algorithm can be implemented easily with \tilde{J} and \tilde{G} calculated using the following identities,

$$\tilde{J}_{ij} = -c_{ij} \dot{\tau}_i, \quad \tilde{G}_{jk} = \sum_i \tau_i \ddot{\tau}_i c_{ij} c_{ik},$$

so that $\tilde{H} = C^T W C$, where W is the diagonal matrix with $W_{ii} = \dot{\tau}_i^2 + \tau_i \ddot{\tau}_i$.

9 Summary and concluding remarks

This chapter has been concerned with applying standard least squares methods to a range of regression problems. The algorithmic requirements of the approaches are modest: all the algorithms described in this chapter can be implemented using the Cholesky and QR factorisations. (The generalised QR factorisation employs the RQ factorisation which can be implemented in terms of a QR factorisation.) Both factorisations can be implemented in a few lines of software; library implementations are freely available in LAPACK, for example. The classes of regression problems considered in the chapter occur frequently in metrology, but are often solved using approximate methods or inefficiently. In the algorithms described here, the computations can be organised so that they require a number of computational steps linear in number of datapoints. We have also described robust approximation methods that represent a useful approach to coping with outlying data within a least squares framework.

Acknowledgements

This work was supported by the Department of Trade and Industry's Software Support for Metrology Programme.

The author is grateful to Dr Ian Smith, NPL, for comments on an earlier version of this paper, and for comments provided by the referees.

References

1. Ahn, S.J., Westkämper, E., W., R.: Orthogonal distance fitting of parametric curves and surfaces. In: J. Levesley, I.J. Anderson, J.C. Mason (eds.) *Algorithms for Approximation IV*, pp. 122–129. University of Huddersfield (2002)
2. Anderson, I.J., Cox, M.G., Forbes, A.B., Mason, J.C., Turner, D.A.: An efficient and robust algorithm for solving the footpoint problem. In: M. Daehlen, T. Lyche, L.L. Schumaker (eds.) *Mathematical Methods for Curves and Surfaces II*, pp. 9–16. Vanderbilt University Press, Nashville, TN (1998)
3. Barker, R.M., Cox, M.G., Forbes, A.B., Harris, P.M.: *Software Support for Metrology Best Practice Guide 4: Modelling Discrete Data and Experimental Data Analysis*. Tech. rep., National Physical Laboratory, Teddington (2007)
4. Bartholomew-Biggs, M., Butler, B.P., Forbes, A.B.: Optimisation algorithms for generalised regression on metrology. In: P. Ciarlini, A.B. Forbes, F. Pavese, D. Richter (eds.) *Advanced Mathematical and Computational Tools in Metrology IV*, pp. 21–31. World Scientific, Singapore (2000)
5. BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML: *Guide to the Expression of Uncertainty in Measurement*. Geneva, second edn. (1995)
6. Björck, A.: *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia (1996)

7. Boggs, P.T., Byrd, R.H., Schnabel, R.B.: A stable and efficient algorithm for nonlinear orthogonal distance regression. *SIAM Journal of Scientific and Statistical Computing* **8**(6), 1052–1078 (1987)
8. Boggs, P.T., Donaldson, J.R., Byrd, R.H., Schnabel, R.B.: ODRPACK: software for weighted orthogonal distance regression. *ACM Trans. Math. Soft.* **15**(4), 348–364 (1989)
9. Box, G.E.P., Tiao, G.C.: Bayesian inference in statistical analysis, Wiley Classics Library Edition 1992 edn. Wiley, New York (1973)
10. Cox, M.G.: Linear algebra support modules for approximation and other software. In: J.C. Mason, M.G. Cox (eds.) *Scientific Software Systems*, pp. 21–29. Chapman & Hall, London (1990)
11. Cox, M.G., Forbes, A.B., Fossati, P.M., Harris, P.M., Smith, I.M.: Techniques for the efficient solution of large scale calibration problems. Tech. Rep. CMSC 25/03, National Physical Laboratory, Teddington (2003)
12. Cox, M.G., Forbes, A.B., Harris, P.M., Smith, I.M.: Classification and solution of regression problems for calibration. Tech. Rep. CMSC 24/03, National Physical Laboratory (2003)
13. Cox, M.G., Harris, P.M.: SSJM Best Practice Guide No. 6, Uncertainty evaluation. Tech. Rep. DEM-ES-011, National Physical Laboratory, Teddington, UK (2006)
14. Dongarra, J.J., Moler, C.B., Bunch, J.R., Stewart, G.W.: LINPACK Users' Guide. Society for Industrial and Applied Mathematics, Philadelphia (1979)
15. Fletcher, R.: *Practical Methods of Optimization*, second edn. John Wiley and Sons, Chichester (1987)
16. Forbes, A.B.: Least squares best fit geometric elements. In: J.C. Mason, M.G. Cox (eds.) *Algorithms for Approximation II*, pp. 311–319. Chapman & Hall, London (1990)
17. Forbes, A.B.: Generalised regression problems in metrology. *Numerical Algorithms* **5**, 523–533 (1993)
18. Forbes, A.B.: Efficient algorithms for structured self-calibration problems. In: J. Levesley, I. Anderson, J.C. Mason (eds.) *Algorithms for Approximation IV*, pp. 146–153. University of Huddersfield (2002)
19. Forbes, A.B.: Asymptotic least squares and Student- t sampling distributions. In: P. Ciarlini, E. Felipe, A.B. Forbes, F. Pavese, C. Perruchet, B. Siebert (eds.) *Advanced Mathematical and Computational Tools in Metrology VII*, pp. 297–300. World Scientific, Singapore (2006)
20. Forbes, A.B.: Structured nonlinear Gauss-Markov problems. In: A. Iske, J. Levesley (eds.) *Algorithms for Approximation V*, pp. 167–186. Springer, Berlin (2006)
21. Forbes, A.B.: Surface fitting taking into account uncertainty structure in coordinate data. *Measurement Science and Technology* **17**, 553–558 (2006)
22. Forbes, A.B.: Uncertainty evaluation associated with fitting geometric surfaces to coordinate data. *Metrologia* **43**(4), S282–S290 (2006)
23. Forbes, A.B.: Least squares approaches to maximum likelihood estimation. Tech. Rep. DEM-ES-019, National Physical Laboratory (2007)
24. Forbes, A.B., Harris, P.M., Smith, I.M.: Generalised Gauss-Markov Regression. In: J. Levesley, I. Anderson, J.C. Mason (eds.) *Algorithms for Approximation IV*, pp. 270–277. University of Huddersfield (2002)
25. Garbow, B.S., Hillstom, K.E., Moré, J.J.: User's guide for MINPACK-1. Tech. Rep. ANL-80-74, Argonne National Laboratory, Argonne, IL (1980)

26. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: Bayesian Data Analysis, second edn. Chapman & Hall/CRC, Boca Raton, Fl. (2004)
27. Gill, P.E., Murray, W., Wright, M.H.: Practical Optimization. Academic Press, London (1981)
28. Golub, G.H., Van Loan, C.F.: Matrix Computations, third edn. John Hopkins University Press, Baltimore (1996)
29. Hammarling, S.: The numerical solution of the general Gauss-Markov linear model. Tech. Rep. TR2/85, Numerical Algorithms Group, Oxford (1985)
30. Helfrich, H.P., Zwick, D.: A trust region method for implicit orthogonal distance regression. Numerical Algorithms **5**, 535 – 544 (1993)
31. Helfrich, H.P., Zwick, D.: Trust region algorithms for the nonlinear distance problem. Num. Alg. **9**, 171 – 179 (1995)
32. Helfrich, H.P., Zwick, D.: A trust region algorithm for parametric curve and surface fitting. J. Comp. Appl. Math. **73**, 119–134 (1996)
33. Huber, P.J.: Robust Statistics. Wiley, New York (1980)
34. van Huffel, S. (ed.): Recent Advances in Total Least Squares and Errors-in-Variables Techniques. SIAM, Philadelphia (1997)
35. van Huffel, S., Vandewalle, J.: The Total Least Squares Problem: Computational Aspects and Analysis. SIAM, Philadelphia (1991)
36. Jenkinson, D.P., Mason, J.C., Crampton, A., Cox, M.G., Forbes, A.B., Boudjema, R.: Parameterized approximation estimators for mixed noise distributions. In: P. Ciarlini, M.G. Cox, F. Pavese, G.B. Rossi (eds.) Advanced Mathematical and Computational Tools in Metrology VI, pp. 67–81. World Scientific, Singapore (2004)
37. Lawson, C.L., Hanson, R.J.: Solving Least Squares Problems. Prentice-Hall, Englewood Cliffs, NJ (1974)
38. MacKay, D.J.C.: Information Theory, Inference and Learning Algorithms. Cambridge University Press, UK (2003)
39. MathWorks, Inc., Natick, MA: Using Matlab (2002). <http://www.mathworks.com>
40. Milton, M.J.T., Harris, P.M., Smith, I.M., Brown, A.S., Goody, B.A.: Implementation of a generalized least-squares method for determining calibration curves from data with general uncertainty structures. Metrologia **43**(4), S291–S298
41. Moré, J.J., Wright, S.J.: Optimization Software Guide. SIAM, Philadelphia (1993)
42. The Numerical Algorithms Group Limited, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR: The NAG FORTRAN Library, Mark 20, Introductory Guide (2002). <http://www.nag.co.uk/>
43. SIAM, Philadelphia: The LAPACK User's Guide, third edn. (1999)
44. Tarantola, A.: Inverse Problem Theory and Methods for Model Parameter Estimation. SIAM, Philadelphia (2005)
45. Watson, G.A.: Some robust methods for fitting parametrically defined curves or surfaces to measured data. In: P. Ciarlini, A.B. Forbes, F. Pavese, D. Richter (eds.) Advanced Mathematical and Computational Tools in Metrology IV, pp. 256–272. World Scientific, Singapore (2000)
46. Wilkinson, J.H., Reinsch, C.: Handbook of Automatic Computation Volume II: Linear Algebra. Springer-Verlag, Berlin (1971)

Frequency and Time–Frequency Domain Analysis Tools in Measurement

Pedro M. Ramos¹, Raul C. Martins¹, Sergio Rapuano², Pasquale Daponte²

¹ Instituto de Telecomunicações, Instituto Superior Técnico, DEEC, UTL, Av. Rovisco Pais 1, 1049-001 Lisbon, Portugal pedro.ramos@lx.it.pt, rcmartins@ist.utl.pt

² Department of Engineering, Università del Sannio, 82100 Benevento, Italy rapuano@unisannio.it, daponte@unisannio.it

Summary. A traditional approach to signal processing has been, for a long time, the frequency domain analysis, in which time or space periodicities can be identified. It is a relatively simple approach which usually carries significant information, suitable both as a first-step approach for further signal processing and for feature extraction. Because this approach carries no time information, frequency and time-domain analysis based on wavelets has become increasingly important. This shares a similar analytical approach making use of time-limited functions as the basis of the transform, allowing for space or time localization of short-lived repeating patterns. These are signal-processing tools which require some good understanding of the underlying theory to avoid common pitfalls and circumvent some limitations. Examples are given to show applicability.

Key words: Frequency analysis, Fourier transform, time–frequency analysis, wavelet transform, wavelet packets, wavelet networks

1 Introduction

This chapter deals with the frequency domain and time–frequency domain representation of signals. Its focus is centered on what is the physical meaning of these two very important domains in instrumentation and signal processing. The first three sections after the introduction concern the frequency domain representation and in the last four sections, the spotlight is aimed at the ever-growing field of time–frequency domain representation.

Time representation of continuous signals is the most traditional way to represent and show signals. However, in many applications (if not in the majority of applications), the time domain representation is insufficient to gain insight as to what constitutes a signal. For example, from the time domain representation (e.g., observed in an oscilloscope) a user can see the basic shape

of a signal, its amplitude, and eventually its frequency. However, some smaller disturbances or distortion of the signal are difficult to spot and even harder to classify and quantify. In these cases, the signal can be, for example, affected by the presence of noise or disturbances at particular frequencies (e.g., at the power grid frequency which affects so many poorly designed or badly conditioned circuits). Some of these situations can be solved by filtering, that is, the selective elimination or amplification of certain regions of the frequency domain. Waveform distortions are another field where the time domain representation of signals can be insufficient. In all these cases (and in many more), the frequency domain representation of signals is a crucial tool and the mathematical tools that enable its determination must be fully known to every researcher in the field, especially for instrumentation researchers.

Frequency domain representation works with the signal frequency spectrum, that is, with the signal expressed as a function of frequency. The spectrum is usually a complex function of frequency, therefore magnitude and phase spectra or real and imaginary parts of the spectrum are used. The possibility of decomposition of a periodic signal into the sum of components at different frequencies was consistently introduced by J. B. Fourier in the beginning of the 19th century [Fou78]. However, at that time, the ideas defended by Fourier met with a barrage of criticism and only later on was the credit for such a fundamental theory attributed to Fourier. The basic idea behind Fourier's theory can be stated: *Any periodic signal is composed of a superposition of pure sine waves, with suitably chosen amplitudes and phases, whose frequencies are harmonics of the fundamental frequency of the signal.* Although Fourier demonstrated how to determine the corresponding amplitudes and phases of the harmonics that constitute signals, it was only in 1829 that P. L. Dirichlet demonstrated the conditions under which the Fourier series converge [Dir29].

The expressions *spectral analysis* and *spectrum* are usually employed in the frequency domain representation of signals. These terms have their historic origin in the spectroscopic decomposition of light.

The frequency domain representation is suited for the stationary signals. However, in many situations the actual signal frequency composition changes with time and this must be measured. To solve this issue, a mixed time–frequency domain representation can be used. It can easily be recognized as the representation of the evolution along the time of the spectrum.

Basic time–frequency representations are the *short-time Fourier transform* (STFT), using the repeated *discrete Fourier transform* (DFT) calculation for the same window shifted along the sequence of samples, the chirplet transform and the wavelet transform together with their modifications. This chapter deals with DFT and FFT (*fast Fourier transform*, an algorithm for fast implementation of the DFT), wavelet transforms, chirplet transform, and wavelet networks.

2 Fourier Analysis

For a periodic continuous signal, the spectral representation is obtained from the *Fourier Series* (FS)

$$s(t) = \sum_{k=-\infty}^{+\infty} c_k e^{jk\omega_0 t}, \tag{1}$$

where $\omega_0 = 2\pi f = 2\pi/T$ is the angular frequency of the periodic signal and the complex coefficients c_k , obtained from

$$c_k = \frac{1}{T} \int_T s(t) e^{-jk\omega_0 t} dt \tag{2}$$

represent the spectral composition of the original periodic signal. Already clear from Equations (1) and (2) is the fact that only a finite number of elements are obtained for the spectral representation of the periodic signal. This is a direct cause of the periodic nature of the signal. Elements c_1 and c_{-1} represent the spectral composition of the signal at its frequency which is typically called the fundamental. The other terms, which correspond to frequencies that are integer multiples of the fundamental are called harmonics. There are other ways to represent the Fourier series, such as

$$\begin{aligned} s(t) &= \frac{a_0}{2} + \sum_{k=1}^{+\infty} a_k \cos(k\omega_0 t) + b_k \sin(k\omega_0 t) \\ &= A_0 + \sum_{k=1}^{+\infty} A_k \cos(k\omega_0 t + \varphi_k). \end{aligned} \tag{3}$$

All these representations are equivalent and one should use the one that best suits each specific problem. Obviously, there are direct relations that enable as to change between the different representations by combining the coefficients of each representation.

Convergence of the Fourier series for periodic signals is restricted to signals that verify three conditions known as the Dirichlet conditions:

- The signal $s(t)$ must be absolutely integrable; that is,

$$\int_T |s(t)| dt < \infty. \tag{4}$$

This condition guarantees that each coefficient c_k is finite.

- The signal must have a finite amount of maximums and minimums during its period.

- The final condition concerns the number of discontinuities in the signal period and states that, for a Fourier series to converge to the original signal, the number of discontinuities must be finite and also each discontinuity must be finite.

Note that, signals that do not meet these criteria, are usually designed specifically in order not to meet at least one of the conditions. These signals are rarely seen in real-world applications and so the conditions, although they are a formal requirement, are very rarely checked or used.

In instrumentation and measurement, the complete continuous signal only exists in the analog domain of the circuitry. Because the signal processing algorithms are in the digital domain, the analog signals must be converted into the digital domain. This transformation is done using *analog-to-digital converters* (ADCs). These devices are the cornerstone of the digital era and have been, for many years, one of the most relevant research and development activities of the manufacturers of integrated circuits. Although many different architectures for the implementation of the ADCs have been developed and implemented over the years, the basic trade-off of the ADCs still remains true: to have faster digitalisation, you lose amplitude resolution; better resolution can usually be achieved only for lower digitalisation speeds. These are the two most basic parameters in every ADC: its resolution (expressed in the number of bits of the digitised word used to represent each value of the input analog voltage) and its sampling rate (corresponding with sampling frequency f_s) which is expressed as the number of samples that the ADC can process per second.

The acquisition process changes two basic properties of the signal: first the signal is no longer continuous; it is now a discrete-time signal; secondly, only a finite number of samples (N) is available. The ADC also changes the amplitude of the signal from the continuous domain into a discrete domain due to the ADC resolution and input range. For discrete-time signals, the *Discrete Fourier Transform* (DFT) is

$$x_n = \frac{1}{N} \sum_{k=0}^{N-1} X_k e^{j(2\pi/N)kn} \quad \text{with} \quad X_k = \sum_{n=0}^{N-1} x_n e^{-j(2\pi/N)kn}. \quad (5)$$

In these equations, some important aspects of the DFT can be observed: (i) because only N samples are acquired, in the time domain, only N spectral values are obtained; (ii) because the spectral components of the signal are present in the form $e^{j(2\pi/N)kn}$ two such terms are required to represent a real-valued signal (one with positive frequency and another with negative frequency (this can also be seen from $\cos(a) = (e^{ja} + e^{-ja})/2$)); (iii) the maximum frequency available from the DFT result corresponds to $f_s/2$; (iv) each sample x_n can be time stamped with $t_n = n\Delta t$ and each spectrum element X_k has a frequency $\Delta f \times [0 \ 1 \ 2 \ \dots \ (N/2) \ (-N/2 + 1) \ \dots \ -2 \ -1]$; (v) in the time domain, the resolution is $\Delta t = 1/f_s$ whereas in the frequency domain, the resolution is $\Delta f = f_s/N$. This is a very important aspect of the DFT: the

frequency resolution is the reciprocal of the acquisition time $\Delta f = 1/(N\Delta t)$. To achieve a 1 Hz resolution spectrum, a 1 s long acquisition is required. To improve spectral resolution (i.e., reduce $\Delta f = f_s/N$) one must either increase Δt (reducing sampling rate and thus reducing time resolution) or increase N (acquiring more samples). Either way, both solutions require reacquisition of the time signal and can only be performed if the signal is stationary.

Another important aspect of the DFT is its relation with the Fourier Series equations (1) and (2). In fact, to assure a direct correspondence of the two, it is necessary that the set of acquired points in the DFT correspond exactly to an integer number of periods of the input signal. Only then can the integral be replaced with the summation.

The DFT algorithm as presented in Equation (5) is not computationally efficient. In the literature, it is stated that the number of operations required to process (5) directly is proportional to N^2 . This made the application of the DFT very troublesome for many years. In 1965, Cooley and Tukey published an efficient algorithm for the computation of the DFT [CT65]. Their algorithm, and others that have surfaced since, are now known as the *Fast Fourier Transform* (FFT) and some of the results used to improve the computational burden of the FFT have been traced back to Gauss in the beginning of the 19th century. The basic idea behind all of these algorithms is that, due to the symmetry and periodicity of the phasor $e^{-j(2\pi/N)kn}$ many of the original computations of the DFT were repeated. The number of operations in optimised FFT algorithms are now proportional to $N \log_2(N)$ (see [JF07] for the latest results). The original algorithm [CT65] required that the number of samples be a product of two integer numbers with the maximum efficiency obtained when the number of samples is a power of two. Later algorithms improved the computational efficiency requiring only that the number of samples be decomposed in the form $2^n 3^m 5^k$. However, these algorithms always depend on the implementation, either already available in algorithm packages included in commercial programs or on the implementation by the final user for its specific platform. In most commercial numerical analysis packages the determination of the DFT is called FFT. It is up to the user to input a number of samples that can be used in the FFT of the package. Otherwise, the package may resort to the slower DFT direct calculation.

A good, well-written textbook on these subjects is Oppenheim et al. [OSB99].

3 How to use the FFT

The DFT as presented in Equation (5) is periodic with N . This means that replicas of the spectrum will appear centered on multiples of the sampling frequency. If the sampled signal contains frequencies higher than half of the sampling frequency these will overlap with the spectrum centred on neighbouring multiples of the sampling frequency. A corollary of this periodicity

is that a bandlimited signal must be sampled at a rate at least twice the highest frequency contained in the signal so as to avoid overlapping. This also becomes evident from a theorem first enunciated by Nyquist in 1928 [Nyq28] and then proved by Shannon in 1949 [Sha49] which states that any bandlimited signal can be uniquely determined by its samples as long as the sample rate is at least twice that of the highest frequency found in the signal. The highest frequency of the signal is usually referred to as the *Nyquist frequency* and twice this, which is the frequency that must be exceeded by the sampling rate, is commonly called the *Nyquist rate*.

3.1 Aliasing

Consider now a bandlimited signal consisting of only a sine wave. If this signal is sampled at a rate 1.9 times higher than its frequency, upon use of Equation (5) what will appear will be only one frequency, which is in accordance with the fact that the signal is a sine wave, but with a frequency 0.9 times the frequency of the signal. The apparent frequency of the signal is called the alias and this result of subsampling is commonly referred to as *aliasing* (see Figures 1 and 2). In Figure 1 it is possible to view the effect of aliasing in the time domain and in Figure 2 in the frequency domain for different waveforms. Figure 1 corresponds to the acquisition output of one period of a synthesized waveform sampled at three different rates. It is easy to visually perceive from Figure 1 that the frequency of the highest component is not the same when the signal is subsampled, appearing in this case to be much lower than it actually is.

The presence of alias frequencies within a spectrum can be detected by comparing two spectra of the same signal sampled at slightly different rates. If some of the frequencies that make up the signal change, upon application of Equation (5), then those that change are being subsampled and are therefore a result of aliasing. Subsampled signals can be reconstructed as long as (i) they are periodic, (ii) there is some idea of the number of times their frequency is higher than half of the sampling rate, and (iii) there is no superposition with other tones.

3.2 Spectral leakage

Given that the spectrum resolution is $\Delta f = f_s/N$, if the sampled signal contains a frequency which is not an integer multiple of this resolution its energy cannot be accounted for by a single bin. It spreads over the bin nearest to the frequency of the signal and neighbouring bins. This is known as leakage, an example of which is shown in Figure 2. Another way of looking at this is perceiving that inherent to Equation (5) is the periodicity of the acquired interval and, as such, to avoid discontinuities between these virtual periods care must be taken to ensure that within the acquired interval there is an integer number of periods of the signal.




Fig. 1. Example of subsampling in the time domain for a multitone signal. Represented is one period of the signal $0.9 \sin(2\pi t) + \sin(2\pi 25t)$, the points that would be acquired if this signal were sampled at 20 Hz ($f_{\max} < f_s/2$) and the aliased signal corresponding to the undersampled points.

3.3 Windowing

Usually it is hard, nearly infeasible, to guarantee such a condition, especially if the signal is multitonal with frequencies whose ratio is not a rational number. In such cases it is common practice to multiply, point by point, the acquired sample vector with a function which is zero, or some constant, at both end-points. This allows for the smooth tapering at each end of the sample vector, considerably reducing the leakage, albeit usually at the expense of lower frequency resolution. This is known as *time windowing*. The resolution decrease can be better perceived by looking at the analytical description of windowing. If we consider a window function, described in the time domain by the sample vector w and in the frequency domain by the vector W , then Equation (5) can be rewritten as

$$X_k = \sum_{n=0}^{N-1} x_n w_n e^{-j(2\pi/N)kn} = X \star W, \tag{6}$$

where \star corresponds to the convolution operator. The rightmost term of the second equality states that the wider the main lobe of the Fourier transform of the window function is, the wider will also be the lines associated with the tones, effectively decreasing the capability to distinguish close tones. When a continuous signal is acquired for a limited time interval, a time window is implicitly used: the rectangular window. The Fourier transform of this window is a sine function, $\sin(\pi x)/x$, whose main lobe is usually much narrower than the lobes associated with the most commonly used time windows, such as the




Fig. 2. Graphical representation of subsampling and leakage for a multitone signal. The original signal has four frequency components. The first three components are acquired without spectral leakage. The fourth component has spectral leakage because the number of periods is not an integer as shown in the inset.

Bartlett (triangular), Hamming, Hanning, and Blackman. This is why usually the spectrum resulting from the use of time windows has lower frequency resolution. The decrease in the leakage stems from the fact that the peak-sidelobe level (the ratio between the peak of the side lobes and that of the main lobe) is usually much lower in the rectangular window than in the aforementioned common time windows. The most important properties of the time window to look for are (i) the width of its main lobe, (ii) the peak-sidelobe level, and (iii) the sidelobe roll-off which is usually a trade-off with the peak-sidelobe level but, naturally, should be as high as possible.

4 Example of spectral analysis application using FFT

The Fourier analysis is commonly used for measurement of two parameters describing signal quality: Total Harmonic Distortion (THD) and Signal-to-Noise Ratio (SNR). It can also be applied to the characterisation of ADCs. Albeit necessarily limited in scope, the concepts dealt with in this example can be readily extended to a huge number of applications.

The experimental procedure consists in generating a spectrally pure sine wave which will then be digitized by the ADC. The frequency content of the sample vector will give an idea of the ADC distortion, its effective resolution, and its noise. These metrics, however, are frequency dependent and as such should be carried out for several different frequencies within the ADC bandwidth.

The first step that needs to be taken is to ensure the amplitude of the sine wave does not exceed about 95% of the ADC input range and is centred. This is to avoid clipping which would severely bias the distortion metrics. The second step should be to make sure the ratio between the signal frequency and the acquisition rate is a noninteger rational number. Ideally the greatest common divisor of the two should be as small as possible because the inverse, commonly referred to as the *Unit Test Period* (UTP), gives the time interval along which there is no information redundancy; that is, all samples correspond to different phases of the signal, $UTP = 1/\text{gcd}(f, f_s)$, where gcd stands for *greatest common divisor*.

This is an important condition to guarantee for three main reasons:

- (i) It maximizes the probability of all codes of the ADC being stimulated, within the input range.
- (ii) It minimizes leakage because the frequency resolution of the spectrum will be equal to $\Delta f = \text{gcd}(f, f_s)$.
- (iii) It makes it easier to assess the harmonic distortion.

The number of samples to acquire is now a simple matter to determine: $N = UTP \cdot f_s$. Even though this configuration minimizes leakage, there is no guarantee leakage won't occur. One way to keep its effects to a minimum is to manipulate the length of the acquired vector. This can be done automatically by assessing the approximate number of points per period, f_s/f , and then keeping the original sample vector evaluating the DFT (5) each time a sample is removed, up to a maximum of f_s/f points. The DFT at which the amplitude is at a maximum corresponds to the sample vector length of minimal leakage. This method, albeit simple and efficient, is suitable only when the main tone has no neighbours with significant amplitude.

For the purpose of ADC characterisation it is enough to consider only one side of the spectrum and, as such, the second half of the vector returned from the DFT can be discarded. The first half is kept and all of its elements are multiplied by 2, except for the first (DC component) and the last if the length of the sample vector is even.

The total harmonic distortion of the ADC is measured by dividing the power of the harmonics by the power of the fundamental of the digitised signal in the ADC output assuming an ideal analog sine wave at the input. This means the bins associated with those frequencies must be identified, $\text{bin}(f_{1H}) = f_{1H}/\Delta f + 1$ and $\text{bin}(f_{xH}) = x f_{1H}/\Delta f + 1$ where the (+1) terms are required only if the indexing of the sample vector starts at 1 (e.g., MATLAB) and f_{xH} corresponds to the frequency of harmonic x . The com-

putation of the total harmonic distortion, computed up to the n^{th} harmonic, is now

$$THD_{dB} = 10 \log_{10} \frac{\sum_{x=2}^n X_{xH}^2}{X_{1H}^2}. \quad (7)$$

where X_{xH} is the spectral amplitude of harmonic x . The n^{th} harmonic to be evaluated in Equation (7) is usually chosen in accordance with the sampling rate and the number of harmonics to consider [IEE01]. When assessing the performance of the ADC for frequencies near half the maximum sampling rate subsampling usually takes place and the identification of the harmonics is still possible. In this case, due to the spectrum repetition that takes place at multiples of the sampling rate, it is fairly simple to determine the alias frequencies

$$f_{xH}(\text{alias}) = \begin{cases} f_s/2 - \text{mod} \left(f_{xH}, \frac{f_s}{2} \right) \leftarrow \text{floor} \left(\frac{f_{xH}}{\frac{f_s}{2}} \right) & \text{is odd} \\ \text{mod} \left(f_{xH}, \frac{f_s}{2} \right) \leftarrow \text{floor} \left(\frac{f_{xH}}{\frac{f_s}{2}} \right) & \text{is even.} \end{cases} \quad (8)$$

When the subsampling is associated with frequency components of the input signal, care should be taken not to remove them in the analog frontend by an anti-alias filter. This is not the case, however, in the previous example of distortion due to the converter.

Another common parameter in the characterisation of ADCs is the signal-to-noise ratio which corresponds to the ratio of the signal power to the power of noise in the digitised signal at the output of the ADC, assuming an ideal sine wave at the input,

$$SNR_{dB} = 10 \log_{10} \frac{X_{1H}^2}{\sum_{k \neq xH} X_k^2}. \quad (9)$$

Still another related parameter is the Signal-to-Noise And Distortion ratio (*SINAD*),

$$SINAD_{dB} = 10 \log_{10} \frac{X_{1H}^2}{\sum_{k > 0, k \neq 1H} X_k^2}. \quad (10)$$

From this last metric the Effective Number Of Bits (*ENOB*) of the ADC can be determined. This metric tries to assess the resolution of an ideal converter which would generate the same level of distortion as the ADC under test. The distortion in that case would only be attributable to the quantisation. Because for an ideal ADC the $SINAD_{dB} = 6.02n_{bit} + 1.76$, then

$$ENOB = \frac{SINAD_{dB} - 1.76}{6.02}. \quad (11)$$

5 Wavelet transform

Transients are generally characterised by a short duration when compared to the observation interval. Their presence in a stationary waveform makes the resulting signal exhibit time-varying properties; that is, its instantaneous spectral contents vary with time. Transients are encountered in various fields such as audio signals, sonar, radar, medical signal analysis, ultrasonic signals, power quality analysis, and computerised vision of moving objects. Very important measurement information is often associated with some kinds of transients. As an example, the measurement of the time interval between ultrasonic echoes coming back from a multilayer structure allows the evaluation of the thickness of each layer; the estimation of transient phenomena (notches, spikes, dips) superimposed on the power line voltage offers the opportunity of monitoring the quality of power.

On the contrary, disturbing sources sometimes produce undesired transients which corrupt the waveform to be analysed. The removal of these transients is necessary for:

- Carrying out accurate measurements on the waveform itself
- Proceeding to a deeper analysis in order to identify probable sources and causes

This is the case, for example, for the transients which occur in high voltage tests and corrupt the desirable shape of the impulses used in these tests. The measurement of a transient requires the evaluation of significant parameters such as duration, amplitude, period of oscillation, and so on; sometimes, it may also be necessary, for identification and classification purposes, to recover the time domain shape of the transient by extracting it from the waveform on which it is superimposed.

Until now, some detection schemes based on both analog and digital solutions have been proposed. The analog solutions exploit sophisticated triggering circuits in support of specialised transient digitizers. The digital solutions, on the other hand, provide for suitable algorithms tuned for the detection of transients of known shapes and unknown arrival times. The possibility of detecting, measuring, and, eventually, classifying transients at the same time and in an automatic way is still being investigated, and could be carried by implementing measurement methods based on the *Wavelet Transform* (WT). This is a powerful theory allowing the processing of transients simultaneously in time and frequency domains. In particular, the WT can assure (i) very good time resolution at high frequencies and good frequency resolution at low frequencies and (ii) reduction of the influence of random noise affecting the signals being analysed. Transient arrival times and duration are automatically evaluated by exploiting the good time resolution at high frequencies assured by the *Continuous Wavelet Transform* (CWT), a particular WT implementation. On the other hand, the decomposition of the signal spectrum in frequency subbands and the subsequent reconstruction in the time-domain of

those subbands containing the transient are performed by means of another WT implementation, the *Discrete Time Wavelet Transform* (DTWT).

In the following, attention is principally put on the WT in its different forms. All the most important theoretical relations are shown along with some guidelines for their practical implementations.

6 The wavelet transform: Theoretical background and implementation

The *Wavelet Transform* acts as a sort of mathematical microscope through which different parts of the signal under analysis may be examined by adjusting the focus. The WT can be seen as the correlation between the signal and a set of functions that are small waves, called wavelets. Each wavelet, also called a daughter wavelet, is generated by scaling and translating one original wavelet, called the mother wavelet or basic wavelet. Scaling implies that the mother wavelet is either dilated or compressed and translation implies shifting of the mother wavelet in the time domain (Figure 3). In the following, two typical implementations of the WT are described: *continuous wavelet transform* and *discrete time wavelet transform*.

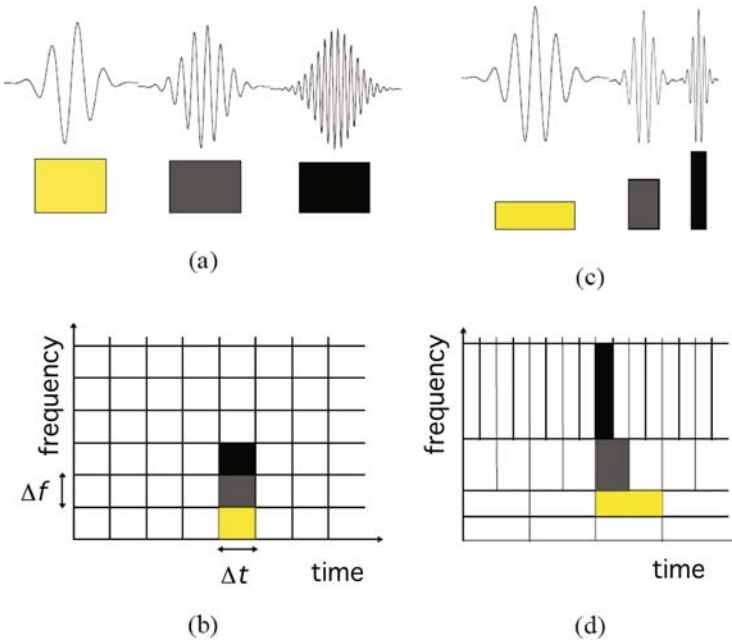


Fig. 3. Basis functions and time frequency resolution of traditional time–frequency representations (a,b) and of the wavelet transform (c,d).

6.1 Continuous wavelet transform

Definition

The expression of the continuous wavelet transform for real signals is

$$\text{CWT}(a, b) = \frac{1}{\sqrt{|a|}} \int h^* \left(\frac{t - b}{a} \right) \cdot s(t) dt, \tag{12}$$

where $h^*(t)$ denotes the complex conjugate of the mother wavelet $h(t)$, $s(t)$ is the signal to be transformed, and a and b are the dilation and translation parameters, respectively. This relation is defined on the open (b, a) half-plane ($b \in \mathbb{R}, a > 0$). The scale parameter is proportional to the reciprocal of frequency; the translation parameter stands for time.

If $h(t)$ is defined as

$$h_{a,b}(t) = a^{-1/2} h \left(\frac{t - b}{a} \right), \tag{13}$$

then Equation (12) can be written as a scalar or inner product of the real signal $s(t)$ with the function $h_{a,b}(t)$

$$\text{CWT}(a, b) = \int_{-\infty}^{+\infty} h_{a,b}^*(t) s(t) dt \text{ or } \text{CWT}(a, b) = \langle h_{a,b}^*(t), s(t) \rangle. \tag{14}$$

Implementation

For practical implementation a discretised version of the original CWT Equation (12) must be considered

$$\text{CWT}(a, i\Delta t) = \Delta t \frac{1}{\sqrt{a}} \sum_{n=0}^{N-1} h^* \left[\frac{(n - i)\Delta t}{a} \right] s(n\Delta t), \tag{15}$$

where N is the number of samples in the signal and Δt is the sampling interval. From this equation, it can be seen that at all values of the scale parameter a , a full set of N samples is generated. Although the time variable is discretised with uniform spacing, the parameter a could be discretised with a uniform or dyadic spacing. When the dyadic spacing is used, Equation (15) defines the discrete time wavelet transform. However, as in this section, Equation (15) is used as an approximation of the CWT, the notation CWT is used here. The scale parameter a is expressed in dyadic form as the series

$$a = 2^{j+m/M}, \tag{16}$$

where j indicates the octave number, m the voice number ($0 < m < M$), and M is the number of voices per octave. The number of octaves J to be adopted for CWT computation is the greatest integer satisfying the relation

$$J = \log_2(N) - 1, \tag{17}$$

and the octave number j ranges from 2 to $(J + 2)$. Moreover, the value of M generally ranges from 0 to 12; the greater M the better the frequency resolution. The direct implementation of Equation (15), especially for a number of voices per octave greater than 8, would require an extensive amount of computation. For this reason, an efficient algorithm, based upon the FFT, may be adopted [BM94]. By analysing the CWT in the Fourier transform domain and referring to the theorem of Parseval, the basic convolution operation of the CWT can be achieved by simple multiplication operations. As a matter of fact, writing the CWT in the Fourier domain gives

$$\mathfrak{F}\{CWT(a, b)\} = a^{-1/2} \cdot H^*(a\omega) \cdot S(\omega), \tag{18}$$

where $\mathfrak{F}\{CWT(a, b)\}$ is the Fourier transform of the CWT, $H(\cdot)$ is the Fourier transform of the mother wavelet, and $S(\cdot)$ is the Fourier transform of the input signal. Equation (18) can be represented graphically as shown in Figure 4

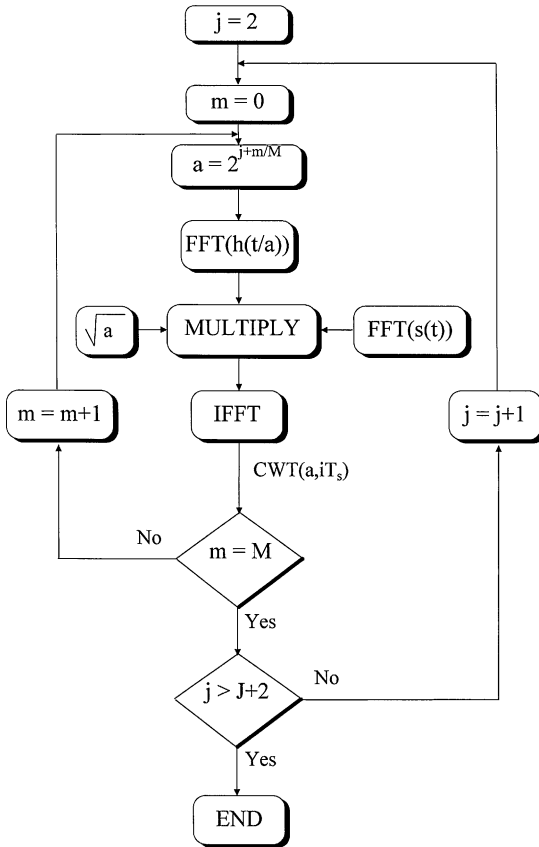


Fig. 4. Discrete WT implementation scheme.

after precalculation of the FFT of the mother wavelet and signal, the CWT is implemented via repeated scale, multiply, and inverse FFT (IFFT) operations [BM94].

Mother wavelet features

There are some conditions that must be met for a function to qualify as a mother wavelet; the function $h(t)$ is said to be a mother wavelet if and only if its Fourier transform $H(\cdot)$ satisfies

$$\int_0^{+\infty} \frac{|H(\omega)|^2}{\omega} d\omega = \int_{-\infty}^0 \frac{|H(\omega)|^2}{\omega} d\omega = C_H < +\infty. \tag{19}$$

In other words, the mother wavelet $h(t)$ must be oscillatory and have amplitude that quickly decays to zero. Furthermore, a mother wavelet $h(t)$ must have d vanishing moments, if and only if, for all nonnegative integers $q < d$, it satisfies

$$\int_{-\infty}^{+\infty} t^q h(t) dt = 0. \tag{20}$$

Each mother wavelet must have at least one vanishing moment; this condition implies that

$$\int_{-\infty}^{+\infty} h(t) dt = 0, \tag{21}$$

and $h(t)$ must therefore be a zero mean value function. Examples of mother wavelets are (Figure 5 upper):

1. Modulated Gaussian (Morlet)

$$h(t) = e^{j\omega_0 t} \cdot e^{-t^2/2}. \tag{22}$$

2. Second derivative of a Gaussian (Mexican hat)

$$h(t) = (1 - t^2) \cdot e^{-t^2/2}. \tag{23}$$

3. Haar

$$h(t) = \begin{cases} +1 & 0 \leq t < 1/2 \\ -1 & 1/2 \leq t < 1 \\ 0 & \text{otherwise.} \end{cases} \tag{24}$$

4. Shannon

$$h(t) = \frac{\sin(\pi t/2)}{\pi t/2} \cdot \cos\left(\frac{3\pi t}{2}\right). \tag{25}$$

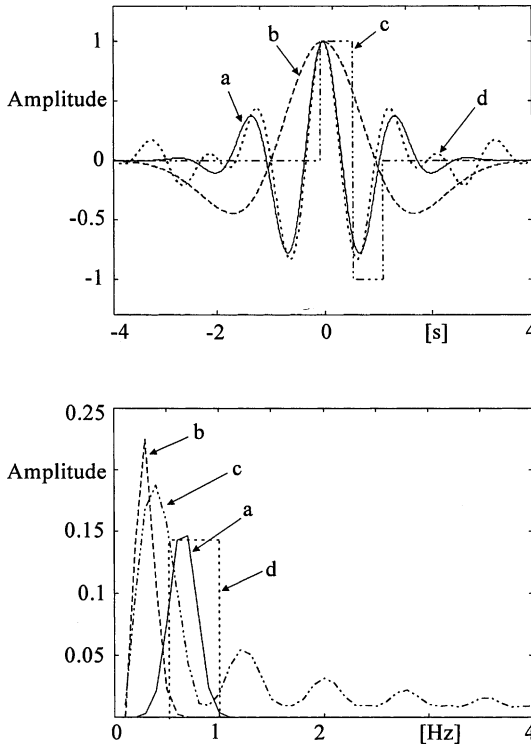


Fig. 5. Time domain shapes (upper) and frequency spectra (lower) of the following mother wavelets: Morlet (curve a), Mexican hat (curve b), Haar (curve c), and Shannon (curve d).

All the aforementioned characteristics allow the mother wavelet to be thought of as a bandpass filter (Figure 5 lower). Furthermore, unlike sines and cosines, individual mother wavelets are quite localised in time; simultaneously, as with sines and cosines, they are localised in frequency. This allows the WT to associate the nonstationary time-domain signal with a representation that is localised not only in frequency but also in time. Such an approach results in a more natural description of the signal under analysis thus giving the opportunity of extracting details and information from it [RV91, Dau92].

CWT modulus maxima properties

The local maxima of the CWT modulus correspond to the sharpest variation points of the signal. Let $h(t) = d\phi(t)/dt$ where $\phi(t)$ is a function whose integral is equal to 1. The CWT modulus at a fixed scale a satisfies

$$\begin{aligned}
 |\text{CWT}(a, b)| &= \left| \sqrt{a} \int \frac{d}{db} \phi^* \left(\frac{t-b}{a} \right) \cdot s(t) dt \right| \\
 &= \left| \frac{d}{db} \left[\sqrt{a} \int \phi^* \left(\frac{t-b}{a} \right) \cdot s(t) dt \right] \right|.
 \end{aligned} \tag{26}$$

The CWT modulus is proportional to the first derivative of $s(t)$ smoothed by $\phi(t)$. The local maxima of the CWT modulus thus correspond to the local maxima of the derivative of $\langle s(t), \phi_{a,b}(t) \rangle$, which are the sharpest variation points of the signal smoothed at the scale a . Furthermore, for a fixed b

$$|\text{CWT}(a, b)| = k(a)^\alpha, \tag{27}$$

where α is the so-called Lipschitz exponent, and k is a constant. It follows that:

- The CWT modulus local maxima decrease when the scale decreases (frequency increases) for discontinuities characterised by positive α .
- The CWT modulus local maxima does not decrease when the scale decreases for discontinuities characterised by nonpositive α .

From these considerations one can discriminate the signal from the noise by looking at the behaviour of the CWT modulus local maxima across scales. As a matter of fact, because actual signal discontinuities characterised by positive Lipschitz exponents and random noise is almost everywhere singular with negative Lipschitz exponents (one can prove that a white noise has a uniform Lipschitz exponent equal to $-1/2$), the CWT of the signal decreases when the scale decreases; on the contrary, the CWT of the noise increases, on average, when the scale decreases. In order to compute the evolution of CWT amplitude across scales, it is possible to relate each modulus maximum at the given scale to a modulus maximum at the successive and lower scale which is as close as possible and with the same sign, thus constructing some sequences called *chains*. At this point one can remove any chain of the modulus maxima whose amplitude increases, on average, when the scale decreases and retaining those chains whose amplitude decreases when the scale decreases; the former are related to a discontinuity which is most influenced by the noise, and the latter are related to an actual signal discontinuity [ADD99, LAD99].

It is worth noting that extensions of the CWT were proposed; for example, in [SML96] an extension based on a moving and scalable localizing Gaussian window was proposed. This offers some desirable characteristics that are absent in the CWT.

6.2 Discrete time wavelet transform

DTWT definition and implementation as bank of filters

The discrete time wavelet transform is defined as

$$\text{DTWT}_{m,n} = \frac{1}{\sqrt{a_0^m}} \sum_{k=0}^{N-1} h(a_0^{-m}k - nb_0)s(k), \tag{28}$$

where $a_0 \neq 1$, N is the number of samples, and n is the discrete-time variable.

As seen before, assuming that a_0 and b_0 are integers, the smallest possible values for a_0 and b_0 ($b_0 = 1$ and $a_0 = 2$) give the condition that all the scale factors are a power of two. This definition is a discrete approximation of the CWT and therefore can be obtained from Equation (15) as described in Section 6.1 with a very high number of discrete scales. However, it can also be obtained by means of analytical relations which regulate a digital filter bank, provided that the filter impulse responses satisfy some conditions described below. In particular, the whole DTWT implementation scheme is, generally, subdivided into two stages: decomposition and reconstruction. The DTWT decomposition is based on two digital filters: a highpass $h_a(\cdot)$, called the discrete mother a wavelet, and its lowpass mirror version $g_a(\cdot)$. The decomposition is carried out by arranging these filters in a tree structure as shown in Figure 6a which, in particular, reproduces a tree structure allowing a seven subband decomposition for the signal $s(\cdot)$ to be analysed. Each step of this scheme consists of the same digital filters. At each step, the input signal is simultaneously low-pass and high-pass filtered. Both the filter outputs are decimated by a factor of 2; then the low-pass filter output is sent to the next step where it is processed in the same way. The combination of the impulse responses of the encountered filters produces as a result, at the points A, B, C, D, E, and F six bandpass filters and at the point G a lowpass filter. All these filters are characterised by the same relative bandwidth (ratio between frequency bandwidth and centre frequency), but the filter obtained at the point A has the centre frequency, and consequently the frequency bandwidth, double with respect to that of the filter obtained at the point B and so on for the filters obtained at the points B and C. The results of the decomposition at the points A, B, C, D, E, and F can be expressed as

$$\text{DTWT}_{m,n} = \sum_{k=0}^{N-1} h_{m,2^{m+1}n-k} s_k, \quad m = 0, \dots, 5, \quad (29)$$

and at the point G as

$$\text{DTWT}_{m,n} = \sum_{k=0}^{N-1} h_{m,2^m n-k} s_k, \quad m = 6, \quad (30)$$

where $h_{m,n}$, for $m = 0, \dots, 6$, are the impulse responses of the seven filters. As can be seen, Equations (15) and (28) to (30) are very similar, except for the conjugation of the discrete mother wavelet function. However, it should be noted that the dyadic decomposition obtained by using the filter bank is incomplete, considering only a branch of the tree at each step.

Similar considerations are valid for the DTWT reconstruction stage, where $h_s(\cdot)$ and $g_s(\cdot)$ are the highpass and lowpass filters, respectively (Figure 6b) where an interpolation by factor 2 is used at each step. Each input signal for




Fig. 6. (a) DTWT decomposition stage; (b) DTWT reconstruction stage. The digital filter $h_a(\cdot)$, also called the discrete mother wavelet, is generally highpass. In orthogonal hypothesis, $g_a(\cdot)$ is the lowpass mirror version of $h_a(\cdot)$, and $h_s(\cdot)$ and $g_s(\cdot)$ are identical to their counterparts, within time-reversal.

the reconstruction stage at the points A, B, C, D, E, F, and G gives its own contribution to the output signal. In particular, each contribution represents the time domain reconstruction of the contents of the related frequency subband.

Discrete mother wavelets’ features

The scheme of Figure 6 is said to be a perfect reconstruction tree if the output signal obtained reconstructing the contents of all subbands is identical to the input signal, within a possible time shift. Perfect reconstruction properties are assured if the filters $h_a(\cdot)$, $g_a(\cdot)$, $h_s(\cdot)$, and $g_s(\cdot)$ satisfy either biorthogonal or orthogonal constraints. The biorthogonality constraints on these filters, also requires that

1. The overall scheme be an identity system or that

$$z^{-2k} H_a(z) H_s(z) = \begin{cases} 1 & k = 0 \\ 0 & k \neq 0 \end{cases}, \quad k \in \mathbb{Z}, \quad (31)$$

and

$$z^{-2k} G_a(z) G_s(z) = \begin{cases} 1 & k = 0 \\ 0 & k \neq 0 \end{cases}, \quad k \in \mathbb{Z}, \quad (32)$$

where $H_a(z)$, $G_a(z)$, $H_s(z)$, and $G_s(z)$ are the z -transform system functions of the filters

2. The information in one branch of each step be independent of the information in the other branch

This last condition, which assures no redundancy, can be expressed in terms of the filters as

$$z^{2k}H_a(z)G_s(z) = 0, \quad k \in \mathbb{Z}, \quad (33)$$

and

$$z^{2k}G_a(z)H_s(z) = 0, \quad k \in \mathbb{Z}. \quad (34)$$

The conditions in Equations (31) to (34) can be imposed directly onto the filter coefficients for implementation purposes thus giving the following relations

$$g_s(j) = (-1)^j h_a(j), \quad j = 0, \dots, L_h - 1, \quad (35)$$

$$h_s(j) = (-1)^{j+1} g_a(j), \quad j = 0, \dots, L_g - 1, \quad (36)$$

where L_h and L_g indicate the length of the filters $h_a(\cdot)$ and $g_a(\cdot)$, respectively. Arbitrary length linear phase filters can be designed. In orthogonal hypothesis, a special case of biorthogonality, a further relation must be satisfied

$$g_a(\cdot) = (-1)^j h_a(L - j - 1) \quad j = 0, \dots, L - 1, \quad (37)$$

where $h_a(\cdot)$ and $g_a(\cdot)$ have the same length $L = L_h = L_g$. This condition simplifies filter design (the highpass and lowpass filters in Figure 4 are practically identical); however, some nice properties have to be relaxed such as linear phase [AH92, VH92, Rio93]. The Daubechies filters are good examples of orthogonal discrete mother wavelets [Dau92] and some biorthogonal examples are given in [VH92].

DTWT properties

A significant and very useful property of the DTWT is the multiresolution subband decomposition (analysis) and reconstruction (synthesis). Its basic concept is to divide the signal spectrum into subspectra or subbands and, then, to treat individually those subbands more useful for the purposes at hand [RV91, Dau92, AH92]. Referring to Figure 6, the result of the reconstruction of the signal obtained at the point G from the decomposition stage can be regarded as a lowpass filtered (smoothed) version of the input signal $s(\cdot)$, and a first approximation of $s(\cdot)$. The result of the reconstruction of the signal obtained at the point F is a bandpass filtered version of $s(\cdot)$ and adds slightly more detail to the output. The same is true for the results of the reconstructions of the signals obtained the points E, D, C, and B. Finally, the result of the reconstruction of the signal obtained at the point A adds

the highest-frequency detail to the output. Each reconstruction result is referred to as a multiresolution component and contributes increasingly finer detail to the output of the synthesis stage. In practice, this DTWT property can be exploited in several application fields. In data compression, for example, the multiresolution components which contribute as much detail as desired are stored or transmitted; the higher-resolution components can be omitted or coded with fewer bits. The DTWT is a powerful tool for removing noise overlapping to signals in the same frequency band. In fact, due to the WT properties, the main part of the noise power is translated to the higher frequencies and separated from the signal. Therefore the noise can be easily highpass filtered out [PL06]. With regard to the analysis of transient signals, an occurred transient can be separated from the stationary waveform on which it can be superimposed thanks to their different and disjointed spectral contents [ADDT98].

An example of filter bank-based DTWT application: Extracting transient waveforms from distorted powerline signals

Power quality monitoring and analysis must be able to detect, locate, estimate, and classify disturbances on the supply lines. As a consequence, it must be supported by suitable measurement methods and systems.

The authors of paper [ADDT98] propose a procedure based on the multiresolution signal decomposition and reconstruction obtained by means of the DTWT. Such a procedure is able to detect, locate, and estimate the peak-to-peak amplitude of a disturbance. The number of subbands, to be used for signal decomposition, is chosen in such a way that the signal at the fundamental frequency f_f is included in the middle of the lowest-frequency subband, in order to limit the effects of fundamental spectral contents on the other subbands. The sampling frequency f_s can be maintained fixed in order to allow an easier implementation of the algorithm in a suitable hardware structure.

The original signal is decomposed in p frequency subbands by means of the DTWT analysis stage, arranging a tree digital filter structure as reported in [ADDT98]. In particular, the Daubechies filter with 16 coefficients has been adopted as a discrete mother wavelet for its revealed suitability to power quality analysis [ADDT98].

Figure 7 shows the DTWT analysis results obtained by decomposing the original signal, a pure sine wave with an overlapped oscillatory transient with a 600 Hz frequency, into 7 frequency subbands, with $f_s = 12.8$ kHz and $f_f = 50$ Hz. The extraction of the disturbance (3) from the fundamental is obtained by reconstructing (DTWT synthesis stage) only the synthesis results coming from the subband including the reciprocal of the transient duration Figure 7 and the nearest two. The reconstructed disturbance gives the possibility of evaluating some relevant parameters (amplitude, rise time, and so on) and, therefore, to classify it.

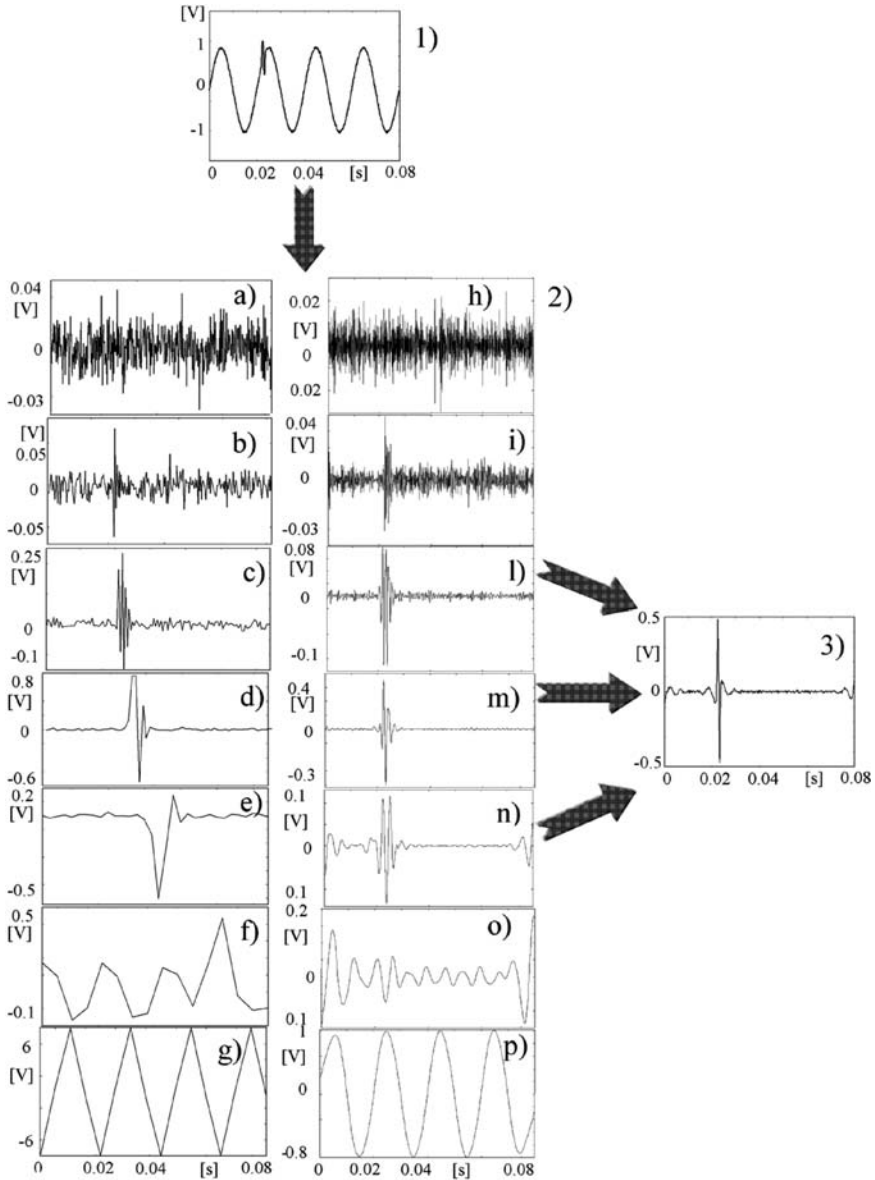


Fig. 7. (1) Powerline signal with a transient. (2) Results of the DTWT analysis and synthesis in the following frequency subbands: (a), (h) 6400–3200 Hz; (b), (i) 3200–1600 Hz; (c), (l) 1600–800 Hz; (d), (m) 800–400 Hz; (e), (n) 400–200 Hz; (f), (o) 200–100 Hz; (g), (p) 100–0 Hz. (3) Extracted transient.

7 Chirplet transform

Like the other *Time–Frequency Representations* (TFRs), the *Chirplet Transform* (CT) projects the input signal onto a set of functions that are all obtained by modifying an original window function $g(t)$, also called the *mother chirplet* (see Figure 8a.) Besides the well-known time and frequency shifting peculiar to the *Short Time Fourier transform* (STFT) – shown in Figure 8b, c – and scaling adopted by the WT –, see Figure 8d – the CT performs other modifications such as chirping both in time and frequency, Figure 8e, f. Thanks to chirping, the CT is capable of rotating each cell of the time–frequency plane as well as shearing it along the time and frequency axes, (Figure 8g) [AD02].

New degrees of freedom in shaping the cells are thus available to the user with respect to the other TFRs, which, ultimately, means new opportunities of optimizing the time–frequency resolution ($\Delta t \times \Delta f$) according to the analyzed signal. As an example, with reference to instantaneous frequency estimation, it is possible to rotate and shear each cell according to the local

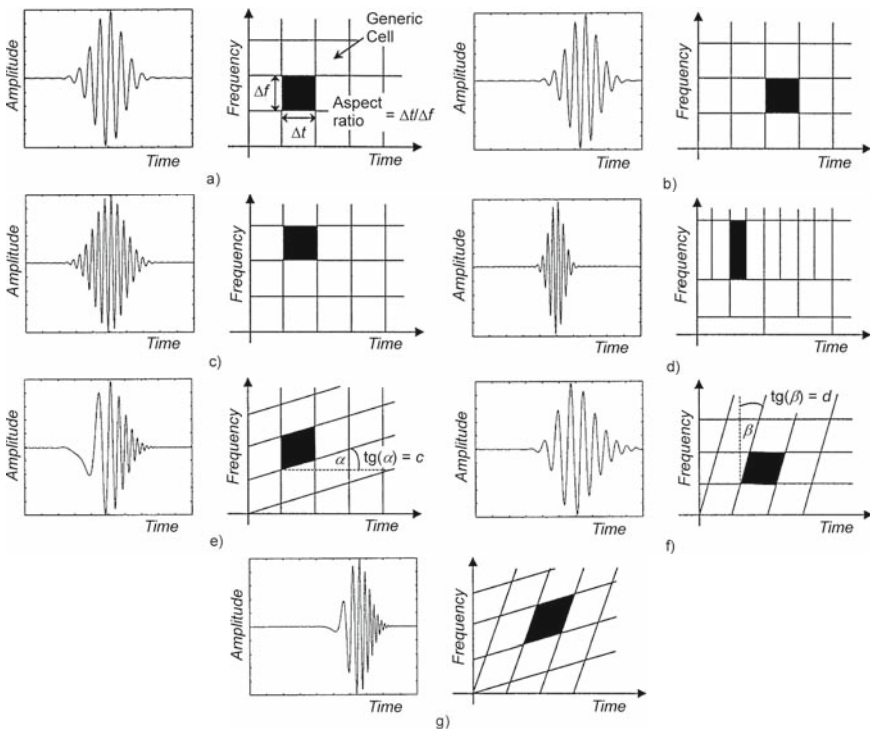


Fig. 8. Modifications of the mother chirplet peculiar to the traditional CT and related effects on the time–frequency plane: (a) original window function, (b) time shifting, (c) frequency shifting, (d) scaling, (e) chirping in time, (f) chirping in frequency, and (g) all the modifications applied [AD02].

slope of the trajectory of the analyzed instantaneous frequency, thus giving the opportunity of best tracking its evolution versus time.

Chirping in time is obtained by multiplying the mother chirplet, $g(t)$, with a linear FM signal, also known as chirp, given by

$$e^{j2\pi(c/2)t^2} \quad (38)$$

where the parameter c is the so-called chirp rate. It causes a rotation of all cells on the time–frequency plane as well as their shear along the frequency axis. Specifically, the slope of the cell is determined by the value of the aforementioned parameter c (see Figure 8e).

Chirping in frequency, on the other hand, is given by the convolution in the time domain between the mother chirplet and another FM linear chirp

$$(-jd)^{-(1/2)} e^{j2\pi(1/2d)t^2}, \quad (39)$$

where the parameter d accounts for the shear amount along the time axis imposed on the cell (Figure 8f). It is worth highlighting that the described operation turns in a product in the frequency domain between the Fourier transform of the mother chirplet $G(f)$ and

$$e^{-j2\pi(d/2)f^2} \quad (40)$$

which can be considered a chirp in the frequency domain, the chirp rate of which is equal to d .

The combined effect of all the aforementioned modifications of the mother chirplet gives rise to the following complete analytic expression of the CT of a signal $s(t)$,

$$CT_s(t, f, a, c, d) = \int s(\tau) h^*(\tau - t, f, a, c, d) d\tau. \quad (41)$$

The kernel $h(\tau - t, f, a, c, d)$ is given by

$$h(\tau - t, f, a, c, d) = \frac{1}{\sqrt{-jda}} \int g \left[\frac{(\tau - t)}{a} - v \right] e^{j\pi \{ 2f\tau + c [((\tau - t)/a) - v]^2 + (v^2/d) \}} dv \quad (42)$$

where t and f account, respectively, for time and frequency shift, and a is the scale parameter.

To improve the efficacy of the traditional CT in instantaneous frequency estimation, some changes in the analytic definition of this transform as well as a proper choice and use of its parameters are proposed in [AD02]. The main advantage of the modified version over the traditional CT is the better aptitude of adapting its resolution, in terms of shape and aspect ratio (see Figure 8a) of the cell, to the local features of the analyzed signal, in any point of the time–frequency plane (Figure 9).




Fig. 9. Proposed modifications of the mother chirplet and related effects on the time–frequency plane (cfr. [AD02]).

8 Wavelet networks

The idea of using Artificial Neural Networks (ANNs) for classification purposes has proved useful for a long time. To achieve best classification performance, a preprocessing of the input samples is advisable. In particular, for nonstationary signals, time–frequency transforms are desirable. Such transforms yield simultaneously two different aspects of the signal: the frequency characteristics as well as the temporal behaviour. From the variety of solutions, the WT shows itself the most suitable for the analysis of transient signals in several application fields.

The *wavelet network* (WN) combines the properties of the WT with the advantages of ANNs. Wavelet networks can be considered as an extended perceptron consisting of two parts (Figure 10). The first part contains the so-called wavelet nodes, in whose the classical sigmoidal activation functions are substituted by mother wavelet functions $h(a, b)$ as in [DMR01], but without the energy normalization factor $a^{-1/2}$. They act as preprocessing units for transient detection and feature extraction. The classification is performed by the second part, a traditional single-layer or multilayer perceptron [DMR01]. The classic backpropagation method is applied to train the network by tuning the weighting coefficients, the sigmoid steepness, and the wavelet node scale and time parameters (a and b) in order to minimize the difference, in the mean squared error sense, between the current and the target output arrays for a given training signal set. As a consequence, during the training stage, the WN is able not only to learn arbitrarily complex decision regions defined by the weight coefficients, but also to look for those parts of the time–frequency plane that are suited for a more reliable classification of the input signals. In [DMR01] an example of using WNs to classify the disturbances on quadrature amplitude modulated signals is described. In the paper some considerations about the correct procedure to choose the mother wavelet, the WN modularity, and the selection of training signals can be found may be found.




Fig. 10. Wavelet Network scheme.

References

- [AD02] L. Angrisani, and M. D’Arco. A measurement method based on a modified version of the chirplet transform for instantaneous frequency estimation. *IEEE Transactions on Instrumentation and Measurement*, 51(4):704–711, 2002.
- [ADD99] L. Angrisani, P. Daponte, and M. D’Apuzzo. A method for the automatic detection and measurement of transients. Part i: The measurement method. *Measurement*, 25(1):19–30, 1999.
- [ADDT98] L. Angrisani, P. Daponte, M. D’Apuzzo, and A. Testa. A measurement method based on the wavelet transform for power quality analysis. *IEEE Transactions on Power Delivery*, 13(4):990–998, 1998.
- [AH92] A. N. Akansu, and R. A. Haddad. *Multiresolution Signal Decomposition*. Academic Press, San Diego, CA, 1992.
- [BM94] P.M. Bentley, and J.T.E. McDonnell. Wavelet transforms: An introduction. *Electronics & Communication Engineering Journal*, 6(4):175–186, 1994.

- [CT65] J. W. Cooley, and J. W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19:297–301, 1965.
- [Dau92] I. Daubechies. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, 1992.
- [Dir29] J. P. G. Lejeune Dirichlet. Sur la convergence des séries trigonométriques que servent à représenter une fonction arbitraire entre des limites données. *Journal de Crelle*, 4:157–169, 1829.
- [DMR01] P. Daponte, G. Mercurio, and S. Rapuano. A wavelet networks-based method for the digital telecommunication system monitoring. *IEEE Transactions on Instrumentation and Measurement*, 50(6):1773–1780, 2001.
- [Fou78] J.-B. J. Fourier. *The Analytical Theory of Heat*. Cambridge Press, London, 1878.
- [IEE01] IEEE. Standard for terminology and test methods for analog-to-digital converters. Standard 1241, IEEE, 2001.
- [JF07] S. G. Johnson, and M. Frigo. A modified split-radix fft with fewer arithmetic operations. *IEEE Transactions on Signal Processing*, 55(1):111–119, 2007.
- [LAD99] L. Angrisani, P. Daponte, and M. D’Apuzzo. A method for the automatic detection and measurement of transients. Part ii: Applications. *Measurement*, 25(1):31–40, 1999.
- [Nyq28] H. Nyquist. Certain topics in telegraph transmission theory. *Transactions AIEE*, 47:617–644, 1928.
- [OSB99] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck. *Discrete-Time Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 1999.
- [PL06] C. Petrarca, and G. Lupo. Wavelet packet denoising of partial discharge data. *Proceedings of 2006 IEEE Conference on Electrical Insulation and Dielectric Phenomena*, pages 644–647, 2006.
- [Rio93] O. Rioul. A discrete-time multiresolution theory. *IEEE Transactions on Signal Processing*, 41(8):2591–2606, 1993.
- [RV91] O. Rioul, and M. Vetterli. Wavelets and signal processing. *IEEE Signal Processing Magazine*, pages 14–38, Oct 1991.
- [Sha49] C. E. Shannon. Communication in the presence of noise. *Proceedings Institute of Radio Engineers*, 37(1):10–21, 1949.
- [SML96] R. G. Stockwell, L. Mansinha, and R. P. Lowe. Localization of the complex spectrum: The s transform. *IEEE Transactions on Signal Processing*, 44(4):998–1001, 1996.
- [VH92] M. Vetterli, and C. Herley. Wavelets and filter banks: Theory and design. *IEEE Transactions on Signal Processing*, 40(9):2207–2232, September 1992.

Data Fusion, Decision-Making, and Risk Analysis: Mathematical Tools and Techniques

Pedro S. Girão¹, Octavian Postolache^{1,2}, José M. D. Pereira^{1,2}

¹ Instituto de Telecomunicações, Instituto Superior Técnico, DEEC, UTL, Av
Rovisco Pais 1, 1049-001 Lisbon, Portugal, psgirao@ist.utl.pt,
poctav@ist.utl.pt

² Escola Superior de Tecnologia, IPS, Setúbal, joseper@est.ips.pt

Summary. Human activity involves sequential decision-making. Activities with alternatives require deciding for one of the alternatives. A rational decision is one that weighs each alternative pros, cons, and risks. The support for decision-making is data that come basically from experience, either previously acquired or gathered for the specific decision-making. The data usually come from different sources and thus have to be fused for a single decision. The core of this chapter is precisely about data fusion. In its subsections, we look namely at some procedures and techniques commonly used in data fusion. Decision-making and risk analysis are briefly discussed

Key words: Sensor data fusion, data fusion, data fusion tools, decision-making, risk analysis

1 Data fusion

In this section we address issues related to *data fusion*. Our concern is applications that can be framed in the engineering domain where experimental data play a decisive role. Procedures and techniques considered are pertinent for data fusion obtained with different sensing systems, independent of size and organization.

It is impossible to present in detail all the techniques and algorithms that we elected to include. We refer the reader to [1] to [4] for details of fusion techniques, and for the mathematically inclined, [5]. In a limited number of pages it is only possible to introduce them, some in more detail than others. The exposition of each subject is complemented with references that in general deal not with the state-of-the-art but rather with the fundamentals and foundations of that subject.

1.1 Definitions, concepts, and terms of reference (terminology). Processing and topological issues

Data or information fusion is currently defined in several ways; some focused on supporting a functional model and architecture [1], others centered in the framework and fundamentals [6]. Nevertheless, we think there is general agreement that the ultimate idea underlying data fusion is to obtain greater quality information (better suited for a specific purpose) by exploiting the synergy of data gathered from different sources.

According to [1], *data fusion is the process of combining data or information to estimate or predict entity states*. Data fusion deals with everything that has to do with this objective and because problems in many domains of application require some decision-making, data fusion also encompasses aspects such as *classification and pattern recognition* used to support decisions.

As an organized research area, data fusion owes a lot to people working in the framework of military applications, target tracking and identification in particular, in which it is crucial not only to fuse data obtained from multiple sensors but also to assess threats and risk. It is thus natural that these issues are included in the data fusion domain by those defending the definition, model, and architecture proposed for instance in [1].

As mentioned in [1], data fusion has or may have many advantages: (a) it increases robustness and reliability and reduces the vulnerability of the system supporting the decision, because it allows decision-making even in the absence of malfunction of some sources of information; (b) it may provide a better and larger coverage of space and time; (c) it reduces ambiguity, because better information provides better discrimination between available hypotheses; and (d) it provides a solution to deal with the large amount of data that may be available.

It is beyond the scope of this chapter to provide an exhaustive and detailed treatment of data fusion. Several texts on the subject are available and we particularly recommend [1] to [7] for those seeking insight of the data/information fusion framework.

Information fusion is based on experimental data output by sensing devices and eventually on information obtained by other means (e.g., the user as a data source for a priori knowledge, experience, model application). Fusion requires or advises all data to have the same representation (e.g., numeric values in the same units, relative values), the realization of which is often the first step of the data fusion process. This process, particularly difficult if data are heterogeneous (noncommensurate), is often referred to as *data alignment* or *data registration* [6].

Measurements, understood as the output of a sensor, form a signal [6] more or less affected by noise whose reliability has to be verified (e.g., malfunction of the sensor, express corruption of sensors' measured quantity, for instance, jamming). *Data filtering* and *data validation* are two common and important tasks in a data fusion. Other basic aspects involved in data fusion involve

addressing: (a) data coming from sources with different levels of quality (e.g., different accuracy); (b) nonindependent data; (c) too much information, which may lead to computational problems; and (d) need to change the context of the observation (e.g., from time to frequency domain) or to extract features or attributes [6].

The above-mentioned issues are of the data-processing type. They address questions such as how to select the proper measurements, determine the relevance of the data, and ultimately, how to select the fusion methods and architectures, once the data are available. Other types of issues that we do not address here but that are important for correct decision-making are of a topological nature: sensors and their spatial distribution, the communication network between sensors and places of processing and decision-making, information exchange, availability and reliability of information at the time of the fusion, and cost of acquiring the information. These topological issues are particular pertinent with the increasing number of applications that use sensor networks, namely of the wireless type. In such applications it is not trivial but extremely important to organize and distribute data fusion tasks in the network. It is a subject also beyond the scope of this chapter and has been researched for a few years now (e.g., [8]).

The data/information fusion domain uses dedicated terminology and terms of reference, some of them with different meanings according to the authors or to the specific domain of application (e.g., association: concatenation of data, or correlation between a measurement and the actual value of interest). Some of those terms were already introduced and some others are used in the paragraphs that follow. Most of them are presented and discussed in detail in [6].

1.2 Data fusion models and architectures

The first data fusion model was developed in 1985 by the U.S. Joint Directors of Laboratories (JDL) Data Fusion Group and is known as the JDL data fusion model. Widely accepted as the base model, it was revised more than once, the later revision dating from 2004 [9]. Also in 2004, the Data Fusion Information Group, which integrates most of the people working in data fusion modeling, proposed an upgrade to the JDL model, the DFIG model [10]. Dasarathy [11], Bedworth and O'Brien [12], and Salerno [13], among others, also proposed models and architectures for data fusion with merits.

In different ways, all data fusion models, either of the functional type, such as the JDL and Dasarathy models or of the process type, such as the Omnibus model of Bedworth and O'Brien, differentiate types of data fusion functions using fusion levels. Those levels are more or less connected to the stages at which the fusion occurs. There are three fundamental stages to consider (Figure 1):

- *Direct fusion of sensors' data or raw data fusion.* If sensors measure the same quantity and the data are registered in time and space, their raw

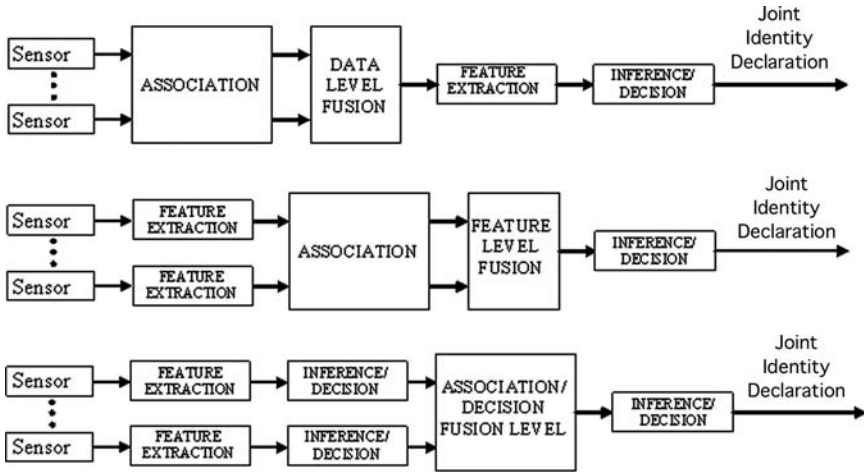


Fig. 1. Sensor data fusion processing architectures.

data can be directly fused; otherwise, either the output of each sensor is converted into relative units (normalized), and special algorithms (e.g., neural networks) used, or data must be fused at the feature/state vector stage or decision stage.

- *Feature or state vector fusion.* A characteristic or a feature vector representing data is extracted before fusion is implemented. This often involves a data-mining problem, that is, the extracting of information from large datasets or databases.
- *Decision fusion.* Data are combined with other data or with a priori knowledge or the data are processed to yield inferences or decisions that are then combined.

1.3 Data fusion techniques and algorithms

Measuring transducer inverse modeling and characteristic approximation. Unit adjustment. Coordinate transformation. Normalization.

Measuring transducers, often simply called sensors, and measuring devices are the main sources of experimental data. Electrical measuring transducers, which are the most used, output an electric quantity (e.g., voltage, current) that informs on the quantity in the input of the transducer sensing device (measurand). Examples include radar, sonar, and cameras. The relation between the transducer’s output and the measurand must be known to calculate the measurand’s value from the transducer’s output value. That relation is called the transducer characteristic and the operation now mentioned is often

called *transducer inverse modeling*. When the transducer characteristic is not linear, it is necessary to approximate it to simplify inverse modeling.

A data fusion algorithm does not usually operate on raw data, but rather on some types of data in some representations. This implies that the first steps in a data fusion process may need what we would call *data formatting*. This may involve, for instance, unit conversion, coordinate transformation, data registration, and data normalization.

Association metrics and similarity functions

A *metric*, d , is a function that satisfies the following properties:

1. Nonnegativity: $d(x, y) \geq 0$.
2. Symmetry: $d(x, y) = d(y, x)$.
3. Identity: $d(x, x) = 0$.
4. Definiteness: $d(x, y) = 0$ if and only if $x = y$.
5. Triangle inequality: $d(x, y) + d(y, z) \geq d(x, z)$.

Association metrics is a general term for assigning a number representing a degree of likeness to data [3, 14]. Particular types of association metrics have special designations (e.g., association measures, matching coefficients, distance measures). The association metrics measures usually the similarity of two elements of a dataset and is of paramount importance in digital-processing techniques for classification and pattern recognition purposes (e.g., clustering).

The selection or definition of an association metrics depends essentially on the structure of the data. Some examples follow.

1. Comparison between two elements $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$ of two numerical datasets: *sample product moment correlation*,

$$C(X, Y) = \frac{\text{cov}(X, Y)}{[\text{var}(X)\text{var}(Y)]^{1/2}} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\left[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 (y_i - \bar{y})^2\right]^{1/2}}, \quad (1)$$

where *var* and *cov* stand for variance and covariance and the $\bar{}$ for mean value.

2. Comparison between elements X and Y of datasets where one or both are not numerical: the range of each variable is divided into classes (intervals),

$$X = \bigcup_{j=1}^p X_j \quad Y = \bigcup_{k=1}^q X_k \quad (2)$$

and correlation coefficients $A_{j,k}/n$ are defined. $A_{j,k}$ is the number of data values for which $x_i \in X_j$ and $y_i \in Y_k$.

3. Comparison between elements X_j and X_k of a numerical dataset: *Minkowski distance function*,

$$D_p(X_j, X_k) = \left(\sum_{i=1}^n |X_{ij} - X_{ik}|^p \right)^{1/p}. \quad (3)$$

For $p = 2$, Equation (3) is the Euclidean metric or Euclidean norm; for $p = 1$ the metric is called the Manhattan or *taxi cab*. If the elements are different quantities or are in different units, they should be normalized using, for instance:

$$\hat{X}_j = \left(\frac{x_{ji} - \bar{x}}{\sigma_j} \right)_{i=1}^n, \quad (4)$$

where σ is the standard deviation.

When p tends to infinity, the *minimax* or *Chebyshev distance* results:

$$D_\infty(X_j, X_k) = \max_{i=1}^n |x_{ij} - x_{ik}|. \quad (5)$$

4. Comparison of datapoints whose components are some or all not numerical: the Minkowski distance function with

$$|x_{ij} - x_{ik}| = \begin{cases} 0 & \text{if } x_{ij} = x_{ik} \\ 1 & \text{otherwise} \end{cases}. \quad (6)$$

A similarity function S is more loosely defined than a metric and satisfies the three following properties.

1. Nonnegativity: $S(x, y) \geq 0$,
2. Symmetry: $S(x, y) = S(y, x)$,
3. The more similar the objects a and b , the greater is $S(x, y)$.

General metrics is a subject covered by the mathematical literature in measure theory. Nonetheless, practical applications in domains such as physics, biology, or object and target tracking and identification do require special metrics and similarity functions, some of which are discussed in [15].

Figures of merit

Generally speaking, data fusion encompasses all activities that ultimately support a decision for a user. The performance not only of the processing involved but also of the final result must have a way to be evaluated. Figures of merit are indicators that allow or simplify that evaluation. Blasch *et al.* [16] propose a minimum set of metric figures that include: accuracy, confidence, throughput, timeliness, and cost. One simple example of two possible figures of merit when evaluating the performance, for instance of a program that implements a routine for pattern recognition: number of floating-point operations involved (execution time) and memory requirements.

Gating

In the context of data fusion, *gating* are correlation-type algorithms used to validate data. Particularly used in tracking systems (a track is a state trajectory estimated from the observations that have been associated with the same object), a *gate* defines a neighborhood of an observation or predicted observation (data element) inside which a new observation should be. Several gate types can be used according to the problem at hand [17].

Rectangular gate

If \hat{O}_i is a predicted observation, an observation O_i is valid if the residual $|O_i - \hat{O}_i|$ satisfies the condition:

$$|O_i - \hat{O}_i| \leq K_{RG}\sigma_r, \quad (7)$$

where

$$\sigma_r = (\sigma_O^2 + \sigma_{\hat{O}}^2)^{1/2} \quad (8)$$

is the standard deviation of the residual.

Elliptical gate

In the case of an elliptical gate, the relation that the residual must satisfy is:

$$[O - \hat{O}]^T \text{cov}(O, \hat{O})^{-1} [O - \hat{O}] \leq K_{EG}. \quad (9)$$

Appropriate values for K_{RG} are in the interval 2.81–3.09 and for K_{EG} in the 9.21–15.09 interval [18].

Dynamic systems often require more complicated gates [4].

FFT and cepstrum

Fourier transform

It is well known that time and frequency are reciprocal quantities and that it is straightforward to transform a signal defined in the time domain into one defined in the frequency domain using the *Fourier transform*.

The Fourier transform $X(f)$ of a continuous time function $x(t)$, which changes the representation from the time to the frequency domain, can be expressed as

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft} dt. \quad (10)$$

To obtain the time representation $x(t)$ of a signal whose frequency representation is $X(f)$ the *inverse transform* formula must be used,

$$x(t) = \int_{-\infty}^{\infty} X(f)e^{j2\pi ft}df. \quad (11)$$

Expressions (10) and (11) reveal that the Fourier transformation is supported on sine and cosine functions whose arguments are integer multiples of a base value. Such functions are orthogonal and thus form a base in which any signal can be represented.

Some authors prefer to write the transform in terms of angular frequency, $w = 2\pi f$, which leads to the introduction of either a factor $1/2\pi$ in either Equations (10) or (11) or $(1/2\pi)^{1/2}$ in both Equations (10) and (11).

In most practical situations, though, the signal is only known at some finite number of time instants t_i (sampled systems). Then, if we consider a complex series $x(n)$ with N samples of the form $x_0, x_1, x_2, x_3, \dots, x_{N-1}$ where x is a complex number, $x_i = x_{\text{real}} + jX_{\text{imag}}$ and that the series outside the range $0, N - 1$ is extended N -periodic, that is, $x_n = x_{n+N}$ for all n , the Fourier transform of this series, the *Discrete Fourier Transform* (DFT) denoted by $X(k)$, will also have N samples and is defined by:

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2kn\pi/N} \quad n = 0, \dots, N - 1. \quad (12)$$

The inverse transform will then be given by:

$$X(n) = \frac{1}{N} \sum_{k=0}^{N-1} x(k)e^{j2kn\pi/N} \quad n = 0, \dots, N - 1 \quad (13)$$

The *Fast Fourier Transform* (FFT) is nothing else but an efficient algorithm to compute the discrete Fourier transform and its inverse. In fact, the evaluation of the DFT requires about N^2 arithmetical operations, $(N - 1)^2$ complex multiplications, and $N(N - 1)$ complex additions. The FFT is an algorithm to compute the same result in only $N \log N$ operations. Because the inverse DFT is the same as the DFT, but with the opposite sign in the exponent and a $1/N$ factor, any FFT algorithm can easily be adapted for it as well.

Many FFT algorithms only depend on the fact that $e^{-j(2\pi/N)}$ is a primitive root of unity. By far the most common FFT algorithm is the Cooley–Tukey algorithm [19]. This is a divide-and-conquer algorithm that recursively breaks down a DFT of any composite size $N = N_1 N_2$ into many smaller DFTs of sizes N_1 and N_2 , along with multiplications by complex roots of unity. In the most well-known use of the Cooley–Tukey algorithm the transform is divided into two pieces of size $N/2$ at each step, and is therefore limited to power-of-two sizes, but any factorization can be used in general. Although the basic idea is recursive, most traditional implementations rearrange the algorithm to avoid explicit recursion.

The FFT is the object of a particular chapter of this book, *Frequency and Time-Frequency Domain Analysis Tools in Measurement*, so we invite the reader to consult it or the abundant bibliography on the subject for detailed information. Nevertheless, and because the DFT is used in the analysis of continuous time signals and systems, the approximations involved when using the DFT and the problems that may arise in the process that may lead to erroneous results merit a few comments here.

As it is well known, *undersampling*, that is, sampling at a frequency lower than twice the maximum frequency of the sampled signal (Nyquist rate), causes frequency components that are higher than half of the sampling frequency to overlap with the lower-frequency components. As a result, the higher-frequency components roll into the reconstructed signal and cause distortion of the signal called *aliasing*. The only solution to the aliasing problem is to ensure that the sampling rate is high enough to avoid any spectral overlap, or to use an anti-aliasing filter.

Signals are observed in a finite time interval. Terminating a signal after a finite number of terms is equivalent to multiplying the signal by a window function. The result is a distortion of the spectrum. There is a spreading or leakage of the spectral components away from the correct frequency, which leads to an undesirable change in the total spectrum. Because leakage results in a spreading of the spectrum, the upper frequency may move beyond the Nyquist frequency, and aliasing may also then result. The best approach for alleviating the leakage effect is to choose a suitable window function that minimizes the spreading.

DFT looks at the spectrum not as a continuous function but rather as through a “picket-fence” because the observation is only at discrete points. If the peak of a particular component lies between two of the discrete transform lines it is not detected without some additional processing. It is convenient to remember here that DFT points are separated in frequency (resolution of the DFT) by the reciprocal of the total sampling time and that thus each line of the DFT corresponds to a frequency whose value is obtained by multiplying the line order by the reciprocal of the total sampling time. One procedure for reducing this picket-fence effect is to vary the number of points in a time period by adding zeros at the end of the original record, while maintaining the original record intact. This process artificially changes the period, which in turn changes the locations of the spectral lines without altering the continuous form of the original spectrum. In this manner, spectral components originally hidden from view can be shifted to points where they can be observed.

In the context of data fusion, the Fourier transform, and the FFT in particular, can be used for many purposes and situations, namely: (a) to convert data from the time to the frequency domain for, for instance, alignment purposes; (b) to digitally (numerically) filter a signal (FFT of the signal, elimination of the undesired components, inverse FFT to obtain the filtered signal); and (c) to extract features of a signal for posterior data fusion.

The Fourier transform provides information of high resolution in what concerns the frequency contents of a signal and is thus very useful when the signal is stationary; that, is its frequency content is time constant. For nonstationary signals, it is often important to know not only the frequency contents but also the time those frequencies occur. Because the Fourier transform involves an infinite integration time, frequency resolution is extremely high but, on the other hand, time resolution is zero. This is a manifestation of the Heisenberg uncertainty principle applied to time–frequency: one cannot know the exact time–frequency representation of a signal; that is, one cannot know what spectral components exist at what instances of times. It is, however, possible to know both representations with a lower bounded uncertainty. To do this, the Fourier transform must be changed so that the integration interval is finite. The solution used in the *Short Time Fourier transform* (STFT) is to divide the signal into small enough segments, where these segments (portions) of the signal can be assumed to be stationary using a window function, and evaluate the Fourier transform of each segment. The width of the window must be equal to the segment of the signal where its stationarity is valid. The result of the STFT is a time–frequency function. As we show, the wavelet transform also provides simultaneous time–frequency information of a signal and because it outperforms the STFT we do not elaborate more about the STFT.

One final but important comment: expressions (10)–(13) consider a time–changing signal (one-dimensional Fourier transform). However, some applications, such as those involving images, require processing multidimensional signals. Expressions (10)–(13) can be extended to higher dimensions ([20], [21]).

Cepstrum

Another mathematical tool of interest for signal analysis and feature extraction is the so-called *cepstrum* (spectrum with letters of the first syllable inverted). The cepstrum of a signal was introduced in [22] and was defined as the Fourier transform of the logarithm, with unwrapped phase (phase between 0 and 2π or between $-\pi$ and π), of the Fourier transform of the signal; that is:

$$\text{cepstrum of signal} = \text{FT}(\log(\text{FT}(\text{signal}))) + j2\pi m, \quad (14)$$

where FT stands for the Fourier transform given by Equations (10) and (12), log for the natural logarithm, and m is the integer required to properly unwrap the angle or imaginary part of the complex natural logarithm function. The independent variable of the cepstrum thus defined was labeled *quefrequency* and operations on cepstra labeled *quefrequency analysis* or *cepstral analysis*.

Nowadays, however, it is common to relate the cepstrum of a signal not with the Fourier transform but with the inverse Fourier transform of the $\log(\text{FT}(\text{signal}))$. In this case, the independent variable of the cepstrum has time units. Software packages, such as MATLAB, have functions that evaluate the cepstrum this way (e.g., MATLAB Digital Signal Toolbox function *rceps*, if the logarithm is calculated using the magnitude of the signal's Fourier

transform (real cepstrum) or *cceps*, if the complex logarithm of the spectrum is used (complex cepstrum)).

By analogy with a filter that is a selective device in the frequency domain, a *lifter* is a filter that operates in the cepstrum domain. Filtering of a signal can be implemented by multiplying the cepstrum by a window in the cepstral domain and then converting back to the time domain.

The cepstrum basically provides information about the rate of change in a signal's spectrum and was originally proposed and used for characterizing and identification of echoes (e.g., radar) [23]. In terms of signal-processing methods taxonomy, the cepstrum belongs to the class of homomorphic deconvolution methods [24].

Nowadays, the domains where cepstral analysis is more used are sound and voice recognition and speech and music classification [25].

Wavelets

The Fourier transform is a powerful tool for signal processing that allows the representation of a signal as the sum of a possibly infinite number of sines and cosines. This Fourier expansion means that we are able to determine all the frequencies present in a signal, but has the drawback that we do not know when they are present in time. The Fourier expansion provides frequency resolution but no time resolution. Because both pieces of information are important in many applications, namely those dealing with nonstationary signals, several solutions have been developed to overcome the Fourier transform limitation and thus to represent a signal in the time and frequency domain at the same time. The *wavelet transform* or *wavelet analysis* is probably the most interesting solution to overcome the shortcomings of the Fourier transform. A large number of more [26, 27] or less [28], [29, 30] mathematical oriented publications exist on the wavelet transform and thus we mention here only its basics.

Continuous wavelet transform

The wavelet transform is based on the representation of a function using base functions generated from a single basic wavelet $\Psi(t)$, called the *mother wavelet*, by *scaling* and *translation*,

$$\Psi_{S,T}(t) = \frac{1}{\sqrt{s}} \Psi\left(\frac{t-\tau}{s}\right). \quad (15)$$

In Equation (15), s is the scale factor, t is the translation factor, and the factor $s^{-1/2}$ is for energy normalization across the different scales. The *Continuous Wavelet Transform* (CWT) is then given by:

$$Y(s, \tau) = \int_{-\infty}^{\infty} f(t) \Psi_{s,\tau}^* dt, \quad (16)$$

where $*$ denotes complex conjugation and $f(t)$ is the signal to be transformed. The inverse wavelet transform is given by

$$f(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} Y(s, \tau) \Psi_{s, \tau} d\tau ds. \tag{17}$$

Contrary to the Fourier transform, the wavelet transform is not anchored in a specific type of base function. The mother wavelet can be any function provided that it satisfies some conditions: admissibility, regularity, and vanishing moments. From Equation (16) it is clear that the wavelet transform of a one-dimensional function is a two-dimensional function whose representation requires a three-dimensional space, amplitude, s and τ , the second variable informing on the frequency content of the transformed signal and the third variable on time. Because of Equation (15), the higher the s is, the lower is the frequency.

In wavelet analysis, the use of a fully scalable modulated window solves the signal-cutting problem (see short Fourier transform in the FFT and cepstrum paragraphs). The window is shifted along the signal and for every position the spectrum is calculated. Then this process is repeated many times with a slightly shorter (or longer) window for every new cycle. In the end the result will be a collection of time–frequency representations of the signal, all with different resolutions. Because of this collection of representations we can speak of a *multiresolution analysis*. In the case of wavelets, we normally do not speak about time–frequency representations but about *timescale representations*, scale being in a way the opposite of frequency, because the term frequency is reserved for the Fourier transform.

Discrete wavelet transform

The continuous wavelet transform defined by Equation (16) has some limitations: (a) the scaled functions used do not form an orthogonal base and thus the information obtained is redundant. Note, however, that the redundancy of the CWT is sometimes useful, for instance, to reduce sensitivity to noise; (b) the number of functions is often unmanageable; (c) for most signals, the wavelet transform has no analytical solution and can be calculated only numerically.

Problem (a) is overcome by using not a continuous set of wavelet functions, but wavelets whose scale values assume a discrete number of values (discrete wavelets). This is achieved by:

1. Modifying the wavelet representation:

$$\Psi_{j,k}(t) = \frac{1}{\sqrt{s_0^j}} \Psi \left(\frac{t - k\tau_0 s_0^j}{s_0^j} \right), \tag{18}$$

where j and k are integers and $s_0 > 1$ is a fixed dilation step, so as to be discrete Discrete Wavelet Transform (DWT). The translation factor τ_0




Fig. 2. Localization of the discrete wavelets in the timescale space on a dyadic grid.

depends on the dilation step. The effect of discretizing the wavelet is that the timescale space is now sampled at discrete intervals. Usually, $s_0 = 2$ so that the sampling of the frequency axis corresponds to dyadic sampling. Where the translation factor is concerned, $\tau_0 = 1$ also means a dyadic sampling of the time axis. Figure 2 shows the wavelet localization in the timescale space in this case.

- Appropriate selection of the mother wavelet. A function is eligible for a mother wavelet if it is oscillatory and its amplitude tends quickly to zero. If a set of discrete wavelets satisfies

$$\int_{-\infty}^{\infty} \Psi_{j,k}(t) \Psi_{m,n}^*(t) dt = \begin{cases} 1 & \text{if } j = m \text{ and } k = n \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

then they form an orthogonal base.

It is interesting to note that: (a) the use of discrete wavelets leads to a series of wavelet coefficients, $\gamma(j, k)$; (b) the coefficients evaluated using an orthogonal base allow the reconstruction of the signal in a way similar to the Fourier series for periodic signals (inverse wavelet transform),

$$f(t) = \sum_{j,k=-\infty}^{\infty} \gamma(j, k) \Psi_{j,k}(t). \quad (20)$$

Even with discrete wavelets, an infinite number of functions are required to calculate the wavelet transform and thus something else must be done to overcome problem (b). Time translation is naturally limited because the signal to transform is also time limited. Where scaling is concerned, it is possible to reduce the number of wavelets and still have a fair result using an additional function called the *scaling function* or *father wavelet*. Figure 3 shows a particular pair, scaling function–mother wavelet.

We invite the reader to find a detailed explanation of how a finite number of wavelets plus a scaling function do provide an efficient solution for the wavelet transform evaluation in the literature (e.g., in [30]). If all the functions are orthogonal, one is in the *multiresolution domain* (MRA). In a very simple




Fig. 3. Daubechies 20 scaling function (left) and Daubechies 20 mother wavelet (right).

way but containing the essence of the explanation, a wavelet has, and must have, a bandpasslike spectrum and the scaling function has a lowpasslike spectrum and it is fairly easy to understand that it is possible to analyze all the signal's spectrum using a lowpass filter and bandpass filters as long as the bands of adjacent filters overlap. The series of dilated wavelets together with a scaling function act as a *filter bank*. This idea is also instrumental for the implementation of a practical algorithm to calculate the wavelets and scaling function coefficients – *wavelet analysis* – that takes care of problem (c).

A detailed procedure of the algorithm due to Mallat that is commonly used and called *Fast Wavelet Transform* (FWT) is presented in the chapter *Frequency and Time-Frequency Domain Analysis Tools in Measurement* in this book. The algorithm uses not bandpass filters but highpass filters for practical reasons. Because the signal to transform (analyze) is by nature bandlimited and because each iteration reduces to half the bandwidth to be analyzed, the two solutions are equivalent. It is an iterative algorithm whose number of iterations would be infinite for a continuous signal. However, signals must be sampled and thus the number of iterations depends ultimately on the number of samples. Because on each iteration the bandwidth of the signal analyzed is halved (the number of samples used is also halved), the number of signal samples should be a multiple of a power of 2, the power of two corresponding to the number of iterations (*decomposition level*) and the value of the multiple corresponding to the number of samples used in the last iteration. If the total number of samples is $2n$, then the maximum theoretical number of iterations is n and the number of filter outputs is $n + 1$ (n details and 1 approximation). Figure 4 shows the process in the frequency domain. Note that with this algorithm, based on the decomposition of the signal spectrum in frequency subbands, the high-frequency components of the signal (detail 1) are obtained first (first iteration, first level of decomposition), being the DC component (approximation) of the last value output.

An alternative to the decomposition process just mentioned is the wavelet packet decomposition algorithm. Whereas in the wavelet decomposition only




Fig. 4. Frequency domain representation of the DWT: d1, d2, d3, details level 1, 2, and 3, respectively; a3, approximation level 3.

the approximation (output of the lowpass filter with the number of samples divided by two) is further decomposed, in the wavelet packet decomposition both the approximation and the detail (output of the highpass filter with the number of samples divided by two) are decomposed, which means that for an n -level decomposition, instead of $n + 1$ one gets more than $2^{2^n - 1}$ outputs. The reduction to half of the samples is usually called *downsampling*.

Mother wavelet

The mother wavelet to use depends on both the type of the signal to be transformed and on the objective aimed with the use of the wavelet transform. It is not possible, at least for us, to provide a clue for the selection of the best mother wavelet and scaling function. We have used Daubechies with success, but it depends a lot on the application. There is a lot of research on the subject but naturally application-oriented (e.g., [31]). Contrary to the Fourier analysis, the wavelet analysis is not supported on a single function type and, from that point of view, is an open technique. This means that it is possible to design a mother wavelet better suited for a specific application (e.g., [32–34]).

1.4 Applications

Wavelets were developed independently in fields such as mathematics, quantum physics, electrical engineering, and seismic geology. Nowadays, the number and domains of applications of the wavelet analysis are huge. It is difficult thus to exhaustively mention them all. Because of the number and importance we, however, refer here to some examples: function approximation (e.g., geophysics and mathematics); general data analysis, namely detection of patterns, trends, and structures in data (e.g., geophysics, medicine, biology, engineering, finance); data compression (e.g., engineering, image); de-noising of noisy data (e.g., engineering, image, medicine); rupture and edge detection (e.g., engineering); and short-time phenomena (e.g., engineering, medicine).

Software

There are several programs for wavelet analysis. Perhaps the most widespread and user-friendly is The MathWorks, Inc. MATLAB toolbox.

Mathematica and Mathcad have similar packages: Wolfram Research, Inc. Mathematica Wavelet Explorer, PTC (former Mathsoft) Mathcad Wavelets Extension Pack. A fairly exhaustive reference to wavelet software is at <http://www.wavelet.org>. In the table available at http://www.sm.luth.se/grip/Research/Publications/Wavelets/SoftwareComparison/Gri04S_Wavelet_Software_Comparison_Table.pdf a partial software comparison is presented.

Kalman filtering

Data fusion, as an organized discipline, owes much to people working in defense, in particular in target tracking and identification. The problematic involved is in the context of dynamic systems and, in particular, of dynamic systems where observations are subject to spontaneous or induced noise. One of the tools that performs particularly well under those circumstances is the *Kalman filter* also known as *Linear Quadratic Estimation* (LQE) in control theory.

The Kalman filter [35] is a discrete-time recursive estimator, which means that only the estimated state from the previous time step and the current measurement together with the system process and measurement models are needed to compute the estimate for the current state. Thus, it has the advantage over batch estimation techniques that no history of observations and/or estimates is required. Contrary to most filters that operate in the frequency domain, it is a *time domain filter*.

Kalman filter underlying model

The Kalman filter model assumes the true state at time k depends on the state at time $(k - 1)$ according to:

$$\mathbf{x}(k) = \mathbf{F}(k) \mathbf{x}(k - 1) + \mathbf{G}(k) \mathbf{u}(k) + \mathbf{w}(k), \quad (21)$$

where

- $\mathbf{F}(k)$ is the *state transition model* that is applied to the previous state $\mathbf{x}(k - 1)$.
- $\mathbf{G}(k)$ is the *control-input model* that is applied to the control vector $\mathbf{u}(k)$.
- $\mathbf{w}(k)$ is the *process noise* that is assumed to be Gaussian discrete white noise with covariance $\mathbf{Q}(k)$.

The bold-italic type indicates that the variables are vectors or matrices. Equation (21) has several labels: *state model*, *state equation*, *plant*, *plant model*, *process equation*, and *process model*, among others.

At time k , an *observation* (or measurement) $\mathbf{z}(k)$ of the true state $\mathbf{x}(k)$ is made according to:

$$\mathbf{z}(k) = \mathbf{H}(k) \mathbf{x}(k) + \mathbf{v}(k), \quad (22)$$

where $\mathbf{H}(k)$ is the observation model that maps the true state space into the observed space and $\mathbf{v}(k)$ is the observation noise, which is assumed to be zero mean Gaussian white noise with covariance $\mathbf{R}(k)$. The initial state, and the noise vectors at each time step $\{\mathbf{x}(0), \mathbf{w}(1), \dots, \mathbf{w}(k), \mathbf{v}(1), \dots, \mathbf{v}(k)\}$ are all assumed to be mutually independent. Equation (22) also has different labels, for instance *sensor model*, *measurement model*, *observation model*, and *measurement equation*.

Many real dynamical systems do not exactly fit the model described by Equations (21) and (22); however, because the Kalman filter is designed to operate in the presence of noise, an approximate fit is often good enough for the filter to be very useful.

Basic Kalman filter algorithm

The Kalman filter is a predictor–corrector algorithm and in each time step it passes two distinct phases: predict and update (correct). The predict phase uses the state estimate from the previous time step to produce an estimate of the state at the current time step. In the update phase, measurement information at the current time step is used to refine this prediction to arrive at a new, more accurate state estimate for the current time step. The state of the filter is represented by two variables: the estimate $\hat{\mathbf{x}}(k+1|k)$ of \mathbf{x} at time $k+1$ given observations up to and including time k , and the error covariance matrix, $\mathbf{P}(k+1|k)$, that is a measure of the estimated accuracy of the state estimate.

In the basic Kalman filter algorithm, the prediction–updating process can be described as follows.

Prediction:

- Prediction of the state a time step ahead, $k+1$:

$$\hat{\mathbf{x}}(k+1|k) = \mathbf{F}(k) \hat{\mathbf{x}}(k|k) + \mathbf{G}(k) \mathbf{u}(k). \quad (23)$$

- Prediction of the state error covariance at instant $k+1$:

$$\mathbf{P}(k+1|k) = \mathbf{F}(k) \mathbf{P}(k|k) \mathbf{F}^T(k) + \mathbf{Q}(k). \quad (24)$$

- Prediction of the observation (measurement) at instant $k+1$:

$$\hat{\mathbf{z}}(k+1|k) = \mathbf{H}(k+1) \hat{\mathbf{x}}(k+1|k). \quad (25)$$

Updating:

- Computation of the *innovation*, \mathbf{J} , that is, the difference between the observation and its prediction, at instant $k+1$:

$$\mathbf{J}(k+1) = \mathbf{z}(k+1) - \hat{\mathbf{z}}(k+1|k). \quad (26)$$

- Computation of the *Kalman gain*, \mathbf{K} at instant $k + 1$:

$$\mathbf{K}(k + 1) = \mathbf{P}(k + 1|k) \mathbf{H}^T(k + 1) \cdot \left[\mathbf{H}(k + 1) \mathbf{P}(k + 1|k) \mathbf{H}^T(k + 1) + \mathbf{R}(k + 1) \right]^{-1}. \quad (27)$$

The factor to the -1 power is the inverse of the *innovation covariance matrix*, $\mathbf{S}(k + 1)$, so Equation (27) can be written in the equivalent more compact form:

$$\mathbf{K}(k + 1) = \mathbf{P}(k + 1|k) \mathbf{H}^T(k + 1) \mathbf{S}^{-1}(k + 1). \quad (28)$$

- Correction of the state estimation at instant $k + 1$:

$$\hat{\mathbf{x}}(k + 1|k + 1) = \hat{\mathbf{x}}(k + 1|k) + \mathbf{K}(k + 1) \mathbf{J}(k + 1). \quad (29)$$

- Correction of the state error covariance at instant $k + 1$:

$$\mathbf{P}(k + 1|k + 1) = \mathbf{P}(k + 1|k) - \mathbf{K}(k + 1) \mathbf{H}(k + 1) \mathbf{P}(k + 1|k). \quad (30)$$

The Kalman gain, a variable that represents the gain which can be taken from the received measurement, relies on the knowledge about the variances of all the previously taken measurements and incorporates the variance of the new measurement in each step. The correction of the state error covariance depends on the Kalman gain definition.

The basic Kalman filter briefly introduced is quite limited because it applies only to *discrete-time processes* and both the state model and the measurement model are *linear*, that is, the process is governed by a linear stochastic difference equation, and the noise of both the process and the measurement is Gaussian. To overcome this limitation, several changes have been introduced. Such is the case, for instance, of the Extended Kalman Filter (EKF), that overcomes the linearity constraint, and the Kalman–Bucy filter, which is a continuous-time version of the Kalman filter. The reader will easily find literature on the different types of Kalman filters (e.g., [36]–[38]).

Applications and software

The estimation of system’s states using past and present observations is a non-linear filtering problem. The Kalman filter is an Minimum Mean Square Error (MMSE) optimal estimator equivalent to the optimal nonlinear estimator for systems described by linear models with Gaussian noise and in this context the Kalman filter is one of the mathematical tools more used to solve problems involving the estimation theory. The domains of applications are thus many, just to name a few: control, telecommunications, macroeconomics, econometrics, medicine and biology, and weather forecasting.

The Kalman filter is also used in multisource (multisensor) estimation. There exist several measurement fusion methods for Kalman-filter-based multisensor data fusion ([39]– [42]). Two examples are (a) merging of the multisensor data through the observation vector of one Kalman filter (centralized Kalman filter fusion); and (b) combination of the observation vectors of the Kalman filters that individually process each sensor’s output data (decentralized Kalman filter fusion).

There is a lot of software for Kalman filtering. MATLAB, in its Control System toolbox, includes the possibility of designing Kalman filters both for steady-state and for time-varying systems or Linear Time-Invariant (LTI) systems with nonstationary noise covariance and also for simulation purposes. Other toolboxes for MATLAB are available from individual contributors (e.g., KALMTOOL, from Magnus Nrgaard). Generic Kalman Filter software, a program written in ANSI C, and O-Matrix with an add-on called Kalman Filter Studio Design from Harmonic Software are alternatives to MATLAB. Mathematica and Excel with add-ons and linkage software can also be used for Kalman filter design and simulation.

Principal components analysis (PCA)

Principal Components Analysis (PCA) is included in a group of techniques that also includes factor analysis and principal co-ordinates analysis. In what follows, we only refer to the more common linear PCA able to capture only the *linear features of the data*, and not to nonlinear PCA that is served by several different algorithms (e.g., [43]–[45]).

One is often confronted with the problem of extracting information about poorly known processes from data. When multivariate data are collected, it is common to find some correlated variables. One implication of these correlations is that there will be some redundancy in the information provided by the variables. In the extreme case of two perfectly correlated variables one is redundant, which means that no significant information is lost if we use only one of them. Principal components analysis exploits the redundancy in multivariate data, enabling us to detect patterns in the variables and to reduce the dimensionality of datasets without a significant loss of information.

The problem of *dimensionality reduction* is often of paramount importance because: (a) the complexity of most algorithms that operate on data increases sometimes exponentially with its size; and (b) the practical implementation of such algorithms requires high-performance processing systems (computers) and even then, the processing time may be prohibitive. Two approaches are available to perform dimensionality reduction: *feature extraction*, consisting in creating a subset of new features by combinations of the existing features, and *feature selection*, which consists in choosing a subset with the most information of all the features. Feature extraction techniques are grouped into two categories: *signal representation*, when the goal of the feature extraction is to represent the data accurately in a lower-dimensional space, and *classification*,

when the goal of the feature extraction is to enhance the class-discriminatory information in a lower-dimensional space. PCA is a *signal representation feature extraction technique*.

Feature extraction refers to identifying the salient aspects or properties of data to facilitate their subsequent use, for instance, in a data fusion process or decision-making. Features are a set of derived variables, functions of the original problem variables, which efficiently capture the information contained in the original data.

Principal components analysis is probably the most common data dimension reduction technique and involves a mathematical procedure that transforms a number of (possibly) correlated variables into a (smaller) number of uncorrelated variables called *principal components*. The first principal component accounts for as much of the variability in the data as possible, and each succeeding component accounts for as much of the remaining variability as possible. PCA is thus a *statistical method* that based on the determination of the eigenvectors and eigenvalues of the data covariance matrix allows the representation of the original data in a lower-dimensional space. The eigenvectors of the covariance matrix are an orthogonal base where the original data can be represented and the space dimension reduction is achieved by ordering the eigenvectors in the order of descending eigenvalues (largest first) and using only the m of the n eigenvectors with larger eigenvalues for data representation.

PCA algorithm

Let us suppose that we have a set of n values of i variables $\mathbf{X} = (x_1^{(1)}, x_2^{(1)}, \dots, x_n^{(1)}, x_1^{(2)}, x_2^{(2)}, \dots, x_n^{(2)}, \dots, x_1^{(i)}, x_2^{(i)}, \dots, x_n^{(i)})$. The PCA algorithm can be summarized as follows.

Step 1 . Mean subtraction: for each variable determine the mean value and subtract it to each i value of that quantity. A new set of values is obtained: $\mathbf{X}' = (x_1^{(1)'}, x_2^{(1)'}, \dots, x_n^{(1)'}, x_1^{(2)'}, x_2^{(2)'}, \dots, x_n^{(2)'}, \dots, x_1^{(i)'}, x_2^{(i)'}, \dots, x_n^{(i)'})$.

Step 2 . Calculate the covariance matrix: each element of this square matrix, \mathbf{C} , is given by

$$C_{j,k} = \frac{\sum_{p=1}^n \left(x_p^{(j)'} - \bar{x}^{(j)} \right) \left(x_p^{(k)'} - \bar{x}^{(k)} \right)}{(n-1)}, \quad (31)$$

where $C_{j,k}$ is the row j , column k element corresponding to the variance between variables j and k , and $\bar{}$ represents the mean value.

Step 3 . Calculate the eigenvectors and eigenvalues of the covariance matrix: the eigenvectors \mathbf{V} and eigenvalues $\boldsymbol{\lambda}$ are such that the following matrix equation is satisfied,

$$\mathbf{C}\mathbf{V} = \boldsymbol{\lambda}\mathbf{V}. \quad (32)$$

Step 4 . Choose components and form a feature vector: according to the eigenvalues, choose the eigenvectors to be used in the representation of the original data. The eigenvectors with higher eigenvalues are the principal components of the original dataset. These eigenvectors constitute the *feature vector*, \mathbf{F} .

Step 5 . Derivation of the new dataset: the new dataset, \mathbf{Y} , is obtained using the selected eigenvectors as follows,

$$\mathbf{Y} = \mathbf{F}^T \mathbf{X}'^T, \quad (33)$$

where T means transposed matrix. Equation (33) reveals how to recover the original data from the new data. Because the inverse of the feature vector is equal to the feature vector transposed, we have:

$$\mathbf{X}'^T = \mathbf{F}\mathbf{Y}. \quad (34)$$

The original set is obtained by adding the mean value of each variable to the corresponding \mathbf{X}'^T values.

PCA applications and software

Principal components analysis is sometimes one of the first steps in data processing. All domains of application can benefit from this processing technique, but image processing is probably one of the most benefited.

A program to implement PCA can be easily developed using any programming language. Nevertheless, software for the purpose is available. MATLAB, for instance, has a Multivariate Statistics toolbox with a function, *princomp*, which can output most of the information and data related with PCA. Many other solutions exist, namely XLSTAT, an add-in for Windows Microsoft Excel, and multivariate statistics packages such as the one available from Kovach Computing Systems.

Artificial neural networks

Artificial neural networks (ANNs) are processing structures made of *artificial neurons*. An artificial neuron is an entity that establishes a pre-defined relationship between its input and its output. Figure 5 represents the so-called scalar neuron with bias. The output o is a scalar given by




Fig. 5. Artificial neural scalar neuron with bias.

$$o = f(i * W + b), \quad (35)$$

where f is neuron's *activation or transfer function* (e.g., hard-limit (McCulloch and Pitts threshold function), linear, sigmoid [46]) whose argument depends on the input scalar i , which is multiplied by a scalar *weight*, W , and on a constant value, b (*bias*).

In artificial neural networks, the neurons' number and transfer function and interconnections between neurons may differ significantly but the overall objective of the network is to materialize a function F between two sets of variables, input \mathbf{X} , and output \mathbf{Y} , $F: \mathbf{X}'\mathbf{Y}$. Figure 6 represents a possible organization of a neural network. The case depicted corresponds to a processing structure whose information moves in only one direction, forward, from the input nodes to the output nodes (one only, in this case). The neurons are organized in layers and the input of each neuron depends on the output value of the neurons to which it is connected. Neurons of the *hidden layer* are of the perceptron type [46, 47]. They are like the one represented in Figure 5 but having several inputs that are multiplied by possible different weights.

The underlying idea of neural networks is that neurons' weights and bias can be adjusted so that the network materializes F or exhibits some desired behavior. The process of adjustment is called *network training* or *learning process*. It consists basically in the following: weights and biases are initialized, for instance, randomly; a set of input values X_t whose output set values Y_t are known, is input to the network and the output of the network, Y , is calculated; Y and Y_t are compared and the weights' and biases' values are changed until a *cost function* is usually minimized. There are three different learning paradigms [46] – *supervised*, *unsupervised*, and *hybrid* – several cost




Fig. 6. Topology of a feedforward neural network. Arrows represent the weights of the neurons.

functions, and several algorithms for training neural network models, most of them employing some form of gradient descent.

In terms of architecture, and taking into account data flow, ANNs can be classified into two types: *feedforward* and *recurrent*. Contrary to feedforward networks (FFANNs), recurrent neural networks (RANNs) are processing structures with bidirectional data flow. Where as a feedforward network propagates data linearly from input to output, RANNs also propagate data from later processing stages to earlier stages. *MultiLayer Perceptron* (MLP), *Radial-Basis* (RBF), and *Kohonen self-organizing networks* are examples of FFANNs [46]; the *Elman Network* or simple recurrent network [48] shown in Figure 7, whose main task is to produce particular output sequences in response to specific input sequences (spatial and temporal pattern learning and recognition capability), is an example of a RANN.

The most adequate ANN to use depends on the problem to solve. Artificial neural networks can be used for many purposes but their main applications may be included in one of the following categories:

- Function approximation: for example, sensor static characteristic modeling, sensor data correction, and disturbance factor compensation, process modeling and control, fault detection and diagnosis.
- Time series prediction/forecasting: for example, dynamic sensor characteristics modeling, electric power forecasting, weather forecasting, water quality forecasting, financial and sales forecasting.
- Classification: for example, acoustic sound recognition (dolphin whistles), target recognition, medical diagnosis, system diagnosis.
- Data mining: for example, clustering, data visualization, data extraction.

Artificial neural networks are processing structures with a huge domain of applications. In many cases they outperform other more traditional processing techniques and algorithms. Theory and practice of neural networks are still under development, but many aspects are already well established. For time-independent applications or when input values are equally spaced in time, and from the several neural network types and architectures, two of




Fig. 7. The Elman network. The output of the network at time t_k depends not only on the inputs I_1, I_2, \dots In at time t_k but also on the output of the hidden layer at time t_{k-1} stored at D . The hidden neuron is usually of the sigmoid type and the output neuron is of the linear type

the feedforward type are clearly favored by users: multilayer perceptron and radial basis function. When the inputs are time-dependent and not equally spaced in time, only networks of the recurrent type are effective for spatial and temporal pattern learning and recognition capability.

MLPs and RBFs are essentially different: (i) MLPs have one or several hidden layers, where as RBFs have only one hidden layer; (ii) MLPs have hidden neurons with nonlinear activation functions (nonlinear hidden layers) and output neurons with linear or nonlinear activation functions (linear or nonlinear output layer), where as RBFs have a nonlinear hidden layer and a linear output layer; (iii) in MLPs, the argument of the hidden neurons is obtained through a scalar product; in RBFs, it is calculated using an Euclidean norm; (iv) whereas MLPs are global approximators, RBFs are local approximators; (v) MLPs use supervised learning and RBFs hybrid learning.

In spite of all these differences MLPs and RBFs are two competing nonlinear function approximation methods because they are both universal function approximators, meaning that given enough neurons they can model any function to the desired accuracy. The tasks they can perform are also basically the same: function approximation, classification, prediction, and control. A question that naturally arises is thus: when should one use an MLP and when should an RBF network be used instead? To the best of our knowledge, there is not a decisive answer to such question. It is well known that some types of applications better fit MLPs and others RBFs but several factors may completely change things: size of the training and testing sets, number of inputs and outputs, number of neurons, or level of confidence required for the output(s). Generally speaking, one can say that RBFs are better classifiers than MLPs when the size of the input data is large; they provide better local approximations with better repeatability but for global approximation the number of neurons must be higher than in an MLP. In terms of training, backpropagation methods used in MLPs, namely the Levenberg–Marquardt algorithm, are particularly efficient but that does not mean that they outperform the hybrid methods of RBFs because these networks tend to be less complex than MLPs.

Some of the major issues of concern today related with artificial neural networks are scalability, testing, and verification. There are several software applications with powerful and user-friendly neural networks toolboxes (e.g., MATLAB) but programs sometimes become unstable when applied to larger problems. Testing and verification is a critical issue for applications in sensitive domains such as the defense, nuclear, and space industries. The mathematical theories used to guarantee the performance of an applied neural network are still under development and no real profit has been taken from the intrinsic parallel processing structure of most of ANNs.

Pulse-Coupled Neural Networks (PCNNs) are a very special type of neural network that resulted from the work in biology mainly of Eckhorn et al. [49]. Later, Johnson [50] and others developed the algorithm proposed by Eckhorn. PCNNs are basically used for image processing and for image fusion (e.g., [51]).

Software

As mentioned before, MATLAB has a Neural Networks toolbox that allows us not only to design and train the most common neural networks but also to design and train user-defined networks. It is a powerful toolbox whose use requires a fair knowledge of artificial neural networks and of the possibilities of each topology. The same can be said about Mathematica and Mathcad solutions.

Clustering

Clustering deals with methods and techniques to find a structure in a collection of *unlabeled data*. A loose definition of clustering is *the process of organizing objects (data) into groups whose members are similar in some way* [53]. A cluster is therefore a collection of objects that are similar among them and are dissimilar to the objects belonging to other clusters. The similarity criterion is either (a) *distance*: two or more objects belong to the same cluster if they are close according to a given distance—distance-based clustering—or (b) *common concept*: two or more objects belong to the same cluster if this one defines a concept common to all those objects—conceptual clustering.

Clustering is included in the unsupervised learning domain. Contrary to supervised classification techniques, in clustering techniques no *a priori information about classes is required*; that is, neither the number of clusters nor the rules of assignment into clusters are known. They have to be discovered exclusively from the given dataset without any reference to a training set.

Cluster analysis is served by many different algorithms. For simplicity, it would be interesting to organize them into types or classes, but, unfortunately, we do not know any classification broadly used or mentioned in the literature. This is an inconvenience in the context of this chapter, because it would simplify the text presentation. Thus, we use one simple classification that probably covers all the known methods, enumerate some of the most common and known techniques, and make more detailed considerations about two of them.

Clustering methods/algorithms

We divide clustering methods into five types: partitioning algorithms, hierarchical algorithms, model-based methods, grid-based methods, and density-based methods [15, 52, 54].

Model-based methods consist in using certain models for clusters and attempting to optimize the fit between the data and the model.

Although generally classified as a special kind of artificial neural network, grid-based methods use *multiresolution grid data structures* and *dense grid cells* to form clusters.

In density-based methods, clustering is based on local criteria such as *density-connected points*.

Partitioning algorithms

Simply put, partitioning algorithms classify n objects into k clusters. Although partitioning is usually associated with nonoverlapping, for simplicity we admit that the clusters may or may not overlap (*exclusive* and *overlapping algorithms*). k -means clustering and its many variations (e.g., adaptive k -means and hard k -means) are of the exclusive type, whereas fuzzy C -means is an example of the overlapping type.

k-means clustering [55] (not to be confused with k nearest neighbor, k -NN, clustering [56, 57]), is probably the most widespread partitioning clustering algorithm. The algorithm can be summarized in the following sequence of steps:

1. Define the number k of clusters in which data (objects) are to be classified.
2. Place k points into the space represented by the objects that are being clustered. These points represent initial group centroids.
3. Assign each object to the group that has the closest centroid.
4. When all objects have been assigned, recalculate the positions of the k centroids.
5. Repeat Steps 3 and 4 until the centroids no longer move.

Figure 8 shows a possible evolution of the cluster centers m_i if a minimum-distance classifier is used, that is, if x is included in cluster i if the distance $\|x - m_i\|$ is the minimum of all the k distances.

Although it can be proved that the procedure will always terminate in a number of iterations smaller than the size of the data, and in spite of its aptitude to accommodate new data, the k -means algorithm does have some weaknesses:

1. It is a numeric data-oriented algorithm. Results depend on the metric used [58]. A possible solution is to normalize each variable by its standard deviation, although this is not always desirable.
2. Results depend on the value of k . To our knowledge, there is no general theoretical solution to find the optimal number of clusters for any given dataset. A possible approach is to compare the results of multiple runs




Fig. 8. Partitioning clustering: cluster centers' possible evolution [53].

- with different k classes and choose the best one according to a given criterion.
3. The initialization is not specified. Usually, k of the objects are randomly chosen.
 4. Results depend on the initial values, and it frequently happens that sub-optimal partitions are found. This problem can be minimized by running the algorithm with different starting points.
 5. It is unable to handle noisy data and outliers.
 6. It is unsuitable when clusters are nonconvex. Fuzzy c -means algorithms [59,60] are often a solution for this problem.

Hierarchical algorithms

Hierarchical algorithms find successive clusters using previously established clusters. Hierarchical algorithms can be further divided into two types: *agglomerate* and *divisive*, the first being by far the most used of the two. Figure 9 shows, in a tree representation (dendrogram), how the original data are processed in an agglomerate hierarchical algorithm. A divisive algorithm would have the same dendrogram but the sequence would be inverse; that is, it would start with the *abcdef* cluster. Cutting a dendrogram at any level defines a clustering and identifies clusters.

Given a set of N items to be clustered and an $N \times N$ distance (or similarity) matrix, the basic process depicted in Figure 9 can be summarized as follows [61]:

1. Assign each item to a cluster (N items, N clusters).
2. Find the closest (most similar) pair of clusters and merge them into a single cluster, so that now you have one cluster less.

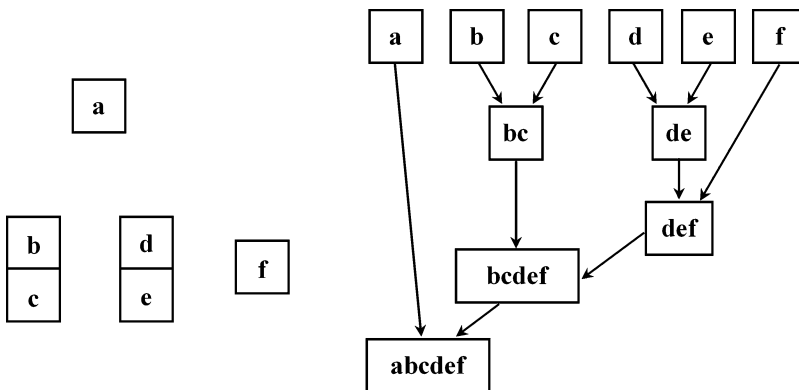


Fig. 9. Original data and dendrogram of hierarchical agglomerate clustering algorithms. The distance metric used is the Euclidian distance.

3. Compute distances (similarities) between the new cluster and each of the old clusters. This can be done in different ways depending on how the distance (similarity) is computed (*single-linkage*, *complete-linkage*, *average-linkage*, *median-linkage clustering*).
4. Repeat steps 2 and 3 until all items are clustered into a single cluster of size N .

The main weaknesses of agglomerative clustering methods are the time complexity, which depends on the square of the number of total objects, the likelihood of the algorithm to be stuck in a local minimum, and the impossibility to accommodate new data.

Applications and software

Clustering is an extremely useful and powerful tool in applications involving pattern recognition (i.e., *automatic identification of patterns in data without human participation in the decision process*, e.g., marketing, planning) and classification (e.g., book ordering in libraries, document classification, and group identification at the World-Wide Web).

Where software is concerned, multivariate statistics packages usually have clustering analysis capability. The Multivariate Statistics toolbox from MATLAB, for instance, is a good example, providing functions for both k -means and hierarchical clustering analysis. Other examples: *statistiXL*, an add-in for Microsoft Excel, Advanced Statistics Module of StatPac, and multivariate statistics packages such as the one available from Kovach Computing Systems.

Support vector machines

Support Vector Machine (SVM) identifies a specific type of algorithms basically for classification purposes that use *supervised learning*. Because they use only part of the data available (vectors)—support vectors—SVMs also help to solve problems associated with the dimensionality, that is, data size. In as much as publications from V. Vapnik (e.g., [62]) and particularly [63] with the help of [64] provide detailed information about SVMs, we restrict ourselves here to an introductory note on the subject.

Although not limited to that case, the basic idea of SVMs is better understood in the context of two-class classification, that is, when data are to be classified in two classes.

SVMs aim at finding the hyperplane separating the two classes that maximizes the margin between it and the examples in the training set, *maximum margin classifier*. To achieve this, it may be necessary to map the training data into another space. In the case of nonlinear classification, that is, when the classes are not separable by hyperplanes, the data should be mapped using kernel functions in a higher-dimensional space where a hyperplane separating the classes exists (see Figure 10).




Fig. 10. Mapping 2D data in 3D transforms the line separating the classes into a plane, simplifying the classification. After classification in 3D, the solution must be mapped back to 2D.

Margin maximization, which uses only part of the data, is important because the higher the margin is, the better the classifier will be, that is, the easier it will be to classify new data. Finding the best hyperplane, that is, the one whose distance to the nearer datapoints is maximal, is an optimization problem usually resolved in the context of *quadratic programming optimization*.

Applications and software

As mentioned before, SVMs are important tools for classification and data size reduction purposes. Applications requiring pattern recognition and decision-making based on data mining or classification (e.g., [65]) are probably those that benefit more from SVMs.

SVMs are currently served by several available software solutions. Some of them, such as the MATLAB Bioinformatics toolbox and LIBSVM, can be found at http://en.wikipedia.org/wiki/Support_vector_machine.

General-purpose software for data mining or classification, such as RandomForests from Salford Systems, also normally has SVM capability.

Self-organizing maps

A *Self-Organizing Map* (SOM) or self-organizing map networks, introduced by T. Kohonen [66], are a special type of unsupervised single layer feedforward artificial neural networks. An SOM is a data reduction and, unintentionally, a clustering technique used basically for pattern recognition, classification, and mainly for *data visualization* purposes. A word about visualization: data with more than three dimensions (3D data) are difficult or even impossible to visualize. SOMs map the input data (feature space) in a 2D space (usually) preserving the distance between clusters in the input data space. This means that neighboring 2D points represent very similar patterns and thus a fairly accurate 2D representation of the input data.

Incremental-learning SOM algorithm

Training 2D-SOMs, those that as mentioned are usually used, consists in the assignment operation from the ensemble of data vectors (input data normalized) to an ordinary two-dimensional array of nodes (output). The number of nodes depends on the discretization level one wants the output representation to have. Nodes are usually arranged in either rectangular or hexagonal configuration. Each node, i has a reference vector \mathbf{m}_i associated. The components of that vector, whose dimension is the same of the input data, are called *weights* and must be initialized, for instance, randomly (not very efficient).

The incremental-learning SOM algorithm is iterative. Upon initialization, the algorithm takes one input vector and looks for the node whose vector is closer, according to some metrics, to that input vector. This node is called the *winner* or *Best Matching Unit* (BMU) and this strategy is called winner-take-all, *competitive learning*. Once identified, the components or weights of the BMU are actualized as well as those of the nodes in its neighborhood more similar to the input vector according to

$$\mathbf{m}_i(t+1) = \mathbf{m}_i(t) + h_{c(x),i}(\mathbf{x}(t) - \mathbf{m}_i(t)), \quad (36)$$

where t is the step index. Index c (“winner”) is defined by the condition:

$$\|\mathbf{x}(t) - \mathbf{m}_c(t)\| \leq \|\mathbf{x}(t) - \mathbf{m}_i(t)\| \forall i. \quad (37)$$

$h_{c(x),i}$ is called the *neighborhood function*. It is like a smoothing kernel that is time-variable (step order) and its location depends on condition (37). It is a decreasing function of the distance between the i th and c th map nodes.

The most used neighborhood functions are:

1. Gaussian:

$$h_{c(x),i} = \alpha(t) \exp\left(-\frac{\|\mathbf{r}_i - \mathbf{r}_c\|^2}{2\sigma^2(t)}\right), \quad (38)$$

where $0 < \alpha(t) < 1$ is the *learning-rate factor* or *function*, which decreases monotonically with the step order, $\mathbf{r}_c \in \mathcal{R}_2^{\mathbb{E}}$ and $\mathbf{r}_i \in \mathcal{R}_2^{\mathbb{E}}$ are the vectorial locations in the display grid, and $\sigma(t)$ corresponds to the width of the neighborhood function, which decreases monotonically with the step order.

2. Circular or bubble:

$$\begin{cases} h_{c(x),i} = \alpha(t) & \text{if } \|\mathbf{r}_i - \mathbf{r}_c\| \text{ is smaller than a given radius} \\ h_{c(x),i} = 0 & \text{otherwise} \end{cases} \quad (39)$$

The assignment of a BMU to one input vector and the actualization of the nodes’ weights go on until all input vectors have been assigned to a node or until a predefined number of iterations has been reached.

A positive aspect of SOMs is their good performance. They classify data well and it is easy to evaluate their own quality so that one can actually calculate how good a map is and how strong the similarities between objects are.

SOMs have, however, some drawbacks or limitations of which one has to be aware. The first problem that may come up has to do with the fact that to generate a map, a value for each component of each input vector is needed. Sometimes this is not possible and often it is very difficult to acquire all of the required data. This limiting feature to the use of SOMs is often referred to as *missing data*.

Another problem is that every SOM is different and finds different similarities among the sample vectors. SOMs organize data so that at the end, the samples are usually surrounded by similar samples. However, similar samples are not always near each other and it is not always obvious that the map is not a good map.

Finally, SOMs are very computationally expensive, which is a major drawback because as the dimensions of the data increase, dimension reduction visualization techniques become more important, but unfortunately then computation time also increases.

Applications and software

As already mentioned, although initially designed as a visualization tool for large-dimensional data, SOMs have been largely used for pattern recognition and classification, namely of high-dimension data. In fact, the properties of these networks make them a fairly good classifier. The classification of a new vector once the SOM is trained consists in the determination of the BMU.

A neural network-based clustering tool in Excel using SOMs is available from Angshuman Saha at <http://www.geocities.com/adotsaha/NN/SOMinExcel.html>. Databionic ESOM Tools is also an interesting free software available at <http://databionic-esom.sourceforge.net>.

Visualization and data-mining-oriented software usually have the capability to train and use trained SOMs. The MATLAB Neural Network toolbox and Mathematica Machine Learning Framework are, however, probably the more accessible solutions for someone wishing to design and use self-organizing maps.

Voting

Voting is probably the oldest fusion decision method. It is also the most basic, the one to which one recurs when having little information about the data to fuse. The method can be used at different stages of the fusion process and, as the name suggests, the basic idea is to obtain an output by weighing inputs in some way. Weights are usually attributed as votes. Common voting methods are:

- Majority: each input has a certain number of votes and the output is derived from the most voted input.
- Plurality: as majority, but the output is derived from the more voted inputs.
- Consensus: an output is obtained only if all the inputs are equal under a certain criterion.
- Weighted: votes or weights are given to each input and the output is derived by a weighed combination of all the inputs.

The votes or weights must be attributed according to some prior knowledge. When possible, the use of probabilities is an interesting solution because it allows the updating of votes/weights (e.g., using Bayes rule) as new data are obtained.

Inference

Inference is the act or process of deriving a conclusion based solely on what one already knows. There are basically two forms of inference: inductive and deductive. In deductive inference one progresses from certain premises to certain conclusions. In inductive inference, one tries to establish general principles from a limited number of observations. Both are important in decision-making and risk analysis and it is also in this context that the present paragraph is included in this chapter.

From the “mathematical approach” point of view, several approaches are possible. We mention here three: the Bayesian approach, the Dempster–Shafer method, and the fuzzy-based approach. They correspond to solutions better suited for different types of problems, depending on the knowledge one has of them and how that knowledge can be quantified. Each approach represents a different way of expressing uncertainty.

Bayesian inference

Bayesian inference is a probability-based statistical inference in which *evidence* or *observations* are used to update or newly to infer the probability that a hypothesis may be true. Bayesian inference is based on the well-known *Bayes’ rule*,

$$P(A|B) = \frac{P(B|A) P(A)}{P(B)}, \quad (40)$$

where $P(A)$ and $P(B)$ are the probabilities of A and B and $P(A|B)$ and $P(B|A)$ are the probabilities of A given B and of B given A , respectively. If variables A and B are independent, $P(A|B) = P(A)$ and $P(B|A) = P(B)$, and Equation (40) is equivalent to saying that $P(A, B) = P(B, A)$. In that case, B does not affect $P(A)$.

It is common to think of Bayes’ rule in terms of updating our belief about a hypothesis A in the light of new evidence B . Specifically, our *posterior belief*

$P(A|B)$ involves the multiplication of our *prior belief* $P(A)$ by the *likelihood* $P(B|A)$ that B will occur if A is true. $P(B)$ is called *evidence*. The power of Bayes' rule is that in many situations where we want to compute $P(A|B)$ it turns out that it is difficult to do so directly, yet we might have direct information about $P(B|A)$. Bayes' rule enables us to compute $P(A|B)$ in terms of $P(B|A)$.

From Equation (40) it is clear that $P(A|B)$ requires all other three probabilities. However, when comparing possibilities (different A) in the light of the same new evidence, $P(B)$ is irrelevant because it is the same for all A . Note also that each time we calculate $P(A|B)$, that value can be used as a prior for posterior updating upon a new evidence.

The objective of Bayesian inference is to support a decision of choosing among a set of N different alternatives A_i , based on an observation B . In this case, for each A_i , $P(A_i|B)$ is calculated and the decision made according to the N values obtained. When the decision is based on multiple observations, M (e.g., multiple sensors measurements) then B is a vector with M elements. For each i of the N alternatives, $P(A_i|(B_1, B_2, \dots, B_m))$, Equation (40) becomes:

$$P(A_i|B_1, B_2, \dots, B_M) = \frac{P(A_i) \prod_{m=1}^M P(B_m|A_i)}{\sum_{i=1}^N \prod_{m=1}^M P(B_m|A_i)} \quad (41)$$

and, as before, it is the N values obtained by calculating Equation (41) for the N alternatives that support the decision.

Dempster–Shafer model

Sometimes, situations cannot be described by a full set of probabilities because some values are unknown. The *Dempster–Shafer approach* is also a statistical one but that takes into consideration situations when the level of uncertainty is not compatible with a description through classic probabilities and that some of the probabilities are thus not known, or are in conflict.

The Dempster–Shafer model ([67]–[69]) is a *belief-based model* but there are other models of this type also identified by *upper and lower probabilities models* [70]. These models' aim is to model someone's degrees of belief and for that purpose a function or functions are introduced.

In any model for belief two components must be considered: one that describes the state of belief and the other that defines how to update the belief given new information. Dempster calls the *plausibility function* the upper probabilities and *belief function* the lower probabilities and the values of these functions can be calculated from a combination of the probabilities assigned to individual events or statements of the problem's state space. As probabilities are the support of a decision in Bayesian inference, plausibility and belief values support a decision when a Dempster–Shafer inference is used.

It is out of the question to give a detailed description of the Dempster–Shafer model here or even an exemplification of how it works (e.g., [70]), the

best way to understand it. To readers interested in the upper and lower probabilities model and in particular in the Dempster–Shafer model we suggest the reading of [71].

Fuzzy inference

Fuzzy inference is the process of formulating the mapping from a given input to an output using fuzzy logic. It tackles problems characterized by vagueness and ambiguity when it is not possible to assign a single value or a probability to a quantity or statement. Fuzzy logic uses fuzzy sets introduced by Zadeh [72, 73], that is, sets without a crisp, clearly defined boundary. They can contain elements with only a partial degree of membership expressed by a *membership function* whose values are reals in $[0,1]$.

Fuzzy inference consists of basically five distinct steps:

- *Initialization*: definition of the fuzzy variables and corresponding membership functions and of the *rules* and *hedges*.
- *Fuzzification*: mapping from “crisp” numerical values to the membership functions of the fuzzy variables.
- *Rule evaluation*: the rules defined are evaluated using a fuzzy set logic.
- *Aggregation*: the results of the rules are aggregated so that they are mapped to the output variables.
- *Defuzzification*: mapping from fuzzy output variables to crisp numerical values.

The *fuzzy set logic* is similar to normal Boolean logic but because variables do not have only values 0 or 1, the definition of operators must be redefined. The three basic operators in fuzzy logic are OR, AND, and NOT defined as follows.

$$\begin{aligned} A \text{ OR } B &= \text{Max}(A, B) \\ A \text{ AND } B &= \text{Min}(A, B) \\ \text{NOT } A &= 1 - A \end{aligned}$$

A *fuzzy rule* is assembled as follows.

IF *antecedent* THEN *consequent*

According to the structure of the consequent, fuzzy inference systems are either of the *Mamdani* (also labeled Zadeh–Mamdani) type [74] or *Sugeno* (also labeled Sugeno–Takagi) [75] type.

A *hedge* is an operator that is applied on a membership function and defined in such a way to reflect their linguistic meaning. Figure 11 shows the change of the membership function associated with the height of a person when hedges “very short” and “very tall” are applied (the operator used was “power of two”).

In fuzzy inference, the value or values obtained upon defuzzification are the support of the decision.




Fig. 11. Change in membership function associated with the height of a person when hedges very short and very tall are applied.

Software

Calculations involved both in basic Bayesian and Dempster–Shafer inferences are fairly easy to implement. Nonetheless, software like Gister from SRI International provides functionalities that may be useful. Fuzzy inference algorithms are less easy to program, but both Mathematica and MATLAB have a well-documented package/toolbox that makes user’s life easier.

1.5 Implementing data fusion

Level and architecture

In Sections 1.1 and 1.2 we saw that data fusion can be made at different levels using different architectures. The consideration of those aspects is obviously important, in particular from the conceptual, taxonomic, and organizational points of view, but practically less pertinent. In fact, data fusion serves applications and normally they “tell you” what should be done. This is evidenced with the two examples that we present next.

Technique/algorithm selection

The choice of the technique or algorithm to use is often not easy because many constraints can exist. Time, programming resources available, technical capacity, and expertise are only some of possible factors to take in consideration. Some tips that can help: (a) investigate those solutions that are more generally followed to resolve the problem in the specific domain of application you face and try to evaluate which is the best; (b) try to use techniques and algorithms that you know by your own experience and whose result you are able to control; and (c) if you are looking for the best solution, you have to compare the results you obtain with each solution and choose accordingly.

Commercial software

In the paragraphs of Section 1.3, and when we judged pertinent, we proposed software that implements the techniques and algorithms considered then. Some of the references are to multipurpose software packages (e.g., Synapse from Peltarion Corporation) and most of the referenced software is accompanied either by a helper or a tutorial to assist the user. Nonetheless, in our opinion, Mathworks' MATLAB (<http://www.mathworks.com>), followed by Wolfram Research Mathematica (<http://www.wolfram.com>), is the best overall solution because it is able not only to answer the user needs, but also because the documentation it provides is exhaustive and usually of high quality. Readers are invited to verify this by themselves. Another advantage of MATLAB is the interfacing capabilities with other software packages, namely with Microsoft Excel (Mathematica also has the same possibility) by means of MATLAB Excel link toolbox and MATLAB Builder for Excel and LABVIEW from National Instruments, a worldwide used software for data acquisition systems.

Examples

The two examples that follow are presented to illustrate the use of some of the techniques of the past paragraphs and to support the comments made then.

Example 1: Air pollution detection using gas sensors

Suppose that we want to detect air pollution based on the measurement of a set of N gases and vapors using low-cost gas sensors. The gas sensors are not very selective (i.e., the output of a sensor depends on more than one gas/vapor) and temperature and humidity are influence quantities. Pollution detection depends on the interpretation of the outputs of all the sensors and thus we are in the presence of a multisensor fusion problem involving (at least) and by this order the following steps:

1. Temperature and humidity correction of sensors outputs—sensor inverse modeling with influence quantities correction (compensation).

With additional temperature and humidity sensors, one could implement either: solution 1—one neural network whose inputs were the normalized values of temperature, humidity, and sensor outputs and whose output would be the N temperature and humidity corrected values of the sensor outputs; or solution 2— N neural networks, one for each sensor. The input of each network would be temperature, humidity, and sensor’s output and network’s output would be the temperature and humidity corrected value of the correspondent sensor. The networks have to be trained using the values output by the sensors under temperature and humidity controlled conditions when submitted to calibrated samples of the gases and vapours at stake.

Due to the poor selectivity of the sensors, solution 2, although more complex and computationally expensive, would be better.

2. Correction of sensors’ poor selectivity. The problem is to know what percentage of a sensor’s output is due to each gas/vapor. This is equivalent to looking for the values of the elements of the matrix α_{ij} that relates the concentrations C_j with sensor outputs O_i ,

$$\begin{bmatrix} O_1 \\ \dots \\ O_i \\ \dots \\ O_N \end{bmatrix} \begin{bmatrix} \alpha_{11} & \alpha_{12} & \dots & \dots & \alpha_{1N} \\ \dots & \dots & \dots & \dots & \dots \\ \alpha_{i1} & \dots & \alpha_{ij} & \dots & \alpha_{iN} \\ \dots & \dots & \dots & \dots & \dots \\ \alpha_{N1} & \alpha_{N2} & \dots & \dots & \alpha_{NN} \end{bmatrix} \begin{bmatrix} C_1 \\ \dots \\ C_j \\ \dots \\ C_N \end{bmatrix}. \tag{42}$$

The experimental values gathered for temperature and humidity compensation mentioned before could be used to calculate either the matrix of Equation (42) or its inverse. This last matrix would allow the fusion of sensor outputs to obtain the gases/vapors concentration in what can be classified as a weighted voting technique.

3. Fusion of gas and vapor concentration values to classify the quality of the air. Three possible solutions would be: *solution a*, fuzzy inference and defining a rule for each gas/vapor measured; *solution b*, a clustering technique; or *solution c*, a neural network with $N + 2$ inputs and 2 or more outputs, depending on the number of classification levels desired. Solutions b and c require sets of values to define the clusters and to train the network, respectively.

Example 2: Autonomic nervous system analysis

The status of the Autonomic Nervous System (ANS) is a general health indicator. Several tests produce values that when isolated allow only reduced and inconclusive knowledge of the ANS condition. Fusing the information one increases that knowledge and medical and clinical conclusions can be drawn. Let us assume that the electrocardiogram (ECG) and the blood pressure (BP)

are two tests that allow the extraction of features that once fused produce a result directly related with the ANS status. Then, to go from the ECG and BP to the ANS status assessment one had to:

1. Extract, from the ECG and from the BP, the features to fuse: heart rate variability (HRV) and blood pressure variability (BPV). This would require the following [76]:
 - a) Determination of beat-to-beat time interval variations (R-R variation, in medical terms) and blood pressure amplitude variation. If the start data are ECG and BP sampled values, this extraction would involve:
 - Filtering of both signals, namely to eliminate DC components. A bandpass digital filter is a possible solution.
 - Identification of the peaks of the ECG and BP filtered signals, and calculation of the time differences between consecutive peaks in the ECG case and of the amplitude difference of consecutive peaks in the case of the BP. Peaks should be defined according to some criterion, for instance using a gating technique or by simply defining a threshold. The time intervals Δt and amplitude variations ΔA would be the new signals to process.
 - b) Decomposition of Δt and ΔA into high- and low-frequency components. For this purpose an FFT could be used, but our experience shows that wavelet techniques perform better [77]. Once decomposed, the ratio between the low-frequency and high-frequency power densities of Δt and ΔA should be calculated and would constitute the two features, HRV and BPV, to be further processed.
2. Feature fusion. How to fuse HRV and BPV is a delicate matter because it depends on how the medical class values each of these indicators, and that is at least controversial. Thus, and to more easily adapt to different opinions, we would propose using fuzzy inference.

2 Decision-making

In domains such as health, business, or engineering, the aspects related to decision-making and risk analysis, although often related, can be dealt with separately using different techniques. Each technique is characterized by a particular way of processing information. However, information understood as organized data, is not always available. Instead, the support of such techniques is data gathered by many different ways: through experience, through measuring, estimated, and so on. In this section, we look into some frameworks for decision-making.

Decision-making may be defined as *the cognitive process leading to the selection of a course of action among alternatives*. Every decision-making process produces a final choice. Decision-making is anchored in many areas of knowledge, *decision theory* being one of them. This section provides a

short incursion into decision analysis by presenting some general-purpose techniques/procedures that because supported in basic mathematics and statistics, are interesting tools for decision-making. Before, though, we have a few comments.

1. When facing a decision problem, two basic situations may occur: (a) all the alternatives are completely known and quantified; (b) alternatives are only identified. In the first case, a mathematical (e.g., based on a model) or statistical based decision is possible and the process of decision-making consists in basically either: (1) the calculation of numerical parameters chosen to support the decision and on the rules of decision (this often involves probabilities); or (2) the resolution of an optimization problem which often consists in finding a solution of an objective function under a set of constraints in the form of a system of equations or inequalities (*linear programming*). When alternatives are only identified, the first step is to assign numerical values to the possible outcomes, which is the domain of the *utility theory* [78]. After that, we are either under the same conditions of alternative (a) or we face what can be classified as an ill-posed problem.
2. One should not identify a good decision with a good outcome. The classification of a decision is meaningful only in the context of the information and data in which it is made. A good decision yields a good outcome only if the framework in which the decision is made correctly describes the problem to solve as identified by the user.
3. In the paragraph on inference of Section 1 three mathematical tools to support decision-making were presented: Bayes, Dempster–Shafer, and fuzzy inference. The values output by these tools are often used to define decision rules or criteria. To account for the risk of each alternative, the rules are based on indicators that weigh both the expected merit of the alternative and its risk.

2.1 Special decision-making tools and techniques

Cost-benefit analysis

Cost/benefit analysis is a powerful, widely used tool for helping decision-making. The underlying idea is to weigh the total expected costs against the total expected benefits of each alternative. Thus, this technique is particularly suited to decide upon alternatives where the factors affecting the decision can be quantified in terms of currency. The procedure involves:

- Calculation of all the costs of each solution
- Calculation of all the benefits of each solution
- Calculation, for each solution, of the time it will take for the benefits to repay the costs (*payback time or break-even point*)

According to the results obtained, the best solution may not be unique. Payback time can be more important than the benefit-cost differential (e.g., when a loan and respective interests are a heavy burden).

BRAND

BRAND is a mnemonic whose letters are the initials of the following words: benefits, risks, alternatives, nothing, decision. These words identify the five steps of a procedure to decide upon a solution:

1. Specification, if possible with quantification, of the benefits of the solution
2. Specification, if possible with quantification, of the risks of the solution
3. Specification, if possible with quantification, of the alternative solutions
4. Evaluation of doing nothing
5. Decision-making

What is really peculiar of this procedure is that “nothing” is an alternative and thus the process of decision-making can lead to the conclusion that the best solution is to do nothing, keeping things as they are.

SWOT analysis

SWOT stands for strengths, weaknesses, opportunities and threats. In SWOT analysis, each alternative solution is characterized by the same four above-mentioned parameters, two related with the solution itself, with its absolute value, positive aspects (strengths), negative aspects (weaknesses), and two that cater to the favorable (opportunities) and unfavorable (threats), aspects that are external to the solution but condition it.

SWOT analysis is mostly used as a strategic planning tool. The implementation may be based on feeling in templates such as the one shown in Table 1 for each alternative. By assigning numeric values and weights to each criterion, a number representing the quality of the alternative is obtained.

Table 1. Example of a SWOT analysis template.

Criteria	Strengths	Weaknesses	Criteria
s1			w1
s2			w2
s3			w3
⋮			⋮
Criteria	Opportunities	Threats	Criteria
o1			t1
o2			t2
o3			t3
⋮			⋮

Pareto analysis

Also known as the 80–20 rule [79], Pareto analysis is a technique based on Pareto’s so-called 80–20 rule that can be expressed in the context of solution-finding as: by doing 20% of work you can generate 80% of the advantage of doing the entire job. Pareto’s distribution whose probability distribution function is given by

$$f(x; k, x_m) = \frac{k}{x_m} \left(\frac{x_m}{x} \right)^{k+1} \quad \text{for } x \geq x_m \quad (43)$$

is the support for phenomena that are observable in different domains. The Pareto distribution, also known as the Bradford distribution, tends to the Dirac function $\delta(x - x_m)$ when k tends to ∞ and can be useful to be fused with other distributions for data approximation (e.g., [80]).

Pareto analysis is mainly oriented to find the solution to a problem by (a) identifying and assigning values (scores) to the various components that contribute to the problem, and (b) selecting a limited number of tasks that produce significant overall effect. The sequence of Pareto analysis is then:

- List the problems or the options available.
- Group options where they are facets of the same larger problem.
- Assign an appropriate value to each group.
- Work on the group with the highest score.

Pareto analysis not only identifies the most important problem to solve, but it also provides values revealing how severe the problem is.

Paired compared analysis

Paired comparison analysis is a technique that allows the relative comparison of alternatives in the absence of objective data. Thus, it is particularly suited to support decisions on completely different alternatives (e.g., acquisition of different types of infrastructures).

The step sequence of implementation of the technique is:

- Listing of the options to compare. Assign a letter to each option.
- Marking of the options as row and column headings on a table, preferably using the same order.
- Blocking of the cells that correspond to the same row and column headings (marked with $(*)$ in Table 2) and of the cells where information would be duplicated (marked with $(**)$ in Table 2).
- In the remaining cells, a comparison between the alternative in the row with the one in the column is introduced under the following format: letter identifying the more important alternative, number, for instance on a scale from 0 to 5, representing the difference (0 no difference) in importance between the two alternatives.
- Addition of the total score for each alternative and eventual conversion into percentage. The results indicate the relative merits of the alternatives.

Table 2. Paired compared analysis. In the example shown, alternative A is the best: Alternative A = 7 (43.75%); Alternative B = 1 (6.25%); Alternative C = 5 (31.25%); Alternative D = 3 (18.75%).

	Alternative A	Alternative B	Alternative C	Alternative D
Alternative A	(*)	A,3	C,1	A,4
Alternative B	(**)	(*)	C,4	B,1
Alternative C	(**)	(**)	(*)	D,3
Alternative D	(**)	(**)	(**)	(*)

Table 3. Grid analysis: unweighed cells (top); weighed cells (bottom).

Factors:		F1	F2	F3	...	Fn	Total
Weights:		S1	S11	S12	S13	...	S1n
		S2	S21	S22	S23	...	S2n
		...	S31	S32	S33	...	S3n
		Sm	Sm1	Sm2	Sm3	...	Smn

Factors:	F1	F2	F3	...	Fn	Total
Weights:	W1	W2	W3	...	Wn	
S1	S11*W1	S12*W2	S13*W3	...	S1n*Wn	$\sum(S1i*Wi)$
S2	S21*W1	S22*W2	S23*W3	...	S2n*Wn	$\sum(S2i*Wi)$
...
Sm	Sm1*W1	Sm2*W2	Sm3*W3	...	Smn*Wn	$\sum(Smi*Wi)$

Grid analysis

Grid analysis (also known as decision matrix analysis, Pugh matrix analysis, or MAUT, which stands for MultiAttribute Utility Theory [81]) is a technique that allows comparing both objective and subjective data. It is particularly effective when the decision involves a large number of good alternatives and many factors have to be taken into account.

Table 3 summarizes the grid analysis procedure:

- Construction of a table with m rows (solutions) and n columns (factors to take into account).
- Filling the table with values $S_{ij}(i = 1, \dots, m; j = 1, \dots, n)$ each representing the degree factor j is satisfied by solution i (unweighed cells).
- Filling the weights row with values representing the importance given to each factor in the final decision.
- Recalculation of the table: $S_{ij} \times W_j$ (weighted cells).
- Calculation of each solution total score adding all the values in a row.
- The solution with highest score is the best one.

Decision trees

Decision trees, like Bayesian networks [82], are *graphical models* that can be constructed to help decision-making. Perhaps the more general implementation of a decision tree involves the following sequence of procedures:

- Decision tree drawing

The tree starts with a decision that is under evaluation. A symbol is drawn to represent this root of the tree. We suppose that the symbol used is a small square.

From this box are drawn lines, one for each possible alternative and the alternative is identified along the line. Lines should be kept apart as far as possible.

At the end of each line, a result must be considered. If the result of taking that decision is uncertain, a symbol to identify it is drawn. We suppose that the chosen symbol is a small circle. If the result is that another decision needs to be made, another square is drawn. If the solution is completed at the end of the line, it is blanked.

Starting from the new decision squares on the diagram, lines representing the options that can be selected are drawn. From the circles, lines representing possible outcomes are drawn again with some identification of their meanings. The process is repeated until all possible outcomes and decisions leading from the original decisions are considered.

- Tree evaluation

A value is assigned to each possible outcome by estimating its merit.

Next, and going up the tree in the direction of the root, each circle (representing an uncertainty point) is identified and the probability of each outcome estimate. Using percentages, the total must come to 100% at each circle. Using fractions, they must add up to one.

- Calculation of tree values

Starting from the end of the tree, each node value is calculated and written down by the node.

- Calculation of the value of uncertain outcome nodes: multiplication of the value of the outcomes by their probability. The total for that node of the tree is the total of these values.
- Calculation of the value of decision nodes: upon writing down the cost of each option along each decision line, the cost is subtracted from the outcome value already calculated.

When all the decision benefits are calculated, the decision falls on the alternative with largest benefit. Figure 12 shows the tree of a multiple hypothetical decision involving two alternatives A and B.

2.2 Final notes on decision-making

Decision-making is something that is in everyday life and affects all fields of human activity. Boosted mainly by economics, decision theory and decision analysis are disciplines that aim to provide the theoretical background for decision-making anchored in different subdomains of mathematics and probability theory. The framework and tools adequate to a decision-making can be quite difficult to handle in particular because of the increasing objective of

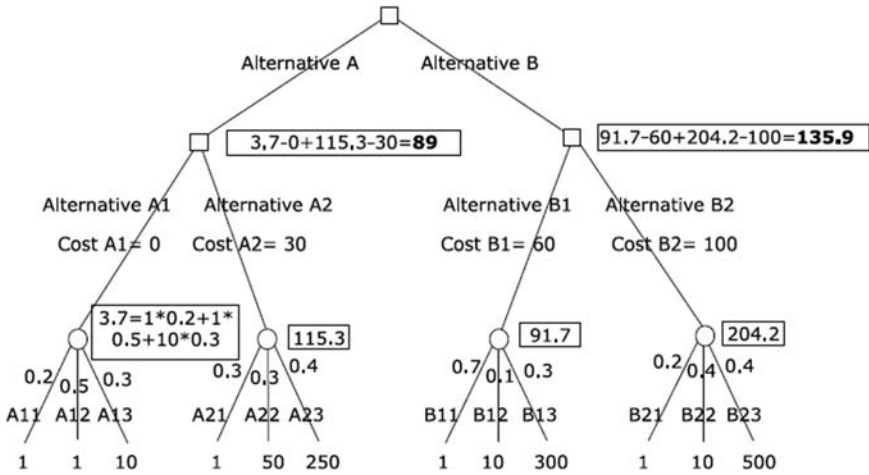


Fig. 12. Decision tree for the selection of one of two alternatives, A and B. Trees draw from left to right are more common.

reaching robust decisions [83]. We mentioned only a few. Many others could be mentioned (e.g., weighing pros and cons, PMI (plus/minus/interesting), force field analysis). One thing they have in common: once the data in which they are supported are known (not always easy), they are fairly easy to implement. Even so, it is possible to find commercial software to help the user (e.g., DPL Fault tree from Syncopation Software) particularly in the evaluation of the influence of the change of the parameters in the decision (*sensitivity analysis*) (e.g., Decision Tool Suite from Palisade). Algorithms and programs operate on data (the user is often a data source for a priori knowledge) and thus one must be aware that decision-making quality is highly dependent on the quality of data that supports it.

3 Risk analysis

Risk is an identified but not certain event leading to negative consequences (i.e., loss). Risk is intimately related with decision-making because risk is an important factor that conditions a decision. This means that if the expected value (mean) of the outcome of a decision can be calculated, it is a possible indicator of the quality of a decision, but it must be complemented with another statistical parameter, the standard deviation, because the value of this parameter is a measurement of risk. For this reason, the quality of a decision taking into account risk is better expressed by the coefficient of variation defined, when the mean is not zero, as the absolute value, expressed in percentage, of the ratio between the standard deviation and the mean.

In engineering terms, the total risk R_{total} of a set of events can be expressed as

$$R_{\text{total}} = \sum_i L_i p_i, \quad (44)$$

where L_i is the cost of or due to the event (magnitude of the potential loss) and p_i is the probability of the event (probability of the potential loss). The summation is extended to all events causing risk.

More generally (Bayes risk), the risk function of an estimator $\delta(x)$ for a parameter θ , calculated from some observables x , $R(\theta, \delta x)$, is defined as the expectation value of the loss function $L(\theta, \delta x)$,

$$R(\theta, \delta x) = \int L(\theta, \delta x) \times f(x | \theta) dx. \quad (45)$$

Risk analysis [84, 85] comprises three components; two result from Equation(44): identification of risks (*risk assessment*), and evaluation of their consequences (*risk evaluation*); the third deals with solutions to minimize risk consequences (*risk management*).

The identification of events that can lead to negative consequences (threats, Table 1) requires a careful analysis of all their potential sources. Some examples by no special order are human (e.g., illness, death); natural (e.g., weather, natural disasters); technical (e.g., technical failure, new technologies); project (e.g., cost overruns, jobs taking too long); operational (e.g., loss of access to essential assets, failures in distribution); reputational (e.g., loss of business partner or employee confidence, reputation damage); procedural (e.g., failures of accountability, organization); financial (e.g., interest rates, unemployment); on political (e.g., changes in tax regimes, foreign influence). The diversity of threats recommends the involvement of more than one person in their identification.

According to Equation (44), risk evaluation involves the estimation of the probability of each threat and of the cost of the potential loss. This estimation can be helped by sensitivity analysis using simulation tools based, for instance, in the Monte Carlo method [85, 86].

Risk management deals with the solutions to minimize risk. Some examples are as follows.

- Contingency plans: the risk is accepted but a plan to minimize its effects if it happens is developed. In a good contingency plan, action should be taken immediately.
- Involvement of additional resources.
- Inclusion of redundancy.
- Risk insurance, which is sharing part of the risk with somebody else.

Final note on risk analysis

Risk analysis may start within the context of a decision process and have many points of contact with decision analysis but has its own identity. Identification

of risk threats and the quantification of their total cost (total risk) is of paramount importance to decide on how to manage risk.

In what concerns software, commercial software for decision-making, such as the one from Palisade mentioned in the decision-making section, usually includes risk analysis capability.

References

1. D.L. Hall and J. Llinas (Editors). *Handbook of Multisensor Data Fusion*. The Electrical Engineering and Applied Signal Processing Series, CRC Press LLC, Boca Raton (FL), 2001
2. R.R. Brooks and S.S. Iyengar. *Multi-Sensor Fusion. Fundamentals and Applications with Software*. Prentice Hall PTR, Upper Saddle River, NJ, 1998
3. D.L. Hall and S.A.H. McMullen. *Mathematical Techniques in Multisensor Data Fusion*. Artech House, Norwood, MA, 2004
4. R.P.S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2007
5. I.R. Goodman, R.P. Mahler and Hung T. Nguyen. *Mathematics of Data Fusion*. Kluwer Academic, Dordrecht, 1997
6. L. Wald. Some terms of reference in data fusion. *IEEE Transactions on Geosciences and Remote Sensing*, 37, 3, 1190–1193, 1999. Also at: <http://www.data-fusion.org/>
7. B. Dasarathy. *Decision Fusion*. IEEE Computer Society Press, Piscataway, NJ, 1994
8. R. Kumar, M. Wolenetz and B. Agarwalla, JunSuk Shin, Phillip Hutto, Arnab Paul, and Umakishore Ramachandran. DFuse: A framework for distributed data fusion. *Proceedings SenSys'03*, pp. 114–125, Los Angeles, November 2003
9. J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz and F. White. Revisiting the JDL Data Fusion Model II. *Proceedings 7th Intl. Conf. on Information Fusion*, pp. 1218–1230, Stockholm, 2004
10. E.P. Blasch. Sensor, user, mission (SUM) resource management and their interaction with level 2/3 fusion. *Proceedings 9th Intl. Conf. on Information Fusion*, Florence, Italy, 2006
11. B. Dasarathy. Sensor fusion potential exploitation-innovative architectures and illustrative applications. *IEEE Proceedings*, Vol. 85, N.1, 1997
12. M. Bedworth and J. O'Brien. The omnibus model: A new model of data fusion? *Proceedings 2nd Intl Conf. Information Fusion*, 1999
13. J. Salerno. Information fusion: A high-level architecture overview. *5th Intl Conf. Information Fusion*, 2002
14. D. L. Hall. *Mathematical Techniques in Multisensor Data Fusion*. Artech House, Norwood, MA, 1992
15. G. Gan, C. Ma and J. Wu. *Data Clustering: Theory, Algorithms, and Applications*, ASA-SIAM Series on Statistics and Applied Probability, Philadelphia, 2007
16. E. Blasch, M. Pribilski, B. Daughtery, B. Roscoe and J. Gunsett. Fusion metrics for dynamic situation analysis. *SPIE 04*, Vol. 5429, pp. 428–438, 2004

17. Y. Kosuge and T. Matsuzaki. The optimum gate shape and threshold for target tracking. *Proceedings Society of Instrument and Control Engineers*, pp. 2152–2157, Fukui, Japan, 2003
18. S.S. Blackman. *Multiple Target Tracking with Radar Applications*. Artech House, Norwood, MA, 1986
19. J.W. Cooley and J.W. Tukey. An algorithm for the machine computation of the complex Fourier series. *Mathematics of Computation*, Vol. 19, pp. 297–301, April 1965
20. T. Theul. Sampling and Reconstruction in Volume Visualization. See at: <http://www.cg.tuwien.ac.at/~theuss1/DA/thesis.html>
21. R. Tolimieri, Myoung An and Chao Lu. *Mathematics of Multidimensional Fourier Transform Algorithms*. Springer-Verlag, New York, 1997
22. B.P. Bogert, M.J.R. Healy and J.W. Tukey. The frequency analysis of time series for echoes: cepstrum, pseudo-autocovariance, cross-cepstrum, and saphe cracking. *Proceedings of the Symposium on Time Series Analysis (M. Rosenblatt, Ed)*, Chapter 15, 209–243. John Wiley and Sons, New York, 1963
23. A.V. Oppenheim and R.W. Schafer. From frequency to quefrequency: A history of the cepstrum. *IEEE Signal Processing Magazine*, Vol. 21, No. 5, pp. 95–106, September 2004
24. L.R. Rabiner and R.W. Schafer. *Digital Processing of Speech Signals*. Prentice Hall Inc., Englewood Cliffs, NJ, 1978
25. G. Lu. *Indexing and Retrieval of Audio: A Survey*. Multimedia Tools and Applications, Volume 15, Number 3, Springer Netherlands, December 2001
26. I. Daubechies. *Ten Lectures on Wavelets - 2nd Edition*. SIAM, Philadelphia, 1992
27. Y. Meyer. *Wavelets and Operators: Volume 1*. Cambridge University Press, Cambridge, UK, 1992
28. C. Blatter. *Wavelets: A Primer*. A.K. Peters, Natick, MA, 1998
29. C.S. Burrus, R.A. Gopinath and H. Guo. *Introduction to Wavelets and Wavelets Transforms*. Prentice Hall, Upper Saddle River, NJ, 1998
30. S. Mallat. *A Wavelet Tour of Signal Processing - Second Edition*. Academic Press, San Diego, CA, 1999
31. K. Saastamoinen and J. Sampo. Use of hybrid method in wavelets bases selection for signal compression, mathematics of data/image coding, compression, and encryption VI, with applications. Edited by Mark S. Schmalz. *Proceedings of the SPIE*, Volume 5208, pp. 88–98, January 2004
32. Z. Xiao-Ping, T. Li-Sheng and P. Ying-Ning. The design of a kind of chirp-like mother wavelet by neural network. *Proceedings of ICSP'96*, pp. 1381–1384
33. S. Qian and Q. Yang. *Graphical system and method for designing a mother wavelet*, United States Patent 6108609. [On-line] at: <http://www.patentstorm.us/patents/6108609.html>
34. E. Jones, P. Runkle, N. Dasgupta, L. Couchman and L. Carin. Genetic algorithm wavelet design for signal classification. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 23, No. 8, pp. 890–895, August 2001
35. R.E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME - Journal of Basic Engineering*, Vol. 82, pp. 35–45, 1960
36. D. Simon. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. John Wiley & Sons, Hoboken, NJ, 2006

37. H. Sorenson. *Kalman Filtering: Theory and Application*. IEEE Press, 1985
38. S.J. Julier and J.K. Uhlmann. A new extension of the Kalman filter to nonlinear systems. *The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Multi Sensor Fusion, Tracking and Resource Management II*, SPIE, Orlando, FL, 1997
39. P.J. Escamilla-Ambrosio and N. Mort. Multisensor data fusion architecture based on adaptive Kalman filters and fuzzy logic performance assessment. *Proceedings of the Fifth International Conference on Information Fusion*, pp. 1542–1549, 2002
40. J.A. Roecker and C.D. McGillem. Comparison of two-sensor tracking methods based on state vector fusion and measurement fusion. *IEEE Trans. Aerospace and Electronic Systems*, 24, 4, pp. 447–449, 1988
41. C. Brown, H. Durrant-Whyte, J. Leonard, B. Rao and B. Steer. Distributed data fusion using Kalman filtering: A robotics application. In M. A. Abidi and R. C. Gonzalez (eds), *Data Fusion in Robotics and Machine Intelligence*, pages 267–309. Academic Press, San Diego, CA, 1992
42. O.E. Drummond. Track fusion with feedback. *Proceedings of SPZE Conference on Signal and Data Processing of Small Targets*, 2759, pp. 342–360, April 1996
43. M.A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChe Journal*, 37, 2, pp. 233–243, February 1991
44. R.J. Bolton, D.J. Hand and A.R. Webb. Projection techniques for nonlinear principal component analysis. *Statistics and Computing*, 13, 3, pp. 267–276, August 2003
45. N. Lawrence. Probabilistic non-linear principal component analysis with Gaussian process latent variable models. *Journal of Machine Learning Research*, vol. 6, pp. 1783–1816, November 2005
46. S. Haykin. *Neural Networks - A Comprehensive Foundation, 2nd edition*. Prentice-Hall. Englewood Cliffs, NJ, 1998
47. F. Rosenblatt. *Principles of Neurodynamics*. Spartan Press, New York, 1962
48. J.L. Elman. Finding structure in time. *Cognitive Science*, 14, 179–211, 1990
49. R. Eckhorn, M.J. Reitboeck, M. Arndt and P. Dicke. Feature linking via synchronization among distributed assemblies: simulation of results from cat visual cortex. *Neural Computation*, 2, pp. 293–307, 1990
50. J.L. Johnson. Pulse-coded neural nets: translation, rotation, scale, distortion and intensity signal invariance for images. *Appl. Opt.*, 33, 26, pp. 6239–6253, 1994
51. M. Li, W. Cai and Z. Tan. A region-based multi-sensor image fusion scheme using pulse-coupled neural network. *Pattern Recognition Letters*, 27, 1948–1956, 2006
52. S. Theodoridis and K. Koutroumbas. *Pattern Recognition, 3rd Edition*. Academic Press, San Diego, CA, 2006
53. M. Matteucci. A Tutorial on Clustering Algorithms. See at: <http://home.dei.polimi.it/matteucc/Clustering/tutorial/html/index.html>
54. B.S. Everitt. *Cluster Analysis*. Edward Arnold, London, 1993
55. J.B. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, pp. 281–297, University of California Press, Berkeley, 1967
56. E. Fix and J. Hodges. Discriminatory analysis: Nonparametric discrimination: consistency properties. *Technical Report 4*, USAF School of Aviation Medicine, Project 21-49-004, 1951

57. B.V. Dasarathy, Ed. *Nearest Neighbor: Pattern Classification Techniques*. IEEE Computer Society Press, Piscataway, NJ, 1991
58. F. Leisch. A toolbox for K-centroids cluster analysis. *Computational Statistics and Data Analysis*, 51, 2, pp. 526–544, 2006
59. J.C. Dunn. A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters. *Journal of Cybernetics* 3, pp. 32–57, 1973
60. J.C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York, 1981
61. S.C. Johnson. Hierarchical clustering schemes. *Psychometrika*, 2, pp.241–254, 1967
62. V. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, New York, 1998
63. C.J.C. Burges. *A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery*, 2, pp. 121–167, Kluwer Academic, Boston, 1998
64. P.H. Chen, C.-J. Lin and B. Scholkopf. *A Tutorial on ν -Support Vector Machines*. See at <http://www.csie.ntu.edu.tw/~cjlin/papers/nusvmtutorial.pdf>
65. N.L.M.M. Pochet and J.A.K. Suykens. Support vector machines versus logistic regression: Improving prospective performance in clinical decision-making. *Ultrasound Obstet Gynecol*; 27: 607–608, 2006. See at <http://www.interscience.wiley.com>
66. T. Kohonen. Self-organizing maps. *Series in Information Sciences*, Vol. 30. Springer, Heidelberg, 3rd Extended Edition, 2001
67. A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Math. Statistics* 38, pp. 325–339, 1967
68. G. Shafer. *A Mathematical Theory of Evidence*. Princeton Univ. Press, Princeton, NJ, 1976
69. G. Shafer. *A Theory of Statistical Evidence in Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*. Harper and Hooker ed. Reidel, Dordrecht, 1976
70. A.O'Neill. *Dempster-Shafer Theory*. See at: <http://www.aonaware.com/binaries/dempster.pdf>
71. P. Smets. *What is Dempster-Shafer's model?* See at: <http://iridia.ulb.ac.be/psmets/WhatIsDS.pdf>
72. L.A. Zadeh. Fuzzy sets. *Information and Control*, 8, pp. 338–353, 1965
73. L.A. Zadeh. Fuzzy algorithms. *Information and Control*, 5, pp. 94–102, 1968
74. E.H. Mamdani. Applications of fuzzy logic to approximate reasoning using linguistic synthesis. *IEEE Trans. on Computers*, 26, 12, pp. 1182–1191, 1977
75. T. Takagi and M. Sugeno. Fuzzy identification of systems and its application to modelling and control. *IEEE Trans. on Systems, Man, and Cybernetics*, 15, 1, pp. 116–132, 1985
76. O. Postolache, L. Silva-Carvalho, G. Postolache and P. Silva Girão, I. Rocha. Portable instrument for autonomic nervous system analysis. *Proceedings IMEKO TC-4 14th Symposium on New Technologies in Measurement and Instrumentation and 10th Workshop on ADC Modelling and Testing*, Vol. I, pp. 312–317, Gdynia, Poland, September 2005
77. G. Postolache, L. Silva-Carvalho, I. Rocha, O. Postolache and P. Silva Girão. A wavelet-based method for estimation of the autonomic balance after experimentally drug administration. *Proceedings 2003 IEEE Canadian Conference on*

- Electrical and Computer Engineering (CCECE 2003)*, Vol. III, pp. 2083–2086, Montreal, Canada, May 2003
78. J. von Neumann, and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944 sec. ed. 1947
 79. J.M. Juran. *Juran's Quality Control Handbook*. McGraw-Hill, New York, 1988
 80. E.P. Blasch and M. Hensel. Fusion of distributions for radar clutter modeling. *Proceedings of the Seventh International Conference on Information Fusion*, pp. 629–636, Stockholm, July 2004
 81. E.L. Waltz and J. Llinas. *Multisensor Data Fusion*. Artech House, Norwood, MA, 1990
 82. I. Ben-Gal. Bayesian networks. In *Encyclopedia of Statistics in Quality and Reliability*. John Wiley & Sons, Hoboken, NJ, 2007. See at: <http://www.eng.tau.ac.il/~bengal/BN.pdf>
 83. D.D. Ullman. *Making Robust Decisions - Decision Management for Technical, Business & Service Teams*. Trafford Publishing, Victoria, BC, Canada, 2006
 84. D. Vose. *Risk Analysis: A Quantitative Guide*. John Wiley & Sons, Chichester, UK, 2000
 85. T. Bedford and R. Cooke. *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, Cambridge, UK, 2001
 86. J. Mun. *Modeling Risk: Applying Monte Carlo Simulation, Real Options Analysis, Forecasting, and Optimization Techniques*. John Wiley & Sons, Hoboken, NJ, 2006

Comparing Results of Chemical Measurements: Some Basic Questions from Practice

Paul De Bièvre

Independent Consultant
on Metrology in Chemistry (MiC), Duineneind 9, B-2460 Kasterlee, Belgium
paul.de.bievre@skynet.be

Summary. A number of basic questions are asked which arise from daily practice in the chemical measurement laboratory “when comparing results of chemical measurements.” Answers to these questions have either not (yet) been given so far, or not made properly known to the users concerned, or the answers are unsatisfactory, or they unveil a variety of opinions, if not controversy. A scientific debate and suitable propagation of its conclusions is needed.

Key words: Metrology in chemistry, VIM, interlaboratory comparison, normal distribution, reference quantity value, measurement uncertainty

1 Introduction

Results of chemical measurements are compared daily. In many practical cases, these comparisons are at the origin of a confrontation between parties such as buyer and seller, inspecting body and inspected person, measurement laboratory and regulatory body, or between declared value and actually measured value. In such cases, only two measurement results are involved. These comparisons usually involve different measurement methods, different measurement laboratories, different measurement instruments, or different analysts, but in the end, they are made between *two measurement results*. There also are cases where comparisons of *more than two* measurement results are made in order to arrive at statistical conclusions describing sets of large numbers of results. We attempt to formulate some questions which arise in both cases, and to which the answers do not seem to have been openly debated sufficiently.

Of utmost importance in discussing a comparison of measurement results, is, of course, a good basic and commonly accepted understanding of the concepts ‘comparison’ and ‘measurement result’ as well as various associated

concepts¹. In this chapter, we exclusively use the concepts as defined in the revised International Vocabulary of Metrology – VIM dGM 200:2008 [1], henceforth termed “VIM3”. The terms labelling these concepts are printed in bold in order to remind the reader that they are defined in VIM3 (sometimes we add the explanatory notes to the VIM3 definitions when they are important).

In daily life, many different kinds of comparisons are made between objects of a similar kind such as books on a given topic (comparison possible), and, sometimes and curiously enough, between objects of a different kind such as apples and oranges (comparison not possible). Here we examine a specific case of the concept ‘comparison’ namely ‘comparison of measurement results’. The former is a *superordinate concept*. The latter is a *subordinate concept* (a special case of the former) and therefore needs a special term to identify it. We term it *metrological comparison* in line with the thinking that underlies VIM3.

2 Why do we compare measurement results?

One of the most important purposes of *measurements*, including measurement in chemistry, is to be able to communicate. We want to exchange knowledge – including quantitative knowledge – about some *measurand*, that is, about some “quantity intended to be measured”, a 2008 definition [1] that constitutes a considerable modification of the previous 1993 definition “quantity subject to measurement” [2]. In order to communicate fruitfully in terms of quantitative knowledge, we need to compare these measurement results. We distinguish two categories of comparisons: comparisons of *two* measurement results, and comparisons of *more than two* measurement results.

2.1 Comparison of two measurement results in chemistry

Every day, two measurement results are compared. There can be very important reasons for such a comparison:

- The wish to know whether two different measurement results (one from a buyer and one from a seller) indicate a real difference or rather an equivalence in the quality of a specified property of a material (e.g., protein concentration in wheat) and therefore in the price (schematically given in Figure 1).
- The wish to know whether the measurement results (e.g., the toxic dioxin content in chicken meat) obtained by an exporter and an importer, and lying above and below an upper regulatory limit for the concentration of dioxin, constitutes (formally) a potential danger to human health in which

¹ Single quotation marks (‘...’) refer to concepts. Double quotation marks (“...”) enclose terms and quotations from other sources.




Fig. 1. Whether two measured quantity values are really different depends on their measurement uncertainties (which are probabilities that the results are located in the indicated interval).

case the trade of chicken meat can be stopped at a border, causing a potential trade dispute; if a regulatory upper limit for the concentration is not exceeded (measured quantity value is lower than the upper limit), no impediment of trade can be invoked; but, if it is (i.e., when the measured quantity value is higher than the upper limit), a permissible barrier to trade can be erected.

Looking at Figure 1, an international normative definition of measurement result is needed which stipulates whether the measurement uncertainty is part of the measurement result or not. As shown in Figure 1, the measurement uncertainties of the results determine whether the measurement results are significantly different. That is a major reason to have a common perception on whether the measurement uncertainty is part of the measurement result. In the VIM3 [1], we find a definition for common global use: “set of quantity values being attributed to a measurand together with any other available relevant information.” It, quite rightly, makes measurement uncertainty de facto part of the result and the explanatory notes to the definition are clear on that [1]. However, two measured quantity values are never identical except by coincidence. More specifically, two *different* measured quantity values for the same measurand in the same sample are the rule rather than the exception. Whether one result can be substituted for another, in other words whether

both results are equivalent, is then determined by the measurement uncertainties associated with the measured values, together with the difference between the two measured quantity values themselves.

Each time measurement uncertainty is used here, it is used exclusively as defined in [1]. Because of its key importance, the definition is quoted:

measurement uncertainty

uncertainty of measurement

uncertainty

non-negative parameter characterizing the dispersion of the **quantity values** being attributed to a **measurand**, based on the information used

NOTES

1 — Measurement uncertainty includes components arising from systematic effects, such as components associated with **corrections** and the assigned quantity values of **measurement standards**, as well as the **definitional uncertainty**. Sometimes estimated systematic effects are not corrected for but, instead, associated measurement uncertainty components are incorporated.

2 — The parameter may be, for example, a standard deviation called **standard measurement uncertainty** (or a specified multiple of it), or the half-width of an interval, having a stated **coverage probability**.

3 — Measurement uncertainty comprises, in general, many components. Some of these may be evaluated by **Type A evaluation of measurement uncertainty** from the statistical distribution of the quantity values from series of **measurements** and can be characterized by standard deviations. The other components, which may be evaluated by **Type B evaluation of measurement uncertainty**, can also be characterized by standard deviations, evaluated from probability density functions based on experience or other information.

4 — In general, for a given set of information, it is understood that the measurement uncertainty is associated with a stated quantity value attributed to the measurand. A modification of this value results in a modification of the associated uncertainty.

Measurement uncertainty must be reflected in a measurement uncertainty budget, whereby the uncertainties must be evaluated according to the ISO Guide for the Expression of Uncertainty in Measurement (GUM) [3].

What we are really interested in, is the

metrological compatibility of measurement results

metrological compatibility

property of a **set of measurement results** for a specified **measurand**, such that the absolute value of the difference of any pair of **measured quantity values** from two different **measurement results** is smaller than some chosen multiple of the **standard measurement uncertainty** of that difference

NOTES

1 — Metrological compatibility of measurement results replaces the traditional concept of “staying within the error”, as it represents the criterion for deciding whether two measurement results refer to the same measurand or not. If in a set of **measurements** of a measurand, thought to be constant, a measurement result is not compatible with the others, either the measurement was not correct (e.g. its

measurement uncertainty was assessed as being too small) or the measured **quantity** changed between measurements.

2 — Correlation between the measurements influences metrological compatibility of measurement results. If the measurements are completely uncorrelated, the standard measurement uncertainty of their difference is equal to the root mean square sum of their standard measurement uncertainties, while it is lower for positive covariance or higher for negative covariance. [1]

A special case of *metrological compatibility* can be characterized by the property that measurement results can be substituted for each other for a specified intended use. This is termed *metrological equivalence*.

metrological equivalence of measurement results

equivalence of measurement results

property of two or more **measurement results** for a given **measurand** that have **metrological compatibility**, and are each acceptable for the same specified intended use [4]

Metrological equivalence of measurement results can vary very much: from very good (two very small measurement uncertainties) to very bad (two very large measurement uncertainties, or one very small and one very large uncertainty). So far, these concepts have not been used very much in the literature although they would assist in this recurrent problem.

Problems of a small number of measurement results (and two is certainly a small number) have recently been discussed in the literature [5], [6].

We still need to address the term “comparing”. A fairly current practice in chemical measurement is to compare various types of measurement results. That is done in a variety of ways, not necessarily governed by clear rules. A recurring example is a certification campaign for the quantity value embodied in a certified reference material (CRM), when two measurement results from sometimes widely diverging origin or type are “compared” in order to arrive at a supposedly better “certified value”. The same usage can occur in a Proficiency Testing Scheme (PTS), when two measurement results are compared, then combined, to yield a reference quantity value which can serve as a reference to evaluate the measurement capability of the participants, supposedly based on the belief that the average of a (preferably) large number of measurement results yields in a value closer to the “true value” which is then considered to be the “best” reference. A normal distribution of the values is assumed either before or after the removal of “outliers”. In most of these cases, the attention goes to the measured quantity values, sometimes without too much consideration of their measurement uncertainty.

The change to a revised definition of measurement result could have consequences, both for field practice as for “high metrology”, entailing the question whether certain rules are not needed for comparing, then “combining” measurement results.

In a practical case, we are interested in, for example, two clinical measurement results for the same patient (one obtained in the home country




Fig. 2. IMEP-17: measurement results for cholesterol concentration in identical human serum samples (991 participating laboratories); certified metrological reference value: $5.111 \pm 0.021 \text{ mmol} \cdot \text{L}^{-1}$ with $U = k \cdot u_c$ and $k = 2$ [7].

and one in a tourist resort) and measured by two laboratories which have shown their measurement capability in an interlaboratory comparison (ILC) by delivering a result obtained on the same sample. One of them could lie in the upper part of group 1 and one from the other laboratory could lie in the lower part of group 2 in Figure 2. It seems important to be able to evaluate the *metrological compatibility of the measurement results*. We must conclude that a difference of as much as 20% to 30% relative may be normal and not indicative for a significant difference. Figure 2 is taken from the International Measurement Evaluation Programme, IMEP, from the Institute for Reference Materials and Measurements, IRMM, and can be consulted on the Web [7].

In addition, it is meaningful, or even necessary, to ask the question whether both results can indeed be compared, that is, are in compliance with the definition of *comparability of measurement results* which stipulates that such results must be traceable to a common reference. For that, it is essential to recall two key definitions given in [1]. One of them is the definition of

metrological comparability of measurement results

metrological comparability

comparability of **measurement results**, for quantities of a given kind, that are metrologically traceable to the same reference

EXAMPLE

Measurement results, for the distances between Earth and Moon, and between Paris and London, are metrologically comparable when they are both metrologically traceable to the same **measurement unit**, for instance the metre.

NOTES

- 1 — See Note 1 to the definition of **metrological traceability**.
 2 — Metrological comparability of measurement results does not necessitate that the **measured quantity values** and associated **measurement uncertainties** compared be of the same order of magnitude.

Because metrological traceability is essential – or even a prerequisite – to understand comparability, it is indicated to look at its definition [1]:

metrological traceability

property of a **measurement result** whereby the result can be related to a reference through a documented unbroken chain of **calibrations**, each contributing to the **measurement uncertainty**

NOTES

- 1 — For this definition, a ‘reference’ can be a definition of a **measurement unit** through its practical realization, or a **measurement procedure** including the measurement unit for a non-ordinal quantity, or a **measurement standard**.
 2 — Metrological traceability requires an established **calibration hierarchy**.
 3 — Specification of the reference must include the time at which this reference was used in establishing the calibration hierarchy, along with any other relevant metrological information about the reference, such as when the first calibration in the calibration hierarchy was performed.
 4 — For **measurements** with more than one **input quantity in the measurement model**, each of the input **quantity values** should itself be metrologically traceable and the calibration hierarchy involved may form a branched structure or a network. The effort involved in establishing metrological traceability for each input quantity value should be commensurate with its relative contribution to the **measurement result**.
 5 — Metrological traceability of a measurement result does not ensure that the measurement uncertainty is adequate for a given purpose or that there is an absence of mistakes.
 6 — A comparison between two measurement standards may be viewed as a calibration if the comparison is used to check and, if necessary, correct the quantity value and measurement uncertainty attributed to one of the measurement standards.
 7 — The ILAC considers the elements for confirming metrological traceability to be an unbroken **metrological traceability chain** to an **international measurement standard** or a **national measurement standard**, a documented measurement uncertainty, a documented measurement procedure, accredited technical competence, metrological traceability to the SI, and calibration intervals (see ILAC P-10:2002).
 8 — The abbreviated term “traceability” is sometimes used to mean ‘metrological traceability’ as well as other concepts, such as ‘sample traceability’ or ‘document traceability’ or ‘instrument traceability’ or ‘material traceability’, where the history (“trace”) of an item is meant. Therefore, the full term of “metrological traceability” is preferred if there is any risk of confusion.

This definition of metrological traceability with explanatory notes means that measurement results that are not traceable to a common reference *cannot*

be compared to each other by definition. The above definitions of basic concepts with associated terms, seem to be essential when formulating questions about “comparing” measurement results, the topic of this chapter.

We can now attempt to formulate and answer some questions.

Question 1

Must an organizer of an ILC, who is going to compare the participants' measurement results, not verify whether they are comparable in principle?

Answer. The answer is clearly positive: the definition of metrological comparability of measurement results requires metrological traceability of measurement results to a common reference regardless of the measurement procedure (which includes measuring systems and measurement method) used.

It is worthwhile to note that the possibility of metrological comparability of measurement results arises only *after* these results have been obtained: it is an operation becoming possible *a posteriori* to the process of measurement by the participants. However, its definition requires metrological traceability of these results (see definition above, Note 2) and that is established (i.e. decided) *prior* to each individual measurement by each participant in the planning stage of his or her measurement, notably when he or she decides on the calibrator(s). That leads to

Question 2

Must the organizer of an ILC not request metrological traceability from each participant for his/her result when he wants to create the possibility for any two laboratories to compare their measurement results and draw conclusions about that pair?

Answer. There is little doubt that the answer is positive because comparability of results is only possible for metrologically traceable results as the definition explicitly requires.

Thus in Figure 3, where measurement results are displayed in four groups and where the measurement results in each group had their metrological traceability only established to a ‘common reference’ for the members of that group [7], a measurement result of one group is *not comparable* to the measurement result of another group and no meaningful conclusion about such a pair of measurement results can be drawn.

More experimental evidence is given in Figure 4 (nine different measurement procedures, each with its own possibly commercial measurement standard as reference). Hence, one measurement result of one group is uncomparable with any measurement result of any other group.

There seems to be another question arising from the figures.




Fig. 3. IMEP-17: measurement results of 950 laboratories for γ -glutamyl-transferase, arranged in method groups (for technical details of the methods, see [7]). Certified [metrologically traceable] reference value: $34.70 \pm 0.93 \text{ U}\cdot\text{L}^{-1}$ with $U = k \cdot u_c$ and $k = 2$ (U is the symbol for international WHO unit) [7].




Fig. 4. IMEP-17: measurement results of 863 laboratories for amylase, arranged in method groups (for technical details of the methods, see [7]). Certified [metrologically traceable] reference value: $56.8 \pm 2.6 \text{ U}\cdot\text{L}^{-1}$ with $U = k \cdot u_c$ and $k = 2$ (U is the symbol for international WHO unit) [7].

Question 3

In view of Figure 1, doesn't an ILC organizer need to require a statement of measurement uncertainty (which automatically means measurement uncertainty according to VIM and GUM) for each individual measurement result declared by any of the participants in order to enable any two participating laboratories to compare their measurement results and draw conclusions about that pair?

Answer. Again, the answer seems to be in the affirmative (see Figure 1) because that is required by the definition of measurement result which makes measurement uncertainty part of the result. That leads to the next question:

Question 4

Is the establishment of metrological traceability not a prerequisite for evaluating measurement uncertainty of each of any two measurement results?

Answer. The answer is positive. It is often claimed in the literature that an uncertainty statement (frequently delivered in many different forms which are not GUM-compliant) automatically entails metrological traceability. That is an erroneous conclusion. The metrological traceability is determined by the choice of a calibration hierarchy [1], a decision which is part of the measurement plan made up *before* the measurement starts. That also fixes the metrological traceability chain which can then be used *after* the measurement to demonstrate metrological traceability (very much as a roadmap has to be decided *before* a journey and be used to guide the journey *after* the roadmap is decided). Measurement uncertainties are attached to the various links in that chain (i.e., each time that a calibration in such link is carried out). That in turn will enable us to evaluate measurement uncertainty of the end user's measurement result as the accumulation of all these uncertainties. See the definition of metrological traceability above.

[However, the end user him/herself only has to worry about the calibration of his/her measuring system and of his/her ensuing measurement. The metrological traceability chain with associated measurement uncertainty of the calibrator used (usually in the form of a certified reference material, CRM), must be provided by the CRM producer/seller with a statement of the metrological reference of the metrological traceability and of the measurement uncertainty of the value embodied in the CRM.]

There is another reason why the answer to Question 4 is positive. The so-called uncertainty statement made by most analysts is mostly not compliant with the definition of measurement uncertainty in VIM and GUM. It is often only a repeatability, or a reproducibility, or it is derived from another very incomplete evaluation of the uncertainty budget.

Question 5

Can we combine two measurement results which are not metrologically comparable and build conclusions on such comparisons, sometimes by combining these results?

Answer. The answer seems to be negative. It is difficult to see why measurement results which are uncomparable by definition, could be compared, then combined. Yet, that is done in practice many times without too much consideration for whether or how metrological traceability was established and, therefore, whether these results were comparable at all.

So far, we have asked questions for cases where only two measurement results are implicated. We now identify a few questions for cases where more than two measurement results are involved.

2.2 Comparison of many measurement results in chemistry

A number ($N > 2$) of measurement results in chemistry are often compared to obtain a measure of the spread of the results amongst a group of laboratories. Or they are pooled in order to obtain some kind of “most probable value”, which is considered as a “best value” and which could be considered for use as a reference quantity value for the measurand under investigation.

Question 6

Must each participant of an ILC not establish him/herself the metrological traceability of his/her own measurement results before measuring the measurand of concern in an ILC, in order to create the possibility that any result in the ILC can be compared to any other result?

(similar to *Question 2* for two measurement results)

Answer. In the light of the above definition of metrological traceability as a prerequisite for metrological comparability of measurement results, the answer is, no doubt, in the affirmative. Metrological traceability of participants' results is established *before* the measurement as it is part of the measurement plan and as the choice (i.e. a decision) of the calibrator to be used, must, of necessity, be made at that stage. Traceability of any measurement result is also required by ISO 17025.

In the study of such sets of measurement results, a normal distribution of the results is “assumed” and some sort of average, mean, weighted mean, median, and so on is calculated after (or without) the removal of “outliers”. That leads to the following

Question 7

Can measurement results in chemistry in an ILC automatically be expected to be normally distributed around a reference quantity value?

Answer. Frequently not, as can be observed in Figures 5 through 7 as well as in other IMEP graphs [7]. It seems clear that normal distribution around a reference quantity value cannot automatically be assumed. An obvious conclusion would be that, in any given case, *a normal distribution must be proven before it can be assumed.*




Fig. 5. Measurement results for amylase in $\text{U}\cdot\text{L}^{-1}$; certified [metrologically traceable] reference value: $56.8 \pm 2.6 \text{ U}\cdot\text{L}^{-1}$ with $U = k \cdot u_c$ and $k = 2$ (U is the symbol for international WHO unit) [7].

Thus follows naturally another question:

Question 8

Can general conclusions which are based on the assumption of a normal distribution around a reference quantity value be drawn when experimental evidence does not substantiate such a distribution?

Answer. The answer to this question must be negative. However, the chemical literature is awash with statements such as “assuming normal distribution of measurement results” around a reference quantity value, without that assumption being proven. In a logical reasoning, any conclusion must be based on premises that have been checked for their validity first. Well-documented cases in the literature (see two examples in Figures 6 and 7) do not exactly display a normal distribution in practice. They prevent conclusions to be drawn that are based on such an assumption.




Fig. 6. Measurement results in IMEP-20 for mass fraction of As in tuna fish according to the criterion: CRM was used (“yes”) or no CRM was used (“no”); certified [metrologically traceable] reference value: 4.93 ± 0.21 mg(As)/kg(material) with $U = k \cdot u_c$ and $k = 2$ [7].




Fig. 7. Measurement results in IMEP-13 for amount-of-substance content of Pb in polyethylene grouped according to the criterion ‘accredited laboratories’ or ‘non-accredited laboratories’ or ‘no statement by participant’; certified [metrologically traceable] reference interval: 0.501–0.539 mmol (Pb)·kg⁻¹(polyethylene) with $U = k \cdot u_c$ and $k = 2$ [7].

Question 9

Do the different measurement results in an ILC belong to a same population?

Answer. Measurement results in an ILC are usually obtained from different measurement laboratories, methods, analysts, and instruments. Most of these characteristics prevent the conclusion “belonging to the same population”. One of the main conditions for being normally distributed is therefore not fulfilled, and any conclusion based on that evidence must therefore remain elusive.

Question 10

If the average, or median, or ‘cleaned’ mean of the participants’ results derived from an ILC, where metrologically traceable and nonmetrologically traceable measurement results are sometimes mixed (“combined”), can that lead to a “reference value”?

Answer. It is very difficult to see how that could be a permissible procedure because the very concept of metrological comparability of measurement results requires metrological traceability of measurement results to a common reference. The presence of nonmetrologically traceable measurement results prevents an immediately positive answer to this question.

One of the subsequent uses of a reference quantity value (average, median, mean, etc.) in a calibrator (CRM) is the calibration of measurement procedures by means of that calibrator. Therefore it is legitimate to ask the following

Question 11

To which extent is a subsequent calibration of measurement procedures, using such a reference quantity value, a circular process as the same measurement procedures were used in the establishment of the reference quantity value itself?

Answer. One cannot escape from the impression that such an approach has at least the presumption of being circular, thus nullifying its scientific value. It is also difficult to see how values which are not normally distributed around a metrologically traceable quantity value and are averaged, could be considered an approach towards an independent measured quantity value equivalent to a metrologically traceable quantity value, as Figures 3 and 8 show clearly. A quantity value calculated from measured quantity values has usually been subjected to various considerations and statistical treatments and hence seems to be dependent on the choice of such treatments (giving the impression of being more a matter of decision rather than of measurement).

Question 12

Who is then guarding orthodoxy in this process?




Fig. 8. Measurement results for mass fraction of Pb in polyethylene; certified (metrologically traceable) reference interval: 0.449–0.491 mmol(Pb)/kg(material) with $U = k \cdot u_c$ and $k = 2$ [7]. (1) and (2) measurement uncertainty encompassing the metrologically traceable reference quantity value with its associated expanded measurement uncertainty ($k = 2$); (3) target measurement uncertainty: measurement uncertainty specified as a goal, in this case put at $\pm 15\%$, and decided a priori on the basis of the intended use of the measurement results.

Answer. This question must probably not be answered by identifying a formal “authority”, but rather by giving an answer based on scientific/metrological grounds in order to prevent any arbitrariness. The answer can come from the concept of metrologically traceable value which shows by its very metrological traceability chain, that it constitutes a connection to our internationally established structure of measurement, the International System of Units/Système International d’Unités, (SI), operating under the formal responsibility of CCQM/CIPM/BIPM and CGPM according to an Intergovernmental Treaty, the Convention du Mètre. Put in simple language: a metrologically traceable value takes automatic precedence over any other value whatever the non-metrological procedure being the basis of its assignment.

The answers to some of the questions above lead to the following:

Question 13

Does the very fact of participating in an ILC not have as logical consequence that the participant who offered a nonmetrologically traceable result a priori cannot use the calculated average (or median or mean) as a–(nother) metrological reference a posteriori?

Answer. It seems so, as the process would look as a self-fulfilling prophecy. If a measurement laboratory contributes to setting a reference without having a metrological reference itself, it seems that the very notion of reference is turned around.

However, it could be argued that the measurement results and the reference quantity value obtained from an ILC, where participants all use the same material as well as a metrologically traceable calibrator value, would not only be comparable but also compatible (or equivalent) amongst themselves. That would/could have the ironic result that an ILC was not necessary to obtain a new metrologically traceable quantity value because any value of any participant is, metrologically speaking, a metrologically traceable quantity value confirming all the others and is being confirmed by all the others (!).

However, it is very important to point out that each of the participants' results may inspire a very increased confidence and trust to a user of the reference value because that value is strengthened very much by these mutual confirmations.

That leads to the next question.

Question 14

Isn't the basic product of an ILC to establish some sort of "degree of equivalence" of any specified pair of results of participants, rather than comparing and possibly combining measurement results?

Answer. It seems so, thus fulfilling an extremely useful role in confrontations as described above in the introduction.

Question 15

Is a 'consensus value' for the certified quantity value(s) carried by a calibrator (CRM) (average, median, mean of selected values) metrologically traceable?

Answer. Usually different methods and different instruments are used by different laboratories and performed by different people in order to have independent (or at least as independent as possible) results enabling us to combine them and to arrive to a suspected "true value" as closely as possible "assuming a normal distribution". However, under these conditions, it is hard to see that such measurement results belong to the same population, condition for the assumption of a normal distribution.

Establishing a reference quantity value for an ILC *after* the measurement (because calculated from the ILC results) seems to violate the (chrono-)logics in the definition of metrological traceability that it be established *before* the measurement.

Question 16

Is not measurement uncertainty sometimes (instinctively) made larger to encompass a perceived uncertainty in the statement of a metrological traceability?

Answer. That operation seems to be acceptable according to the GUM which permits a Type B evaluation based on skill and judgement of the analyst. It can be carried out by *deciding* on a coverage factor (which must then be mentioned).

Question 17

Is it possible to state or establish a (nonmetrological) form of traceability for a ‘consensus value’ of an ILC?

Answer. Probably. Maybe that could just be called “traceable to a given and described decision”. However, some careful thinking is needed here. This form of “traceability” in fact suggests a sort of new unit: it is not traceable to a metrological reference but to a decision. That would prohibit the use of the expression “traceable to an SI unit.” The impression one gets is that a new, arbitrary unit is created, giving the impression of traceability to the definition of an SI unit, whereas in reality it only carries the *name* of an SI unit (i.e., the result is just stated in an SI unit). The metrological traceability chain seems to be broken at the moment of the decision about the choice of the statistical treatment of measurement results belonging to different populations. The same unit cannot be claimed anymore above the break in a metrological traceability chain because the definition of metrological traceability prescribes metrological traceability all the way to, for example, the definition and realization (or embodiment) of the unit chosen (SI or other) before the start of the measurement. That is especially important if one wants to claim ‘metrological traceability to the SI’. It looks like this traceability is claimed by just expressing a measurement result in a given measurement unit. But that is not establishing an unbroken metrological traceability chain to a realization (embodiment) of that unit. It is just a “declaration” of a metrological traceability chain to a(n SI) unit, implied by writing the result in terms of this unit.

However, there is no doubt that such a mean, average, or selected mean, after a ‘cleaning up’ process by a competent assembly of experts, increases the confidence and trust the end user will have in the value. ‘Consensus values’ contribute to a huge increase of trust and confidence building; that is, they confirm trustworthiness of metrologically traceable values embodied in CRMs. If IRMM and NIST have obtained mutually confirming values on certified quantity values carried by a material, or if IRMM, LGC, and BAM do that as they sometimes do in their cooperation towards realizing European Reference Materials (ERM), that process mutually confirms each other’s certified values, giving them additional trust and confidence. One could even state that it confirms trust and confidence in the metrological traceability of each of the three values (or of the value of the Institute which would take full responsibility for the CRM), because the other two confirm by a strict metrological traceability, the value of the third measurement.

That, of course, leads to an interesting question formulated as follows:

Question 18

Is the measurement uncertainty claimed for a “certified measurement result” which is generated by the combination of several measurement results with full (GUM) measurement uncertainties, obtained by for example, a few National Metrology Institutes results, smaller or larger than the individual measurement uncertainty of either of the two individual measurement results?

Answer. The answer is not as straightforward as usually implicitly or explicitly thought. First, only a low number of values are available which makes statistical conclusions very dangerous. Second, presumed carefully drafted measurement uncertainty budgets result in a full (GUM) evaluated measurement uncertainty of the certified measurement result in each of the cooperating institutes. How could that measurement uncertainty be reduced just by combining it with another measurement uncertainty, presumed to be equally carefully drafted and resulting in an equally carefully full (GUM) evaluated measurement uncertainty of the certified measurement result? Would one not naturally and intuitively expect that the measurement uncertainty of the “combined” result be larger rather than smaller, in order to cover all likely values?

3 Final remarks

We did not intend – on purpose – in this chapter of this book to indicate how to compare measurement results and provide statistical procedures for it. We rather concentrated on the description of *practical cases* where comparisons of measurement results are performed for an important intended use. We did so by means of well-documented and available “comparisons” of measurement results. Both for the cases where two measurement results or a large number ($N > 2$) of measurement results are compared, a number of questions have been identified which need to be openly debated and which need answers that are professionally formulated and supported by convincing transparent arguments. In the first category, the answers must be applicable to the case of “comparing” basically two measurement results. In the second category, the answers must not automatically “assume” normal distributions. It may be necessary that new – and simple – forms of treatments of measurement results have to be drafted for practical use, to replace present, sometimes very complicated calculations, either not applicable to a very small numbers of results, or based on a possible fallacy of an unproven normal distribution.

References

- [1] BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP, OIML: The International Vocabulary of Metrology, Basic and General Concepts and Associated Terms. VIM, IdGM 200:2008, BIPM (2008); www.bipm.org

- [2] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML: International Vocabulary of Basic and General Terms in Metrology. ISO, Geneva (1993/1995)
- [3] BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML: Guide for the Expression of Uncertainty in Measurement, ISO/IEC Guide 98. ISO, Geneva (1993) and Supplement 1. (2007)
- [4] IUPAC Recommendations on Metrological Traceability of Measurement Results in Chemistry, section 6. To be published (see provisional draft on <http://www.iupac.org/reports/provisional/abstract07/fajgelj\290208.html>)
- [5] Hibbert, D.B.: Compliance of analytical results with regulatory or specification limits: a probabilistic approach. *Accred. Qual. Assur.*, **6**, 346–351 (2001)
- [6] Belli, M., Ellison, S.L.R., Fajgelj, A., Kuselman, I., Sansone, U., Wegscheider, W.: Implementation of proficiency testing schemes for a limited number of participants. *Accred. Qual. Assur.*, **12**, 391–398 (2007)
- [7] <http://www.imep.ws>

Modelling of Measurements, System Theory, and Uncertainty Evaluation

Klaus-D. Sommer

Physikalisch-Technische
Bundesanstalt, Bundesallee 100, Braunschweig, Germany, D-38116
klaus-dieter.sommer@ptb.de

Summary. Modelling of measuring systems or processes is an indispensable prerequisite of modern uncertainty evaluation. Per se, any modelling remains imperfect; it is always reduced to relevant influences, system parameters, and behaviour. Therefore, in uncertainty evaluation, it is important to consider the effects of imperfect modelling. Derived from the classical theory of signals and systems, this contribution explains the basic approaches to systematic modelling of continuous measuring systems. The approach used is based on the cause–effect analysis and a subsequent design of the model of the measurement by means of standardised components. Emphasis is put on modelling of nonlinear and also time-variant systems and, consequently, on the effects of disregarding nonlinearities and time-variant (dynamic) behaviour. The chapter demonstrates that, based on carefully analysing and understanding the measurement, it is possible to reach proper modelling of measurements and one may sufficiently describe the relevant effects of imperfect modelling.

Key words: System modelling, effect of modelling

1 Introduction to the modelling of measurements

1.1 Modelling tasks in uncertainty evaluation

For evaluating the measurement uncertainty it is necessary to mathematically describe both the measurement process as well as all influencing quantities and parameters of the measuring system [1, 2]. Henceforth, these quantities and parameters are termed input quantities [1, 2].

In accordance with the Bayesian probability theory and the so-called principle of maximum information entropy (pme) [2–4], the unavoidably incomplete knowledge about the input quantities is expressed by appropriate state-of-knowledge probability density functions (pdf). This commonly accepted theory, for example, yields

- A rectangular pdf if one knows that the possible values ξ_i of the quantity X_i are contained in an interval, and

- A Gaussian pdf if one knows the standard deviation σ_j or its sample estimate s_j or the uncertainty u_{x_j} associated with the expectation of the quantity X_j .

Bayesian probability theory and the pme also apply to the modelling of given measurement data that, for example, are obtained from repeated observations without knowing much about the nature of their source. In combination with analytical sampling statistics, they provide the basic justification for making assumptions on appropriate types of distribution, such as, for example, of normally, Student- t , or, in case of multivariate output quantities, χ^2 distributed datasets.

Besides these quantity and data models, evaluation of a measurement essentially needs to know the interrelation between the input quantities, data, and the measurand(s) effectuated by the measurement process.

It is the primary target of the evaluation of a measurement, based on the knowledge of both the input quantities and the (inverse) measurement process, to determine the best estimate y for the measurand Y , together with the measurement uncertainty u_y . This uncertainty, in essence, is a measure of the ‘trustworthiness’ of the ‘estimated’ result value y .

In modern uncertainty evaluation, knowledge about the measurement process is represented by the so-called model (Equation (1)) [1, 2, 5]. It mathematically expresses the interrelation between the measurand(s) \mathbf{Y} and the relevant(!) input quantities $\mathbf{X} = (X_1, \dots, X_N)^T$:

$$\mathbf{Y} = f_M(\mathbf{X}). \quad (1)$$

Usually in classical metrology one deals with one (univariate) measurand only. In uncertainty evaluation, the final result of a measurement is either expressed by a pdf for the output quantity, which theoretically can be calculated using the so-called *Markov*-formula,

$$g_Y(\eta) = \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} g_{X_1, \dots, X_N}(\xi_1, \dots, \xi_N) \cdot \delta(\eta - f_M(\xi_1, \dots, \xi_N)) d\xi_1 \cdots d\xi_N \quad (2)$$

or by stating the best estimate $y = \int_{-\infty}^{+\infty} \eta g_Y(\eta) d\eta = E[Y]$ for the output quantity together with an associated measurement uncertainty $u_y = \sqrt{\text{Var}[Y]}$.

In both cases, one needs to know the model Equation (1) to calculate the measurement result [1, 2, 5]. Modern uncertainty evaluation utilises numerical methods to calculate a pdf for the output quantity [2].

1.2 Basic categories of models

As everywhere in modelling, any model for the description of the measurement can hardly be established completely and perfectly. In practice, models

should pragmatically reflect the ‘purpose-relevant’ behaviour of the measurement to be evaluated. Nevertheless, modelling of a measurement process appears to be the most difficult problem in uncertainty evaluation. It requires a sufficient understanding of the measurement including all influences and relationships.

Basically, there are three means available to set up and express models: language, graphics, and mathematics. Accordingly, one may distinguish the following categories of models.

- Verbal models:
Verbal models describe the facts of the case by means of the language. The language seems to be the most flexible tool to describe a model. It is less suitable in structuring and quantifying systems and processes.
- Graphical models:
The category of graphical models is, for example, comprised of signal-flow charts, block diagrams, structure diagrams, state graphs, and so on. They are suitable for comprehensible depiction of interconnections and flows. They are less suitable to represent quantitative relationships.
- Mathematical models:
Mathematical models may be structured into data models, such as data sequences and frequency distributions, analytical models (see Section 2), and connective models, such as neuronal networks.

1.3 Models for representing measurements and for evaluating the measurement uncertainty

Generally, a model serves to analyse the original system or to draw conclusions from its behaviour. In measurement, usually the measurand and other (system-disturbing) influence quantities may be seen as causative signals which by the measuring system are (physically) transformed into effects, such as indications or output signals. In this way, the measuring system assigns values to the measurand(s), and the system is influenced by system-disturbing influence quantities.

The cause–effect approach is the most commonly used and comprehensible methodology for the representation of basic relationships in modelling of measurements. It is explained in detail in [5, 6]. It is based on the constitution of the path of the measurement signal from cause to effect. Example 1 illustrates this approach.

Example 1: Temperature measurement of a liquid bath

The temperature measurement example is illustrated with Figure 1 [5, 6]: The indication $X_1 = t_{\text{IND}}$ depends on both the bath temperature t_{B} , (i.e., measurand) and the instrumental error Δt_{Th} of the thermometer used [5, 6]. Furthermore, the indication is affected by the temperature inhomogeneity of the bath, the effect of which is represented by the deviation δt_{B} . Then, in




Fig. 1. (a) Simplified example of a bath-temperature measurement. (b) Graph of the respective cause-and-effect relationship. (c) Cause-and-effect relationship expressed in mathematical terms. Symbols: t_B , (average) bath temperature (measurand); δt_B , deviation due to the temperature inhomogeneity of the bath; Δt_{Th} , indication(scale) error of the thermometer used; t_{IND} , indicated value [5].

accordance with the symbols used in Figure 1, the cause–effect relationship for this example can mathematically be expressed by $X_1 = h(Y, X_2, X_3)$.

A model that describes the cause–effect behaviour of a measuring system or sensor is often termed a ‘measurement equation’ or ‘sensor equation’.

In contrast to this, for evaluating the measurement uncertainty, usually an ‘inverse model’ is needed that establishes the relationship between the measurand (often termed output quantity) and all relevant input quantities, such as, for the above example (see Figure 1), $Y = f_M(X_1, X_2, X_3)$. Usually, this model category is termed the ‘model equation’ [1] or ‘model for the evaluation of measurement uncertainty’ (see Figure 2) and is generally expressed by Equation (1).

Figure 2 illustrates the difference between the two model categories. The knowledge about all involved quantities and parameters Y and X are represented by means of pdfs (see Section 1.1) the expectation values of which are the best estimates of the quantities. Consequently, these best estimates represent the operating point of the measuring system that may be expressed by the operating-point vector $x = (x_1, \dots, x_N)^T$. In accordance with the ISO-GUM [1], the unavoidable incompleteness of knowledge about the (true) values of the quantities X_1, \dots, X_N is expressed by means of the uncertainties associated with the above best estimates. These uncertainties are the standard deviations of the above-mentioned state-of-knowledge pdfs for the input quantities (see Section 1.1).

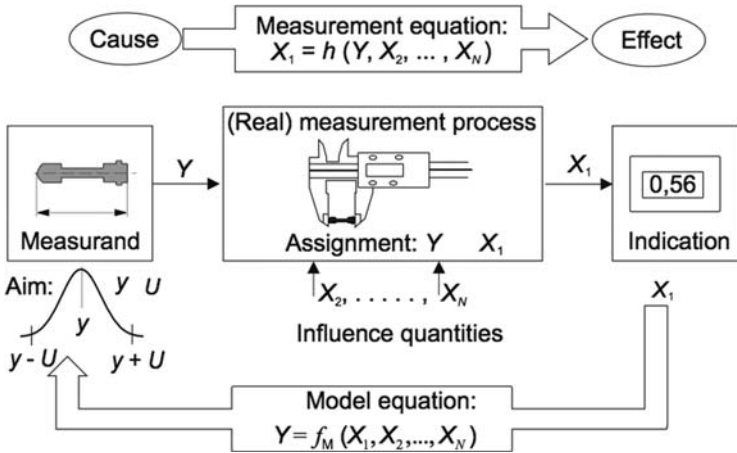


Fig. 2. Comparison of model categories: ‘Measurement equation’ and ‘Model equation for the uncertainty evaluation’ [5].

2 Describing systems and signals

Measuring systems are usually modelled the same way as any other technical information system. First, the system is decomposed and modularised into functional elements. Then, the transmission behaviour of any functional element is mathematically described, usually by means of its transfer behaviour. Equation (2) and Example 2 illustrate the meaning of the transmission behaviour and the transfer function which relate the output signal to the input signal(s):

$$X_{OUT} = h(\mathbf{X}_{IN}), \tag{3}$$

where $\mathbf{X}_{IN} = (X_{IN1}, \dots, X_{INn})^T$, input signal(s); \mathbf{X}_{OUT} , output signal; h , transfer function.

Figure 3 shows both the general graphic depiction of a transmission element and its application to the particular case of Example 2.

Example 2: Air buoyancy correction of a mass

A solid body having the mass m_0 and the density ρ_B is exposed to ambient with the air density ρ_A . The transfer function reads $W = h(m_0, \rho_A, \rho_B) = m_0(1 - \rho_A \cdot \rho_B^{-1})$, where W is the air buoyancy effect in terms of mass (output quantity/signal); $\mathbf{X}_{IN} = (m_0, \rho_A, \rho_B)^T$ is the (vector of the) input quantities/signals.

Referring to the appropriate mathematical description of the transfer behaviour, measuring systems are usually divided into time-invariant (including steady-state/static) systems, time-variant (dynamic) systems, and




Fig. 3. General transmission element: (a): General depiction. (b): Example: air buoyancy correction. Symbols: h , transfer function; $(X_{IN1}, \dots, X_{INn})^T = \mathbf{X}_{IN}$, input signal(s); \mathbf{X}_{OUT} , output signal(s); W , air buoyancy correction in terms of mass; ρ_A , air density; ρ_B , density of the body; m_0 , uncorrected (true) mass.

parametric-distributed systems. Table 1 gives a survey on static dynamic systems along with the corresponding mathematical description tools.

2.1 Linear time-invariant systems

In measurement, the great majority of systems is treated as being linear and time invariant. Therefore, proper description of this system category is of great importance in metrology and industrial measurement. Moreover, in analytical metrology, it is current best practice to treat slightly nonlinear and time-variant systems this way, but accounting for additional uncertainty contributions mainly owing to nonlinearity and dynamic effects.

In signal and system theory, a system is called linear and time invariant (lti system) if it meets the following properties and conditions.

- Superposition: $h(X_1 + X_2) = h(X_1) + h(X_2)$
- Homogeneity: $h(a \cdot X) = a \cdot h(X)$
- Time invariance: $X_{OUT}(t) = S\{X_{IN}(t)\} \Rightarrow X_{OUT}(t-t_0) = S\{X_{IN}(t-t_0)\}$

where X_1, X_2, X_{IN} are quantities acting at an (physical) input of the system or an element of the system; X_{OUT} is the quantity at the (physical) output of the system or an element of the system; t is time; and S is the transfer function.

A system or functional element which is only comprised of the mathematical operations addition, subtraction, multiplication by constant factors, and/or differentiation is always a linear and time-invariant system. The transfer function of a time-invariant system can be represented by an algebraic equation (see Table 1).

For a linear system, the transfer function consists of constant transmission factors, $h(\mathbf{X}_{IN}) = \mathbf{A} = (A_1, \dots, A_m)^T$. This is illustrated in Figure 4 for both the general case and an example.

Table 1. Survey on system classes and tools for their mathematical description.

System Class		Equation for General Systems	Equation for Linear Systems
Static Systems	Algebraic equation	Equation without time dependencies $Y = f(X)$	Linear system of equations $Y = AX$
	Dynamic Systems	Linear differential equations	
Description in time domain State space model		$Y(t) = f(Y^{(1)}, Y^{(2)}, \dots, X, X^{(1)}, Y^{(2)}, \dots, t)$	$\sum_{v=1}^n a_v Y^{(v)} = \sum_{\mu=1}^m b_\mu X^{(\mu)}$ $\dot{Z}(t) = AZ(t) + BX(t)$ $Y(t) = CZ(t) + DX(t)$
Description in frequency domain	Transfer function	$Y(s) = G(s)X(s)$ $G(s) = \frac{\sum_{\mu=0}^m b_\mu s^\mu}{\sum_{v=0}^n a_v s^v}$	




Fig. 4. Linear transmission element. (a) General depiction. (b) Example: $X_{OUT} = kX_1 + X_2$; $A = (k1)$. Symbols: X_1, X_2, X_{IN} , input signals; other symbols see Figure 3 and text.

2.2 Linearity of measuring systems

For uncertainty propagation, the ISO-GUM [1] uses the Gaussian law of error propagation which only can deal with linearisable models. Nevertheless, it should be mentioned that the ISO-GUM also considers second-order approximations. Linearisation is implemented by first-order Taylor-series expansion

of the model equation about the best estimates x_1, \dots, x_N of the input quantities X_1, \dots, X_N (see Equation (1)). Any model is considered to be linearisable if it fulfils the following condition,

$$\eta = f_M(\mathbf{m}\xi) = f_M(\mathbf{x}) + \sum_{i=1}^N c_i(\xi_i - x_i), \quad (4)$$

where η are the possible values of the output quantity Y (i.e. the measurand); $\mathbf{m}\xi = (\xi_1, \dots, \xi_N)^T$ are the possible values of the input quantities $\mathbf{X} = (X_1, \dots, X_N)^T$; $\mathbf{x} = (x_1, \dots, x_N)^T$; $x_i = E[X_i]$ are the expectation and best estimate of the quantity X_i . The c_i are called sensitivity coefficients [1] and given by

$$c_i = \left. \frac{\partial f_M(x_1, \dots, x_N)}{\partial X_i} \right|_{x_i = \xi_i}. \quad (5)$$

For a linear system or model, the best estimate of the measurand becomes [1]

$$y = E[Y] = f_M(x_1, \dots, x_N). \quad (6)$$

In practice, a measuring system is hardly ever described and analysed in a holistic way. Usually, a measuring system or process is described and analysed from cause to effect and considered as a series of nonreactive functional elements (or a sequence of operational steps) to carry out measurements. In practical measurement, the above linearity condition (see Equation (3)) will usually be satisfied if the following assumption for the individual functional elements can be made [5–7].

- In narrow ranges around the operating points of the input variables, the functional elements or steps of a measurement may be regarded as having approximately linear characteristics and may, therefore, be described by a first-order Taylor-series expansion. Figure 5 (which is taken from [5]) illustrates this linearisation and the limits of its applicability.
- The (steady-state) transmission behaviour of the fictitious unperturbed measuring system is related to well-known operating points given by the operating point vector $\mathbf{x} = (x_1, \dots, x_N)^T$ representing the expectation values of the (pdfs for the) input quantities.

For practical modelling of measurements, on the above assumptions, the ‘real world of measurement’ may be taken into consideration by allowing for slight deviations of the real influence quantities and other parameters from the above ‘idealised operating conditions’. The unavoidably incomplete knowledge about these deviations, and therefore about the quantities X_1, \dots, X_N , may be represented by appropriate pdfs and uncertainties, respectively (see Section 3).

2.3 Nonlinear systems

The polynomial model is a natural way to represent nonlinear transfer behaviour of measuring systems or elements of a system:




Fig. 5. Linearisation of the functional dependence of the output quantity Y on the input quantity X . For linearisation one replaces the model function $Y = f_M(X)$ by a relation for its values $\eta = f_M(x) + c(\xi - x)$, where c is the first partial derivative of the model function at its expectation values x . ξ and η are the possible values of X and Y . For a linear model the following hold: $y = f_M(x)$, expectation value of the output quantity; $u_y = c \cdot u_x$, standard uncertainty associated with y ; and c , sensitivity coefficient. Case (a): linear treatment is possible. Case (b): linearisation would lead to erroneous uncertainty propagation and here even underestimate the value for y [5].

$$X_{OUT}(t) = A_1 X_{IN}(t) + \dots + A_n X_{IN}^n(t). \tag{7}$$

In case of lti systems, in practical physical measurement one usually deals with slight nonlinearities only, and in practice, the real characteristics are replaced with idealised characteristics. The idealised linear characteristic will, in a given measuring range (see Figure 6), approximate the real characteristic without causing significant nonlinearity errors.

2.4 Time-variant signals and systems

A system is called dynamic or time variant if its response depends not only on the instantaneous value(s) of the (physical) input quantities/signals but also on elapsed values [8]. Correspondingly, the response of a nondynamic system depends on the instantaneous value(s) of the input signal(s) only.




Fig. 6. Nonlinear measurement characteristic (example). Symbols: $Y = X_{IN}$, (physical) input quantity/signal; $X = X_{OUT}$, (physical) output quantity/signal; x_L, y_L , beginning of the measuring range; x_H, y_H , end of the measuring range.

Generally, in the time domain, the system behaviour may be described by means of ordinary differential equations,

$$X_{OUT} = h \left(X_{OUT}^{(1)}(t), X_{OUT}^{(2)}(t), \dots, X_{IN}^{(1)}(t), X_{IN}^{(2)}(t), \dots, t \right), \quad (8)$$

where $X^i(t) = (d^i X(t))/dt^i$; $X_{OUT} = (X_{OUT1}, X_{OUT2}, \dots)^T$.

For uncertainty evaluation, the inverse response characteristic of a system or element of a system is needed (see Figure 2). Usually, only the forward-response characteristic of a measuring system is known; and its inversion might be awkward.

Therefore, in particular in case of dynamic systems, it might be useful to describe a system in the state space by means of the so-called observation equation [8, 9],

$$\begin{aligned} \dot{Z}(t) &= f_S [Z(t), X_{IN}(t), t], \\ X_{OUT}(t) &= g_{OUT} [Z(t), X_{IN}(t), t], \end{aligned} \quad (9)$$

where $Z(t)$ is the vector of the so-called internal state variables which represent the foretime knowledge about the system.

For a lti system, Equations (9) become

$$\begin{aligned} \dot{Z}(t) &= A \cdot Z(t) + B \cdot X_{IN}(t) \\ X_{OUT}(t) &= C \cdot Z(t) + D \cdot X_{IN}(t). \end{aligned} \quad (10)$$

A, B, C, D , are vectors representing constant factors, $A = (A_1, A_2, \dots)^T$ and so on.

Example 3: Cantilever of an atomic force microscope (AFM)

The dynamic behaviour of the cantilever of an AFM may be modelled as a spring-mass-damper-system (see Figure 7), $mX^{(1)} = F_a - cX - dX^{(1)}$, where X is the deflection (output quantity), F_a being the excitation force, c being the spring constant, and d being the clamping factor. Using velocity $X^{(1)}$ and acceleration $X^{(2)}$ as state variables, state-observation equations read

$$\begin{pmatrix} X^{(1)} \\ X^{(2)} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -\frac{c}{m} & -\frac{d}{m} \end{pmatrix} \cdot \begin{pmatrix} X \\ X^{(1)} \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1}{m} \end{pmatrix} \cdot F_a$$

$$X = (1 \quad 0) \cdot \begin{pmatrix} X \\ X^{(1)} \end{pmatrix} + 0 \cdot F_a.$$

For proper treatment of real stochastic influences and processes, for example, noise, it might be useful to describe continuous systems and signals in the frequency and *Laplace* domain, respectively.

For linear or *linearised* systems, *Laplace* transformation [8,9] offers a suitable way to straightforwardly characterise a transmission system by its so-called transfer function,

$$\mathcal{L}[X_{\text{OUT}}(t)] = \mathcal{L}[X_{\text{IN}}(t)] \cdot G(s), \tag{11}$$

where s is the *Laplace* operator, $\mathcal{L}[X_{\text{OUT}}(t)]$ the *Laplace* transformed output signal of the system or element, and $\mathcal{L}[X_{\text{IN}}(t)]$ is the *Laplace* transformed input signal. $G(s)$ is the so-called system function,

$$G(s) = \frac{\sum_{\mu=0}^m b_{\mu} \cdot s^{\mu}}{\sum_{v=0}^n a_v \cdot s^v}, \tag{12}$$




Fig. 7. Example: (a): Cantilever of an atomic force microscope; (b) Spring-mass-damper model of the cantilever. See text for symbols.

which unambiguously defines the behaviour of a system. This is a useful and proven way of characterising continuous systems in the frequency domain, that originates in automatic control. The decisive limitation is that this methodology only applies to linear(ised) systems with one input and one output. Nevertheless, it is applicable to a large number of problems in measurement.

2.5 Description of stochastic signals

Usually, measurement signals cannot be represented by deterministic signal models. Therefore, it seems to be more appropriate to use stochastic signals. Analogous to statistical methodology, one may define an amplitude density of a stochastic signal (see Figure 8) [10]:

$$h(X) = \lim_{\substack{\Delta X, \Delta t_i \rightarrow 0 \\ T \rightarrow \infty}} \frac{1}{\Delta X \cdot T} \cdot \sum_{i=1}^n \Delta t_i \tag{13}$$

and, hence, compute mean and variance:

$$\bar{X} = \frac{1}{T} \int_0^T X(t) df = \int_{-\infty}^{+\infty} h(X) X dX, \tag{14}$$

$$V[X] = \frac{1}{T} \int_0^T [X(t) - \bar{X}]^2 df = \int_{-\infty}^{+\infty} h(X) [X(t) - \bar{X}]^2 dX. \tag{15}$$

Neither the mean nor the variance contains information about the conservation tendency of the signal. This is provided by the so-called autocorrelation function which expresses the intrinsic coherence of the signal. It reads:

$$\Phi_{XX}(\tau) = \lim_{T \rightarrow \infty} \int_{-T/2}^{+T/2} X(t) \cdot X(t + \tau) df, \tag{16}$$




Fig. 8. Stochastic signal [10].




Fig. 9. Examples of time-variant stochastic signals together with their corresponding autocorrelation functions in time and frequency domains. Picture source: Profos [10].

where τ means a period $t_i - t_{i-1}$. Figure 9 shows three examples of time-variant signals $X(t)$ together with their corresponding autocorrelation functions.

Fourier transformation of the autocorrelation function yields the so-called (auto) power-spectrum density $S_{XX}(\omega)$ which expresses a stochastic signal in the frequency domain:

$$S_{XX}(\omega) = \int_{-\infty}^{+\infty} \Phi_{XX}(\tau) \cdot e^{-j\omega\tau} \cdot d\tau. \tag{17}$$

The power-spectrum density may, for example, be used to characterise an external random disturbance.

The reverse transformation,

$$\Phi_{XX}(\tau) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} S_{XX}(\omega) \cdot e^{j\omega\tau} d\omega \tag{18}$$

is known as *Wiener-Chinchine*-theorem.

3 Systematic approach to the modelling of measurements

Establishing a model for a measurement at least to practitioners often appears to be the most difficult task in uncertainty evaluation. It requires a complete understanding of the underlying physics and the measurement strategy under study. Neither the GUM [1, 2] nor other basic documents provide any assistance on that task. Moreover, from the scientific point of view it does not seem possible to develop a general theory that allows for designing a measurement model stringently. Nevertheless, it has been demonstrated that there is a possibility to achieve systematic modelling for uncertainty analysis based on the idea of the cause–effect propagation of measurement signals.

The modelling concept introduced here in brief is explained in detail in [5]. It is mainly based on the idea of the classical measuring chain and, therefore, refers to the ISO-GUM procedure [1] that only applies to linear or linearisable models satisfying Equation (4).

The second idea refers to the fact that the method of measurement [11] employed is reflected in the structure of the model. This results in only a few generic model structures [5]. For example, it is known that direct measurements (deflection method) always result in a linear unbranched chain reaching from cause to effect. Substitution always results in conjoining cause–effect chains, having a comparing measuring instrument at its end. Compensation needs to be modelled by means of a closed loop [5].

3.1 Standard modelling components

Based on the practical assumptions made in Section 2.2, a generic linear or linearised transmission element k that is part of a measuring system or process can generally be expressed by the scalar product of an input vector \mathbf{X}_{kIN} and a transmission vector \mathbf{A}_k (compare Equation (2) and Figure 4):

$$X_{kOUT} = \mathbf{X}_{kIN}^T \cdot \mathbf{A}_k \quad (19)$$

The vector of the transmission factors $\mathbf{A}_k = (A_{k1}, \dots, A_{ki}, \dots, A_{kn})^T$ basically represents the internal parameters of the element k . For practical reasons, it is split into a (constant) expectation vector \mathbf{a}_k and a parameter-deviation vector $\delta\mathbf{A}_k$, $\mathbf{A}_k = \mathbf{a}_k + \delta\mathbf{A}_k$. The expectation vector represents the best estimates or nominal values of the internal parameters of the element, such as, for example, the so-called k -factor of a strain gauge stated by the manufacturer. The parameter-deviation vector represents the incompleteness in knowledge about the (relevant) system parameters, such as, for example, about a piston area of a piston gauge.

The individual deviation δA_{ki} are called multiplicative or deforming deviation. They influence the transmission factors. The (incomplete) quantitative knowledge about them is usually expressed in terms of uncertainties [5].

If a transmission element serves to multiplicatively combine two or more input quantities/signals, such as, for example, to multiply an electrical current and an electrical resistance, at least one of transmission factors A_{ki} becomes an additional input quantity $A_{ki} = Y_{ki}$. With the view to maintain linearity around the operating point, it is required that the expectation of this second input quantity is constant; $y_{ki} = \text{constant}$.

The input vector $\mathbf{X}_{k\text{IN}}^T$ is comprised of one or more input signals which are processed or simply influenced by the transmission element k in any way. Some of the input quantities $X_{ki} = Z_{ki}$ might be understood as additive influence quantities that, for example, represent the effect of any external influence, (e.g., temperature). In this case, the corresponding transmission factor would become equal to unique; $A_{ki} = 1$.

In classical measurement, we usually deal with only one or two input quantities, one transmission factor, and have to account for additive deviations (offsets). This directly leads from Equation 12 to the so-called generic linear transmission element introduced in [5] (see Figure 10). The transfer equation, consequently, reads:

$$X_{k\text{OUT}} = X_{k\text{IN}}(y_{k1} + \delta Y_{k2}) + (z_{k2} + \delta Z_{k2}) \cdot 1. \quad (20)$$

Linearity of the system and its time invariance allow for writing

$$X_{k\text{OUT}}(t) = X_{k\text{IN}}(t) \cdot (y_{k1} + \delta Y_{k2}) + (z_{k2} + \delta Z_{k2}), \quad (21)$$

where both the input and the output signal are time-dependent.

Practical modelling of technical and other processes are usually based on graphical depictions of cause-to-effect propagations. The same is suggested for systematically modelling for uncertainty analysis [5]. It is a modular procedure which consists of only three generic modelling components and five elementary modelling steps. First approaches to establishing practical procedures of modelling have been published by Bachmair [13], Kessel [14], and Kind [15]. For graphical depiction of the cause-and-effect relationships of the measurements to be analysed and modelled (see Section 3.2) it has been suggested to use only three standard modelling components [5] as follows:

- *Parameter sources* (SRC): They provide or reproduce a measurable quantity, for example, the measurand (see Figure 11).




Fig. 10. Generic linear time-invariant transmission element [5]. Symbols: see text.




Fig. 11. Graphical depiction of the standard modelling components: (a) parameter source (including material measures) and (b) indicating unit. Symbols: X_{RC} , quantity provided or reproduced by SRC; X_{SRC0} , nominal value of X_{SRC} ; $\Delta X_{SRC}(\mathbf{P}) = X_{SRC0} - X_{SRC}(\mathbf{P})$, instrumental error of the material measure depending on parameters \mathbf{P} ; $\delta X_M(\mathbf{P})$, (additive) disturbances/deviations depending on measurement conditions and external influences \mathbf{P} ; ΔX_{INSTR} , instrumental error of IND; δX_{IND} , deviation due to the limited resolution and X_{IND} , indicated quantity.

- *Transmission units* (TRANS): They represent any kind of signal processing and influencing, for example, amplification (see Figure 10).
- *Indicating units* (IND): They indicate their (physical) input quantities. If the errors and corrections that appear cannot be singled out and assigned to particular physical causes, the (summary) error (of indication) of the measuring instrument [11] ΔX_{INSTR} (see Figure 9) is allocated to the indicating unit.

3.2 Stepwise modelling procedure

The suggested modelling procedure consists of the following five elementary steps.

1. Describing the measurement, identifying the causative quantities (measurement, influence quantities) and the method of measurement [11] employed.
2. Analysing the measurement, decomposing it into its functional constituents, and, in turn, establishing graphically the cause-and-effect relationship for the fictitious ideal (unperturbed) measurement in terms of the above-described standard modelling components (see Section 3.1).
3. Incorporating all imperfections, effects of incomplete knowledge about quantities, and influences that may perturb the fictitious ideal measurement.

Establishing graphically and, in turn, mathematically the cause-and-effect relationship for the real (perturbed) measurement.

Note: the incorporation of imperfections, such as external influences, incomplete knowledge about parameters, and instabilities, is carried out by utilising correction factors and deviations that are related to the fictitious operating point which is represented by the operating-point vector.

4. Identifying mutually dependent input quantities and including resulting correlations.

Note: there are three possible ways to include correlation [5, 15].

5. Reversing the mathematical cause-and-effect relationship with a view to explicitly deriving the model equation.

Example 4: Calibrating a nonautomatic scale

This example including the stepwise modelling procedure is described and explained in detail in [5]. Here, Figure 12(a) illustrates the calibration of a non-automatic scale by direct measurement of a standard weight. Figure 12(b) shows the respective graphical cause-effect relationship for the real (disturbed) measurement. For the sake of simplification, possible correlation has been neglected.

By stepping through this graphical model from cause to effect, from Figure 12(b) the following mathematical cause-effect relationship can be derived:

$$W_{IND} = (W_{S0} - \Delta W_S)k_B + \delta W_{CPL}(\mathbf{P}) + \delta W_M(t_a) + \Delta W_{INSTR} + \delta W_{IND}. \quad (22)$$

Reversing this cause-effect equation yields the model equation for evaluating the measurement uncertainty:

$$\Delta W_{INSTR} = W_{IND} - \delta W_{IND} - \delta W_M(t_a) - \delta W_{CPL}(\mathbf{P}) - (W_{S0} - \Delta W_S)k_B. \quad (23)$$

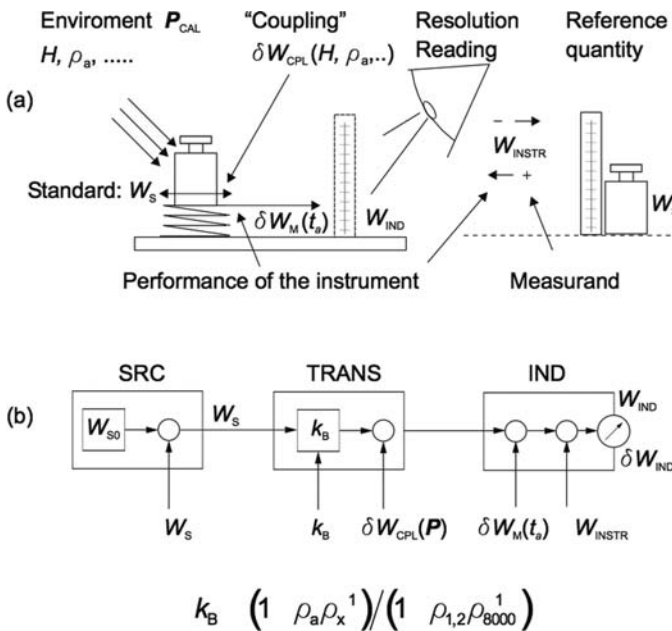


Fig. 12. Simplified example of a calibration of a scale. (b) Respective cause-and-effect relationship modelled for the real measurement. Symbols; W_S , weighing value provided by the standard; W_{S0} , nominal value of the standard and $\Delta W_S = W_{S0} - W_S$, instrumental error of the standard, (e.g., given in a calibration certificate); W_{IND} , indication; δW_{IND} , resolution; k_B , air-buoyancy correction factor (see Figure 12(b)).

4 Effects of imperfectly modelled measurements

4.1 Disregarding nonlinearities

Disregarding nonlinearities usually means that the real (nonlinear) characteristic of a measuring system or element is replaced with an idealised (linear) characteristic, along with neglecting the difference between them. With regard to Figure 6, this yields a nonlinearity error (of indication)

$$\Delta X_{NL} = X - X_{ID}. \tag{24}$$

Related to this indication span, one obtains the relative error

$$\Delta X_{NLrel} = \frac{X - X_{ID}}{x_{ID} - x_L}. \tag{25}$$

The value of this (relative) nonlinearity error depends on the nonlinearity of the real characteristic and the particular measurement value.

Let the real characteristic be given by the function

$$X = f(Y, \mathbf{Z}_0),$$

where \mathbf{Z}_0 is the operating point vector. Thus, Taylor-series expansion of this equation reads

$$X = f(y_L, \mathbf{Z}_0) + \left. \frac{\partial f(Y, \mathbf{Z}_0)}{\partial Y} \right|_{Y=y_L} (Y - y_L) + \frac{1}{2} \left. \frac{\partial^2 f(Y, \mathbf{Z}_0)}{\partial Y^2} \right|_{Y=y_L} (Y - y_L)^2 + \dots \tag{26}$$

Because the sensitivity is generally defined as $c(Y, \mathbf{Z}) = \partial f(Y, \mathbf{Z})/\partial Y$, Equation (26) becomes

$$X = f(y_L, \mathbf{Z}_0) + c(y_L, \mathbf{Z}_0) \cdot (Y - y_L) + c^{(1)}(y_L, \mathbf{Z}_0) \cdot (Y - y_L)^2 + \dots, \tag{27}$$

where

$$c^{(1)}(y_L, \mathbf{Z}_0) = \left. \frac{\partial c(Y, \mathbf{Z}_0)}{\partial Y} \right|_{Y=y_L}.$$

By introducing the sensitivity of the idealised characteristic (see Figure 6) $c_{ID} = (x_H - x_L)/(y_H - y_L)$, Equation (25) may be transformed into

$$X - x_L = c_{ID}(Y - y_L) \cdot (1 + \Delta X_{NLrel}). \tag{28}$$

Therefore, one obtains the following approximation for the relative nonlinearity error [7],

$$\Delta X_{NLrel} \approx \frac{c(y_L, \mathbf{Z}_0) + \frac{1}{2}c^{(1)}(y_L, \mathbf{Z}_0) \cdot (Y - y_L)}{c_{ID}} - 1. \tag{29}$$

For determining the measurement uncertainty associated with the nonlinearity error X_{NLrel} , one has to distinguish the following two cases.

- (a) The value of the nonlinearity error is not taken into consideration but it is considered as an additional uncertainty contribution.
- (b) One corrects the result for the value of the linearity order but wants to account for the uncertainty associated with its estimation/approximation.

In the usual case (a), one simply has to take the squared nonlinearity error as an additional variance which has to be implemented in uncertainty propagation [16,17].

In case (b) a standard uncertainty has to be determined that can be attributed to the estimated nonlinearity error (see Equation (29)). This means that Equation (29) might be interpreted as a mathematical submodel, the constituents of which, such as, for example, the sensitivities $c(y_L, \mathbf{Z}_0)$ and c_{ID} , are to be evaluated statistically by assigning appropriate state-of-knowledge pdfs to them [1,2].

4.2 Dynamic uncertainty contributions

Basically, one might define an instantaneous dynamic error $\delta X_{DYN}(t)$ as the deviation between an influenced or disturbed or processed time-variant signal $X_D(t)$ and the original unaffected or undisturbed input signal $X(t)$ at a given time t :

$$\delta X_{DYN}(t) = X_D(t) - X(t). \tag{30}$$

Figure 13 illustrates this instantaneous dynamic error.

Basically, as in signal and system theory, the dynamic error may be expressed in the frequency and *Laplace* domain as well. For a linear system in the *Laplace* domain holds $X_D(s) = X(s) \cdot G(s)$ where $G(s)$ is the transfer function of the nonideal system or element. The dynamic error might be recognised as being just the deviation caused by the system:

$$\delta X_{DYN}(s) = X(s) \left[\frac{X_D(s)}{X(s)} - 1 \right] = X(s) [G(s) - 1]. \tag{31}$$




Fig. 13. Instantaneous dynamic error. Picture source: Profos [10].

Consequently,

$$\delta X_{\text{DYN}}(t) = \mathcal{L}^{-1} \left\{ X(s) [G(s) - 1] \right\}. \quad (32)$$

In case of deterministic signals, for example, a step function, this error is knowable in principle if two out of three variables (see Equation (18)) are known.

In practical measurement, the definition and use of an instantaneous dynamic error is useful for the treatment of slow transition processes only, such as, for example, temperature equalisation of a thermometer (see Example 5).

Example 5: Dynamic error of an immersed liquid-in-glass thermometer

A liquid-in-glass thermometer that indicates the ambient temperature ϑ_a plus its (statical) instrumental error, $\vartheta_{\text{IND}} = \vartheta_a + \Delta\vartheta_{\text{INSTR}}$, is at time t_a being immersed into a water bath having the temperature ϑ_B (see Figure 13).

The cause-effect relationship of this measurement and temperature equalisation process may be expressed by

$$\vartheta_{\text{IND}} = \vartheta_B + \Delta\vartheta_{\text{INSTR}} - T \frac{d\vartheta_{\text{Th}}}{dt}, \quad (33)$$

where

$$T = \frac{m \cdot c}{\alpha \cdot A}, \quad \frac{d\vartheta_{\text{Th}}}{dt} \approx \frac{d\vartheta_{\text{IND}}}{dt},$$

m is the mass of the temperature, c is the specific heat capacity, α is the thermal transition factor, and A is the immersed surface area of the thermometer.

Consequently, the model equation for evaluating the measurement uncertainty becomes

$$\vartheta_B = \vartheta_{\text{IND}} - \Delta\vartheta_{\text{INSTR}} + T \frac{d\vartheta_{\text{Th}}}{dt}, \quad (34)$$

where $T(d\vartheta_{\text{Th}}/dt) = \delta\vartheta_{\text{DYN}}(t)$ being the dynamic error component, the expectation of which is approximately

$$\delta\vartheta_{\text{DYN}}(t) = (\vartheta_B - \vartheta_a) \cdot \exp\left(-\frac{t - t_0}{T}\right).$$

The uncertainty associated with this expectation may be derived from a rectangular probability density distribution accounting for the incomplete knowledge of the immersion process and the above heat exchange parameters.

In the case of quickly changing processes or stochastic signals, the use of an instantaneous dynamic error is not very useful because it is permanently changing with time, may transitionally disappear and so on. Therefore, similar to statistics, one prefers mean values, such as, for example, a *dynamic mean-square error* [10]. This is given by

$$\overline{\delta X_{\text{DYN}}^2} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T ((X_{\text{D}}(t) - X(t))^2) dt. \quad (35)$$

In the frequency domain, it is usually not possible to describe measurement processes by means of deterministic signals. Instead of this, one uses the (auto) power-spectrum density of a stochastic signal (see Section 2.4). This yields the following transfer relationship,

$$S_{XD XD}(\omega) = S_{XX} \cdot |G(j\omega)|^2. \quad (36)$$

Consequently, the dynamic mean-square error becomes [10]

$$\overline{X_{\text{DYN}}^2} = \frac{1}{\pi} \int_0^{\infty} S_{XX}(\omega) \cdot |G(j\omega) - 1|^2 d\omega \quad (37)$$

and

$$\overline{X_{\text{DYN}}^2} = \frac{1}{\pi} \int_0^{\infty} S_{XD XD}(\omega) \cdot |1 - G(j\omega)^{-1}|^2 d\omega, \quad (38)$$

respectively.

For uncertainty evaluation, this dynamic mean-square error of a stochastic signal basically and apart from spectrum dependencies might be used as a squared standard uncertainty. Example 6 explains this further.

Example 6: Accounting for external stochastic disturbances

Suppose the physical input of an electrical measuring system such as, a digital multimeter is being disturbed by induced noise. Furthermore, it is supposed that the measuring system can be described by a linearised cause–effect model (see Figure 14) and the induced noise is quantified by its power-spectrum density $S_{ZZ}(\omega)$. Thus, a separate treatment of static and dynamic errors (see Figure 15) yields the effect of the stochastic disturbance at the physical output of the system, for example, in indication. The effect is the system-filtered noise




Fig. 14. Illustration of Example 5: dynamic error of an immersed liquid-in-glass thermometer. Symbols; see text.

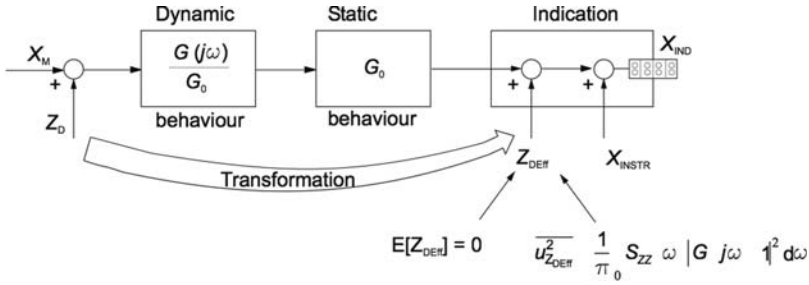


Fig. 15. Example: Taking external stochastic disturbances into account (for explanation see text). Symbols: X_M , measurand; $G(j\omega)$, dynamic transfer behaviour; G_0 , static transmission factor; Z_D , induced disturbing signal; Z_{DEFF} , transformed filtered signal Z_D ; ΔX_{INSTR} , static error of indication; X_{IND} , indication.




Fig. 16. Filter effect of the measuring system. For symbols see Figure 15 and text. Picture source: Profos [10]

(see Figure 16). In the case of a power-limited signal, the principle of maximum information entropy (pme) yields Gaussian distribution of the values of the signal Z_{DEFF} , an expectation $E[Z_{DEFF}] = 0$ and an average variance

$$\overline{u_{Z_{DEFF}}^2} = \frac{1}{\pi} \int_0^\infty S_{ZZ}(\omega) |G(j\omega) - 1|^2 d\omega. \tag{39}$$

This uncertainty can easily be transformed back to the measurand as one of the uncertainty contributions to be taken into account.

Acknowledgements

The author wishes to thank Stefan Heidenblut, Bernd Siebert and Albert Weckenmann for valuable discussions on modelling of measurements in preparing this chapter.

References

1. ISO 1993 *Guide to the Expression of Uncertainty in Measurement* (GUM) 1st edn, corrected and reprinted 1995 (Geneva: International Organization for Standardization (ISO))
2. BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP and OIML Evaluation of measurement data – Supplement 1 to the *Guide to the Expression of Uncertainty in Measurement* – Propagation of distributions using a Monte Carlo method (Draft)
3. Sivia, D. S.: 1996 *Data Analysis – A Bayesian Tutorial* (Oxford: Clarendon)
4. Lira, I.: *Evaluating the Measurement Uncertainty*. Institute of Physics Publishing Bristol and Philadelphia, 2002
5. Sommer, K.-D.; Siebert, B.R.L.: Systematic Approach to the Modelling of Measurements for Uncertainty Evaluation. *Metrologia* 43 (2006) 4
6. Sommer, K.-D.; Siebert B.R.L.: Praxisgerechtes Bestimmen der Messunsicherheit nach GUM/*Practical Determination of Measurement Uncertainty under GUM*. Ein Technisches Messen 71 (2004) 2, 52 - 66 (ISSN 0178-2312)
7. Kiencke, U.; Eger, R.: 1995 *Messtechnik – Systemtheorie* für Elektrotechniker (Berlin: Springer)
8. Kiencke, U.; Jäkel, H.: *Signale und Systeme*. Oldenbourg Verlag München Wien 2005
9. Lunze, J.: *Regelungstechnik 1. Systemtheoretische Grundlagen Analyse und Entwurf einschleifiger Regelungen*. Springer Berlin Heidelberg New York 2007
10. Hemmi, P.; Profos, P.: *Dynamische Messfehler*. in Profos, P.; Pfeifer, T. (Hrsg.): *Grundlagen der Messtechnik*. R. Oldenbourg Verlag München Wien 1997
11. *International Vocabulary of Basic and General Terms in Metrology (VIM)*. 2nd edn, ISO, Geneva 1994
12. Sommer, K.D.; Siebert, B.R.L; Kochsiek, M.; Weckenmann, A.: 2005 A systematic approach to the modelling of measurements for uncertainty evaluation. *7th Int. Symp. on Measurement Technology and Intelligent Instruments, ISMTII 2005* (University of Huddersfield (UK), 6-8 September 2005)
13. Bachmair, H.: 1999 A simplified method to calculate the uncertainty of measurement for electrical calibration and test equipment. *Proc. 9th Int. Metrology Congress* (Bordeaux 18-21 October 1999) (Nanterre Cedex: Mouvement Français pour la Qualité)
14. Kessel, W.: 2001 Messunsicherheit – einige Begriffe und Folgerungen für die messtechnische Praxis. *PTB Mitteilungen* 111 226–44
15. Kind, D.: 2001 Die Kettenschaltung als Modell zur Berechnung von Messunsicherheiten. *PTB Mitteilungen* 111 338–41
16. Lira, I.; Wöger, W.: 1998 The evaluation of the uncertainty associated with a measurement result not corrected for systematic effects. *Measurement Science and Technology*, 9, 1010–1011
17. Bich, Walter; Callegaro, Luca; Pennechi, Francesca: Non-linear models and best estimates in the GUM. *Metrologia* 43 (2006), 196–199

Approaches to Data Assessment and Uncertainty Estimation in Testing

Wolfram Bremser, Werner Hässelbarth

Federal Institute for Materials
Research and Testing (BAM), Unter den Eichen 87, 12205 Berlin, Germany
wolfram.bremser@bam.de

Summary. The following chapter presents a selection of feasible approaches to data assessment and uncertainty evaluation for single measurands (univariate case), multiple measurands with and without a functional relationship (multivariate case), and some aspects of semi-qualitative and qualitative testing.

Key words: Uncertainty, testing, data assessment, data fusion

1 Introduction

Although the following chapter deals with approaches to data assessment and uncertainty evaluation which are applicable to, and frequently used in, the field of (chemical) analysis and testing, it is not the intention to create an artificial barrier between the fields of metrology, calibration, and testing. There are, instead, many common approaches, and as far as type A evaluation is concerned, the same basic statistical tools as presented in the preceding chapters are used. The *Guide to the Expression of Uncertainty in Measurement* (GUM) [1] sets out the principles and rules, and a helpful compendium of tools for type A evaluation is [2], presented from a practitioner's point of view.

Nevertheless one may identify some peculiarities of the majority of measurement tasks in testing. First of all, due to the nature (e.g., destructive testing) or the complexity of the measurement process, the number of (repeated) observations is usually small (process control may be an exemption). Secondly, testing quite often deals with more than one measurand or one single property. Multiple measurands may either be intrinsically independent from each other, or depend on each other according to a certain (functional) relationship. In some cases, the measurement process introduces relationships (correlation) between intrinsically independent measurands such that the estimates of the measurands' values have to be considered correlated. Thirdly, it is not uncommon in testing to determine the same measurand using different,

one hopes independent, methods all having their own bias and delivering different estimates for the value of the measurand.

Data assessment in testing therefore is strongly related to combination of results, and should be based upon a sound estimation of measurement uncertainty. It should answer the question about which results should be combined in which way given the uncertainties of the results to be combined, and it should deliver the best estimates available for both the value and the corresponding uncertainty.

Finally, large areas in the field of testing deal with qualitative properties where the well-established concepts of the GUM [1] meet (serious) problems in application.

The following chapter presents a selection of feasible approaches to data assessment and uncertainty evaluation for single measurands (univariate case), multiple measurands with and without a functional relationship (multivariate case), and some aspects of semi-qualitative and qualitative testing.

2 Assessment and combination of quantitative results

One of the most intriguing and persistent problems in the evaluation of measurement data is how to combine the results obtained in several measurements of the same quantity. Here the problem is not so much how to distil from the data a good estimate of the quantity itself: most often some kind of mean will do, and different means (unweighted, weighted, etc.) will generate similar values. The big question, however, is how to estimate the uncertainty of the respective mean value. When combining discrepant results, any reasonable uncertainty estimate will need to account for the bias observed, and this task requires an appropriate mathematical model as a basis for data analysis. Depending on the model and the data evaluation technique, uncertainty estimates may vary considerably, leaving the user with a difficult decision. But even when the results agree within specified uncertainties, it is often unclear whether the data are sufficiently independent to justify the use of the factor of $1/\sqrt{N}$ in the calculation of the uncertainty of a mean value, and how to proceed if not.

2.1 Univariate case: Approaches and tools for two measurements

General

Consider the most simple case of two measurement results $a \pm u(a)$ and $b \pm u(b)$ where a and b may be single values or mean values from replicate measurements, and $u(a)$, $u(b)$ are standard uncertainties obtained from an evaluation of measurement uncertainty for a and b . The task is to determine a combined result $c \pm u(c)$.

In the absence of any specific reasons suggesting another estimate, the default choice for c is the common mean:

$$c = \frac{a + b}{2}. \quad (1)$$

Straightforward application of uncertainty propagation according to the GUM gives a combined standard uncertainty $u(c)$ as follows,

$$u(c) = \frac{1}{2} \sqrt{u^2(a) + u^2(b)}, \quad (2)$$

provided a and b are independent. Given measurement results values 100 ± 3 and 140 ± 4 , the mean value $c = 120$ obtained from Equation (1) may still be a reasonable joint estimate of the measurand. But considering the difference $\Delta = (b - a) = 40$, the standard uncertainty $u(c) = 2.5$ obtained from Equation (2) obviously is not.

This is the *discrepancy problem*: the measurement results are not consistent within specified uncertainties. Thus at least one of the uncertainties $u(a)$ and $u(b)$ was largely underestimated. Because one or both of the data $u(a)$ and $u(b)$ are invalid, the combined uncertainty $u(c)$ is invalid, too.

Another, less commonly recognised source of trouble is the *correlation problem*: the combination rule according to Equation (2) is only valid for independent measurements. If measurements have major uncertainty sources in common, a correlation term has to be included in the combined uncertainty of a mean value. As a rule, this will give a larger uncertainty than Equation (2). The uncertainty of a weighted mean requires a modification of Equation (2), and the uncertainties thus obtained are often substantially smaller than for the common mean. Last but not least, when the task is to combine the data from two sets of replicate measurements, and uncertainty evaluation is restricted to variability, the simplest approach is to pool the data.¹

Handling of discrepant results

In an ideal world the solution to the discrepancy problem would be to re-examine and amend the uncertainty budgets for the two measurement results. In practice, however, this approach is most often not applicable. The data are what they are, and the alternative is either to declare a meaningful combination impossible, or to estimate a combined uncertainty from the available data. Accommodation of the observed discrepancy by way of amended uncertainties requires some kind of model. Two models are feasible.

¹ We assume $b > a$ throughout this section to get around the use of absolute values $|b - a|$.

(a) Minimum coverage of observed differences

Consistency of measurement results $a \pm u(a)$ and $b \pm u(b)$ may be examined by comparing the difference $\Delta = (b - a)$ with its expected magnitude, as given by the combined standard uncertainty of the difference

$$u(\Delta) = \sqrt{u^2(a) + u^2(b)}. \quad (3)$$

One disregards significance levels and uses $\Delta > u(\Delta)$ as a criterion for lack of consistency, requiring action (see [4]). If $(b - a)^2 > u^2(a) + u^2(b)$, a pragmatic approach to account for the lack of coverage by the specified uncertainties is to take the difference as a combined excess uncertainty according to $(b - a)^2 - u^2(a) - u^2(b) = u_{\text{ex}}^2(a) + u_{\text{ex}}^2(b)$. Next, amended uncertainties are defined by $u_{\text{am}}^2(a) = u^2(a) + u_{\text{ex}}^2(a)$ and $u_{\text{ex}}^2(b) = u^2(b) + u_{\text{ex}}^2(b)$. Now the amended uncertainty of the mean value is calculated according to $u_{\text{am}}^2(c) = 1/4 [u_{\text{am}}^2(a) + u_{\text{am}}^2(b)] = 1/4 [u^2(a) + u^2(b) + u_{\text{ex}}^2(a) + u_{\text{ex}}^2(b)] = 1/4 [u^2(a) + u^2(b) + (b - a)^2 - u^2(a) - u^2(b)] = 1/4 (b - a)^2$. In summary the result of this uncertainty evaluation is as follows.

For compatible results

$$u(c) = \frac{1}{2} \sqrt{u^2(a) + u^2(b)}. \quad (4)$$

For discrepant results

$$u_{\text{am}}(c) = \frac{1}{2} (b - a). \quad (5)$$

Following the approach by Lyons [4], the calculation above may be justified by a measurement error model as follows. Assume that both measurements are affected by an unknown bias and model this by a random shift with zero mean and standard deviation s_{shift} . Accordingly, amended uncertainties $u_{\text{am}}(a)$ and $u_{\text{am}}(b)$ are obtained as root sums of squares of $u(a)$ and s_{shift} or $u(b)$ and s_{shift} , respectively. Assuming further that the shifts for measurements a and b are uncorrelated, and equating the expected magnitude of the difference $b - a$ with the observed value, the shift standard deviation is estimated as $s_{\text{shift}}^2 = 1/2 [(b - a)^2 - u^2(a) - u^2(b)]$. This gives a final result of $u_{\text{am}}^2(c) = 1/4 (b - a)^2$.

(b) Extended coverage of observed differences

In the approach by Levenson et al. [5], the observed difference $(b - a)$ is accommodated by an additional uncertainty contribution based on a rectangular probability distribution. The uncertainty of the mean value is directly amended according to $u^2(c) = 1/4 [u_{\text{am}}^2(a) + u_{\text{am}}^2(b)] + u_{\text{ex}}^2(c)$, where the additional contribution is taken to be the variance of a rectangular distribution with zero mean and a width of $(b - a)$. The standard deviation of this distribution is $u_{\text{ex}}(c) = (b - a) / 2\sqrt{3}$. This gives an amended uncertainty of

$$u_{\text{am}}(c) = \frac{1}{2} \sqrt{u^2(a) + u^2(b) + \frac{(b - a)^2}{3}}. \quad (6)$$

Handling of Correlated Result

If measurements have major uncertainty sources in common, a correlation term has to be included in the combined uncertainty of a mean value. Otherwise the uncertainty may be substantially underestimated. For a mean of two results, this is not a major issue because the correct uncertainty is at most by a factor of $\sqrt{2}$ larger than the value calculated from Equation (2). However, for a mean of $N \geq 10$ results the correct uncertainty may be up to a factor of \sqrt{N} larger than the value obtained without correlation, and this certainly is a major issue.

Uncertainty of Weighted Means

For a weighted mean of two measurement results,

$$c_w = \frac{w_a \cdot a + w_b \cdot b}{w_a + w_b} \quad (7)$$

with weights w_a, w_b , the combined standard uncertainty is given by

$$u(c_w) = \sqrt{\frac{w_a^2 u^2(a) + w_b^2 u^2(b)}{(w_a + w_b)^2}}. \quad (8)$$

Equation (8) also applies if the inverse variances are used as weights ($w_a = 1/u^2(a)$, $w_b = 1/u^2(b)$). Given discrepant results $a \pm u(a)$ and $b \pm u(b)$, it may appear doubtful to use an inverse-variance weighted mean. This would only be reasonable if it could be taken for granted that smaller uncertainty correlates with smaller bias. Assuming this, one may use the weighted mean with variances amended for lack of consistency according to $u_{\text{am}}^2(a) = u^2(a) + u_{\text{ex}}^2$, $u_{\text{am}}^2(b) = u^2(b) + u_{\text{ex}}^2$ with $u_{\text{ex}}^2 = 1/2 [(b - a)^2 - u^2(a) - u^2(b)]$. This gives

$$c_w = \frac{1}{2} \left[(a + b) - \frac{u^2(b) - u^2(a)}{b - a} \right] \quad (9)$$

and

$$u(c_w) = \frac{1}{2} \sqrt{(b - a)^2 - \frac{[u^2(b) - u^2(a)]^2}{(b - a)^2}}. \quad (10)$$

Compared with the common mean $c = 1/2 (a + b)$, the weighted mean is shifted towards a if $u(b) > u(a)$ and towards b if $u(b) < u(a)$.

Pooling replicate data

Given data $A = \{a_1, a_2, \dots, a_m\}$ and $B = \{b_1, b_2, \dots, b_n\}$ from two sets of replicate measurements, the most simple combination is to pool the data, take

the grand mean (c_{pool}), and calculate the standard deviation of the pooled data (s_{pool}). If uncertainty evaluation is restricted to variability, a promising candidate for the standard uncertainty of the grand mean is

$$u(c_{\text{pool}}) = \frac{s_{\text{pool}}}{\sqrt{m+n}}. \quad (11)$$

It may be doubtful whether such a combination makes sense. If the data A and B are supposed to be for the same measurand, there should be some overlap between the range of A and the range of B . Otherwise there is a significant bias between A and B . Similar to the case of two results $a \pm u(a)$ and $b \pm u(b)$, a decision has to be taken as to whether this bias is still acceptable and data combination, accommodating the bias, should be undertaken.

If so, taking the grand mean (c_{pool}) is an appropriate default combination. However, if there is a bias between A and B , Equation (11) will not provide a valid standard uncertainty of the grand mean. The problems of assessing bias between A and B and estimating uncertainty for the grand mean accounting for *between* groups bias are most conveniently addressed by one-factorial ANOVA which is exhaustively covered in textbooks (see, e.g., [3]). ANOVA delivers an estimate for the *between* groups variance $s_{\text{between}}^2 = 1/2 [(\bar{b} - \bar{a})^2 - s^2(\bar{a}) - s^2(\bar{b})]$ which is positive given the observed bias. The amendment to $u(c_{\text{pool}})$ is

$$u_{\text{am}}^2(c_{\text{pool}}) = \frac{s_{\text{within}}^2}{2n} + \frac{s_{\text{between}}^2}{2} = \frac{(\bar{b} - \bar{a})^2}{4} \quad (12)$$

which is similar to Equation (5).

2.2 Univariate case: Approaches and tools for multiple measurements

With just two measurement results $a \pm u(a)$ and $b \pm u(b)$ there is only very limited scope for statistical data processing. However, if a larger number N of measurement results $x_i \pm u(x_i)$ have to be combined into a mean value and an associated standard uncertainty, $\bar{x} \pm u(\bar{x})$, statistical tools for analysing and processing the data become available. One assumes reasonably well-behaved data such that the common mean or a weighted mean can be used. Robust means, accommodating outliers or skewed data, are addressed in a later part.

According to GUM principles, the expanded uncertainty U should cover a large fraction of the distribution of values that could reasonably be attributed to the measurand. As a rule, $U = k \cdot u$ is taken to effect about 95% coverage. Given the case that none of the results $x_i \pm u(x_i)$ is obviously invalid, the interval $\bar{x} \pm U(\bar{x})$ should cover the entire range. Targeting for minimum 95% coverage, we would put $\bar{x} = 1/2 [x_{\text{max}} + x_{\text{min}}]$, that is, use the midrange value as an estimate for the mean, $U(\bar{x}) = 1/2 [x_{\text{max}} - x_{\text{min}}]$, and therefore estimate the standard uncertainty of the mean as

$$u(\bar{x}) = \frac{x_{\max} - x_{\min}}{4}. \quad (13)$$

Interestingly enough, this uncertainty estimate is rather optimistic for $N = 2$ results due to the fact that one targets at 95% coverage whereas previously the target was “one-standard uncertainty” coverage or better. On the other hand, for large numbers N of measurement results, the uncertainty estimate above would rather be considered as unduly pessimistic, because cancellation of measurement errors is completely ignored. In other words, Equation (13) assumes total correlation among the measurements results. An alternative approach would assume uncorrelated measurements results and introduce the factor $1/\sqrt{N}$ to account for error cancellation. This would give

$$u(\bar{x}) = \frac{x_{\max} - x_{\min}}{4\sqrt{N}}. \quad (14)$$

Most often the ‘truth’ will be somewhere in between.

Approaches based on analysis of variance

As already mentioned, analysis of variance (ANOVA) is an extremely versatile and powerful tool for data analysis and data consolidation. Among numerous other applications, ANOVA has been, and still is, used as the preferred tool for the evaluation of interlaboratory studies for value assignment of certified reference materials. Comprehensive guidance on this issue is given in ISO Guide 35 [6]. A definitive treatment of the use of ANOVA for combination of measurement results was given by Cochran [7].

The main objectives of data analysis by ANOVA are (i) to decide whether the different groups of replicates are mutually consistent and can be pooled, and (ii) to determine an appropriate estimate for the standard deviation of the grand mean, depending on the outcome of step (i). For a one-factorial layout, the data analysis is based on a statistical model, where data variability is described by combination of two effects: random *withingroup* variations and a *between*groups influence of the factor (e.g., different laboratories) according to

$$x_{lm} = \mu + \beta_l + \varepsilon_{lm}. \quad (15)$$

In this equation μ is the actual value of the measurand, β_l the bias of measurements made in the l th laboratory, and ε_{lm} the random error in the m th replicate measurement of that laboratory. The β 's and ε 's are supposed to be random variables with expectation zero and variance σ_{bias}^2 and σ_{rep}^2 , respectively. As a consequence of the bias term β shared by measurements made in the same laboratory, replicates are correlated, with a covariance of $\text{cov}(x_{lp}, x_{lq}) = \sigma_{\text{bias}}^2$ whereas measurements from different laboratories are independent. The variance of the grand mean is

$$\text{var}(\bar{x}) = \frac{\sigma_{\text{bias}}^2}{L} + \frac{\sigma_{\text{rep}}^2}{L \cdot M} \quad (16)$$

and takes up an appropriate estimate for *betweengroups* bias provided it exists, hence $s_{\text{bias}}^2 > 0$. In the case that the estimate $s_{\text{bias}}^2 \leq 0$, this means that the *withingroup* variability covers the overall data variability, and there is no need to introduce a *betweengroups* bias as an additional source of variability.

An ANOVA can also be performed without individual replicate data provided experimental standard deviations are substituted by GUM-type standard uncertainties. This may not be fully compatible with the statistical model used in the original ANOVA but should be good enough as a pragmatic approach. Let $\bar{x}_l \pm s(\bar{x}_l)$ be the laboratory mean values and the associated standard deviations: then the decision criterion for data consistency reads

$$s_{\text{means}}^2 \leq \frac{1}{L} \sum_l s^2(\bar{x}_l) \tag{17}$$

and for inconsistency

$$s_{\text{means}}^2 > \frac{1}{L} \sum_l s^2(\bar{x}_l) \tag{18}$$

with s_{means} being the standard deviation of the laboratory means. According to common understanding, ANOVA assumes comparable uncertainties $u(x_l)$; chi-squared techniques do not require that.

Approaches based on the chi-squared technique

In the ANOVA approach, consistency of results was evaluated using the average difference of two squares, $\langle (x_k - x_l)^2 - u^2(x_k - x_l) \rangle_{k \neq l}$. Instead of this, one may consider the average quotient of the squares concerned, $\langle (x_k - x_l)^2 / u^2(x_k - x_l) \rangle_{k \neq l}$. Inserting the general expression for the variance of a difference, $u^2(x_k - x_l) = u^2(x_k) + u^2(x_l) - 2u(x_k, x_l)$, the average quotient of squares is obtained as

$$\chi_{\text{std}}^2 = \frac{1}{L(L-1)} \sum_k \sum_{l \neq k} \frac{(x_k - x_l)^2}{u^2(x_k) + u^2(x_l) - 2u(x_k, x_l)}. \tag{19}$$

Given mutually consistent results, the expected value of χ_{std}^2 is 1. Therefore $\chi_{\text{std}}^2 \leq 1$ would be taken to indicate consistency and $\chi_{\text{std}}^2 > 1$ to indicate inconsistency. A comprehensive treatment of the pair-difference chi-squared technique for assessing the mutual compatibility of key comparison data is given in [8].

As before, an excess variance may be introduced to account for lack of consistency, that is, lack of coverage of the differences $(x_k - x_l)$ by the specified uncertainties $u(x_k)$ and $u(x_l)$ concerned. This cannot be done by difference as before. Instead, the excess variance is introduced in the expression for χ_{std}^2 , and the expression so obtained is equated to 1. This gives a nonlinear equation for the excess variance u_{ex}^2 as follows,

$$\frac{1}{L(L-1)} \sum_k \sum_{l \neq k} \frac{(x_k - x_l)^2}{u^2(x_k) + u^2(x_l) + 2u_{ex}^2 - 2u(x_k, x_l)} = 1, \quad (20)$$

which may be solved for u_{ex}^2 using numerical methods. With this excess variance, for the case of inconsistent data the variance of the mean value is

$$u_{am}(\bar{x}) = \frac{1}{L} \sqrt{\sum_l u^2(x_l) + L \cdot u_{ex}^2} \quad (21)$$

whereas for consistent results it is given by

$$u(\bar{x}) = \frac{1}{L} \sqrt{\sum_l u^2(x_l)}. \quad (22)$$

The estimate for the excess variance could also be used if the common mean \bar{x} is replaced by a weighted mean. For many purposes, the variances and covariances relating to the mean value may be neglected, and the amended variances of the deviations from the mean are approximated according to $u_{am}^2(x_l - \bar{x}) \approx u^2(x_l) + u_{ex}^2$. This gives a modified equation for the excess variance as follows,

$$\sum_l \frac{(x_l - \bar{x})^2}{u^2(x_l) + u_{ex}^2} = L - 1. \quad (23)$$

This approach was introduced by Paule and Mandel [9], and has been recognised as one of the leading approaches for the combination of measurement results. It was in particular considered by NIST for in-house characterisation of reference materials [10]. The statistical basis of the Paule–Mandel approach has been a subject of mathematical investigations [11], and its performance was assessed by simulation studies (see, e.g., [12]). Reference [13] compares various approaches for combining interlaboratory data based on random effects models, including ANOVA and the chi-squared approach of Paule and Mandel, from a generic statistical perspective.

The Birge approach to account for lack of consistency

An alternative approach was introduced by Birge [14] to account for discrepancies in interlaboratory studies for the determination of fundamental constants. It is based on the assumption that all participants have underestimated their uncertainties by about the same factor. Accordingly a factor is introduced to amend the uncertainties according to $u_{am}(x_l) = f \cdot u(x_l)$, and fixed at the smallest value effecting consistency. If consistency is assessed on the average differences $\langle (x_k - x_l)^2 - u^2(x_k - x_l) \rangle_{k \neq l}$, the Birge factor is given by $f^2 = \langle (x_k - x_l)^2 \rangle_{k \neq l} / \langle u^2(x_k - x_l) \rangle_{k \neq l} = s_x^2 / \langle u^2(x) \rangle$. Here, s_x stands for the

standard deviation of the laboratory means, and $\langle u^2(x) \rangle$ for the average of the $u^2(x_k)$. For the common mean, this gives an amended uncertainty of

$$u_{\text{am}}(\bar{x}) = \frac{1}{L} \sqrt{\sum_l f^2 \cdot u^2(x_l)} = \frac{f}{L} \sqrt{\sum_l u^2(x_l)} = \frac{s_x}{\sqrt{L}}, \quad (24)$$

This is just the ANOVA result. If consistency is assessed on the average quotients, $\langle (x_k - x_l)^2 / u^2(x_k - x_l) \rangle_{k \neq l}$, the Birge factor is given by $f^2 = \langle (x_k - x_l)^2 / u^2(x_k - x_l) \rangle_{k \neq l} = \chi_{\text{std}}^2$. This gives an amended uncertainty of the common mean as follows,

$$u_{\text{am}}(\bar{x}) = \frac{\sqrt{\chi_{\text{std}}^2}}{L} \sqrt{\sum_l u^2(x_l)}. \quad (25)$$

The Birge approach can also be used for weighted means. In this case, the Birge factor f is fixed such that

$$\sum_l \frac{(x_l - \bar{x}_w)^2}{f^2 u^2(x_l)} = L - 1 \quad (26)$$

and this factor is then used in the amended uncertainty of the weighted mean according to

$$u_{\text{am}}(\bar{x}_w) = \frac{f}{\sqrt{\sum_l 1/u^2(x_l)}}. \quad (27)$$

Approaches based on bias correction

Instead of amending the uncertainties $u^2(x_l)$ such that compatibility is achieved between the differences $(x_k - x_l)$ among the various measurements, or the deviations $(x_l - \bar{x})$ from the mean value, and the uncertainty attributed to these differences/deviations, discrepancies may be utilised to derive corrections, either of the mean value \bar{x} or the individual values x_l . In addition to this, the uncertainty on these corrections is evaluated, and this uncertainty is included in the uncertainty of the mean value obtained, either by direct correction or from the corrected individual values. A procedure for simultaneous correction of individual biases of the values x_l and amendment of individual uncertainties $u(x_l)$ is described in [15]. Reference [16] specifies a procedure for the correction of bias in mean values (e.g., based on a comparison of the mean value \bar{x} and the midrange value $x_{\text{mid}} = 1/2 (x_{\text{min}} + x_{\text{max}})$) and reviews other approaches from this perspective.

Robust statistics and related tools

Combination of different results $x_l \pm u(x_l)$ for the same quantity is a two-step process. In the first step the data x_l are combined into some kind of mean

value. In the second step the uncertainty of this mean value is evaluated, and data consistency checked. If inconsistency is revealed, either the uncertainties are amended, or the first step has to be revisited so as to decide whether another common average value performs better. Two alternatives have been offered thus far: the common mean value was used by default, supplemented by the inverse-variance weighted mean to account for major differences among the uncertainty data. This approach was also motivated by the fact that, for example, with just two measurement results $a \pm u(a)$ and $b \pm u(b)$, there is only very limited scope for statistical data processing.

However, if a major number L of measurement results $x_l \pm u(x_l)$ have to be combined, the common mean value may not be the best choice. This is a good estimate if the data x_l behave as a sample from a normal distribution. But if the data distribution is skewed, other estimates perform better. As another major problem, contamination with outliers will spoil the performance of the common mean. Thus outliers should be removed or, preferably, robust means should be used.

A well-known robust mean is the median. This is obtained by arranging the data in an increasing sequence, $x_1 \leq x_2 \leq \dots \leq x_L$. If L is an odd number, $L = 2N + 1$, the median is the middle value x_{N+1} . If L is an even number, $L = 2N$, the median is the mean of x_N and x_{N+1} . The median is obviously insensitive to outliers. On the other hand, for well-behaved data from a normal distribution the median gives similar results to the common mean. These are highly attractive features, but unfortunately this medal has a backside, too: it is impossible to evaluate the uncertainty of the median by propagation of the uncertainties $u(x_l)$, a problem shared by many other robust means. Thus the uncertainty data $u(x_l)$ cannot be used in the estimation of uncertainty for the median; that is, a major part of the information contained in the data $x_l \pm u(x_l)$ is thrown away.

If the median is used, the standard uncertainty of the median is estimated from the data x_l alone, discarding the uncertainty data $u(x_l)$. For this purpose the median of the absolute deviations $|x_k - \text{med}\{x_l\}|$ may be used with an appropriate numerical factor (to the effect that for normally distributed data the standard deviation of the mean is recovered).

Reference [17] gives an introduction to the median and its use for a robust evaluation of interlaboratory comparisons. A more detailed explanation of the techniques briefly described above can be found in [18], including worked-out examples and result discussion.

Uncertainty estimation

Assessment of data and data consistency as well as the application of combination strategies as presented above require the availability of sound estimates for the uncertainty of the results to be assessed, compared, or combined. Since the publication of the GUM [1], much has been done to interpret the basic concepts, and facilitate their application, in most different areas of both calibration and testing. Numerous publications were, and are still

being published, and various international associations and sector committees have published guides which target their respective scientific communities. Most prominent is probably the EURACHEM/CITAC Guide [20] which (i) implements the methodology of the GUM for the field of chemical analysis, and (ii) introduces concepts for using (available) data from different quality assurance measures commonly installed in the laboratories (method validation, proficiency testing, control charts, and the use of reference materials) for uncertainty estimation. These concepts are commonly called the top-down approach. Characteristics used in the 'top-down approach' which account for the influence and uncertainty contributions of groups of factors include, inter alia, repeatability, reproducibility, and recovery estimates. Estimation and use of the former are comprehensively covered by ISO 5725 [21], and interesting aspects and treatment of recovery and the uncertainty arising from recovery can be found in [22]. Because the topic is well covered by standards, documents, and guidelines, and because nowadays powerful software supporting uncertainty budget calculation is available, it is not elaborated further, but not without saying that the issue still exists. Quite recently, NORDTEST published guidelines for implementation of the concepts in environmental analysis [23], and the Codex alimentarius Commission for the field of food analysis [24]. EUROLAB developed the principles of the top-down approach into concrete handling instructions [25]. Interestingly enough, the latter make intensive use of the combination strategies, and follow the data consistency approach discussed above.

An example

The example refers to the establishment of traceability for a conventional method. The water (OH bond) content in glasses is usually determined using FTIR spectroscopy, a method which is conventional in the sense that it uses pre-defined extinction coefficients in a two-band model developed by Scholze [26]. For RM certification and establishment of traceability, a completely different primary method, namely nuclear reaction analysis (NRA), was applied. The NRA actually determines hydrogen which can easily be recalculated into molar concentrations of water provided all hydrogen is found in this form. Many years of experience in the glass-manufacturing industry prove that this is essentially the case, at least for these types of soda-lime silica glasses considered here. The nuclear reaction used, and the measurement equations for both methods are given in Figure 1.

Twenty-four samples from three glass disks were meticulously analysed in replicate. The results were combined into within-method means. For the corresponding uncertainties, amendments were made accounting for contributions from within-disk (where applicable) and disk-to-disk inhomogeneities, the repeatability of the actual measurements, and the known systematic effects of the method. For the FT-IR measurements, for example, the within-method uncertainty was estimated as (expressed for the mole fraction of water)




Fig. 1. Model equations for the determination of the value of the measurand for the two independent methods used for water content determination.

$$u^2(x_{\text{OH}}) = \frac{s^2}{n} + u_{\text{in hom disk-to-disk}}^2 + u_{\text{in hom on-disk}}^2 + u_{\text{method}}^2 \quad (28)$$

using the approaches developed in the preceding clauses. In Equation (28), the first term is the contribution from characterisation measurements, assessed as the standard deviation of the mean value of a series of measurements. The other three terms are contributions derived from a two-factorial ANOVA for the factors between-disk variability, on-disk (mostly radial) inhomogeneity, and method repeatability. This established data consistency within a method. Inconsistency was however detected between the means of both methods. It is obvious that besides the factors already considered, the model is a potential source of data incompatibility, mainly due to two factors, namely

1. The use of a two-band model with pre-defined extinction coefficients in FT-IR. The coefficients may be an acceptable, but by far not the best choice, and they come without any uncertainty.
2. A residual risk of measuring hydrogen other than OH⁻ by NRA. This risk is very low such that no correction should be applied, but an allowance made in the uncertainty budget.

An appropriate model accounting for the above factors would be

$$\begin{aligned} c_{\text{NRA}} &= c_{\text{H}_2\text{O}} + \delta' \\ c_{\text{IR}} &= f \cdot c_{\text{H}_2\text{O}} = (1 + \alpha) \cdot c_{\text{H}_2\text{O}} \\ &= c_{\text{H}_2\text{O}} + \alpha \cdot c_{\text{H}_2\text{O}} = c_{\text{H}_2\text{O}} + \delta'' \end{aligned} \quad (29)$$

The probable bias caused by hydrogen other than OH⁻ is additive. The bias caused by nonoptimal extinction coefficients is, in general, complex and non-linear. Within concentration ranges not exceeding certain limits, it may be approximated by a simple multiplicative bias term ($f = (1 + \alpha)$ in Equation (29)). Because one assumes that a result which uses the best estimates for the extinction coefficients does not differ very much from the one obtained using the Scholze coefficients, it holds $\alpha \ll 1$. For a given $c(\text{H}_2\text{O})$, one finally obtains Equation (29).




Fig. 2. Water content in mol/L (consolidated mean per disk and method) determined on 3 disks (24 samples in total) by NRA (light-grey bars) and FT-IR spectroscopy (grey bars), and consolidated mean per method including model error contribution.

The unknown δ' and δ'' are set to zero, and their effect is covered by an excess uncertainty different from zero. One assumes $u(\delta') = u(\delta'') = u(\delta)$, that is, amends the uncertainties according to Equation (28) symmetrically for both methods.

The amendment is calculated such that the within-method means become fully consistent. The 'nominal error' E_n consistency criterion is used which takes the same form for both standard and expanded uncertainties with the only difference that the critical value is 2 for standard, and unity for expanded uncertainties. Figure 2 visualises the effect. It displays the inconsistency of the within-method means (here separately for the three disks) within their uncertainties, and the total within-method means with their amended uncertainties, now fully consistent.

Based upon data consistency, the certified value is derived as the weighted mean of the within-method means, and the corresponding weighted uncertainty is assigned, resulting in a molar concentration of water in glass of $c(\text{H}_2\text{O}) = 0.0340$ mol/L with a combined standard uncertainty of $u(c) = 0.0015$ mol/L.

2.3 Multivariate case: Approaches and tools

In all cases when more than one single value is to be determined for the sample, the test or reference object (e.g., a couple of analyte concentrations for one matrix), and the values are suspected to be correlated.

Here, a joint consideration of a property vector \mathbf{p} (instead of a scalar property) and a variance/covariance matrix \mathbf{M} (instead of a standard deviation) should be carried out [19]. The problem becomes evident when functionally related properties are to be determined (by measurements with an independent parameter varying within certain limits in a domain) and tested with regard to the compatibility of measured values and a reference curve.

Before deciding on the kind of data treatment (i.e., uni- or multivariate) in situations where an assumption of possible correlations cannot be rejected, an analysis of possible correlations between the data obtained for the different property values can be helpful. The hypothesis of uncorrelated (independent) data sets can be adopted on statistical grounds when 1st order Pearson coefficients do not exceed 0.1, and possible correlations can still be neglected or disregarded as long as a limit of 0.5 is not exceeded provided further analytical evidence can be given that the method is sufficiently selective and interferences (both due to the nature of the analytes and the measuring device) are unlikely to appear. When a relationship (and thus correlation in the data) can be assumed which is generated by the underlying physics (or chemistry), multivariate data treatment is a must.

Procedures for establishing consolidated values for distributed properties shall provide the more-dimensional analogue to tools and techniques well established in univariate data assessment, that is, enable

- the definition of an average of a certain number of single measurements (applied to both a grand mean and a laboratory mean) plus corresponding standard deviation(s)
- the definition of a mean of means plus corresponding standard deviation(s)
- analysis of variance for homogeneity testing and/or *between* laboratories bias determination
- outlier tests
- tests on significant differences between determined properties.

Some multivariate data assessment tools and their possible applications to evaluation and combination of data for multiple measurands are listed in Table 1.

The multivariate E_n criterion reads $(\mathbf{c}_1 - \mathbf{c}_2)^T \cdot (\mathbf{M}_1 + \mathbf{M}_2) \cdot (\mathbf{c}_1 - \mathbf{c}_2) < E$, where \mathbf{c}_1 and \mathbf{c}_2 are vectors of values to be compared, \mathbf{M}_1 and \mathbf{M}_2 the corresponding variance/covariance matrices of the vectors, and E the expectation value at the corresponding number of degrees of freedom for an appropriate, most often a chi-squared distribution. Because most of the procedures mentioned in Table 1 refer to the bare data and generate the ‘classical’ uncertainty estimates in a pure Type A estimation step, they are not fully compatible, but adaptable to the GUM approach. It must be admitted that most of the adaptation still has to be done. Besides this fact, from an application of the univariate approaches presented in the last chapter the problem of assessing

Table 1. Possible multivariate tools for assessment and combination of data.

Tool	Description and Application
Pearson correlation matrix	Reveals possible intrinsic or extrinsic correlations.
MANOVA	Reveals significant dataset incompatibilities (between-set variations) and allows a decision on whether all single values in the sets ('pooling') or only the means of the single values in each set ('no-pooling') may be averaged.
Cluster analysis	Detects closeness of and groups in datasets and thus may indicate distant (outlying) datasets.
Discriminant analysis	Detects the (maximum) differences in datasets and classifies sets into groups. This may possibly be used as an outlier test (if groups discriminate, i.e., are distant enough). The normal classification scheme is likely to correspond to a univariate Grubbs, the jackknifed classification to a Nalimov test.
E_n criterion	Confirms or rejects dataset compatibility within the stated uncertainties.
Kernel estimation	Interpolation of generically related data
Principal component analysis	Data reduction for generically related data
Nonlinear regression	Data reduction for functionally related data

data consistency within the limits of full variance–covariance matrices arises, including the problem of comparing these matrices with each other.

Multiple measurands

Measurement results sharing major input uncertainties are correlated. This correlation has to be accounted for in any joint application, for example, if these measurement results are jointly used as input in a subsequent measurement or as input to the calculation of a derived quantity, and therefore have to be stated (see GUM (7.2.5) [1]). Even beyond the case of simultaneous measurement of several measurands considered in the GUM, correlation matters are important in any joint application of multiple correlated measurement results. A prominent application field is (natural) gas analysis where gas properties such as the calorific value, the density, or the Wobbe index are derived from gas composition. Even in the case that component interferences can be disregarded, the commonly used normalisation procedure inevitably causes strong correlations between the up to 17 components measured [27].

The handling of the variance–covariance matrix is by no means obvious. Two standalone (standard) uncertainties, obtained for two measurement results, are easily compared as to whether one is larger than the other. Extension of such comparison to two sets of (standard) uncertainties, obtained for

two sets of measurement results, is rather straightforward in the absence of correlations. Given this condition, the uncertainties would be compared one by one. However, if there are significant correlations within the two sets of measurements, the relevant covariances or correlation coefficients should be included in any such comparison. The question then is how to perform such comparison, that is, how to adequately compare two uncertainty matrices (variance–covariance matrices), in particular if these matrices are significantly nondiagonal.

A feasible comparison might be motivated by the fact that correlation between different measurement results is only important if the results are jointly utilised, and works as follows. Let x_1, x_2, \dots, x_N be the values of N different measurands X_1, X_2, \dots, X_N obtained on some measuring object, for example, the lengths of the sides of a plane triangle or the concentrations of the various components (hydrocarbons, carbon dioxide, and nitrogen) in a sample of natural gas. Furthermore let P be a quantity of the measuring object that is related to the measurands X_1, X_2, \dots, X_N by a linear equation. Given values x_1, x_2, \dots, x_N , the corresponding value $p = P(x_1, x_2, \dots, x_N)$ is obtained as

$$p = p_0 + \sum_{i=1}^N p_i x_i. \tag{30}$$

When emphasis is made on the handling of the variance–covariance matrix of the measured quantities, the uncertainty of the parameters and the uncertainty arising from errors in the model equation may be neglected, at least temporarily. For the uncertainty of p one gets

$$u^2(p) = \sum_{i=1}^N \sum_{k=1}^N p_i p_k u(x_i, x_k) = \sum_{i=1}^N \sum_{k=1}^N p_i U_{ik} p_k, \tag{31}$$

where $u(x_i, x_i) = u^2(x_i)$ represent the ‘normal’ variances, and $u(x_i, x_k)$ with $i \neq k$ the covariances. For two variance–covariance matrices \mathbf{U}, \mathbf{U}' of two different estimates for the same set of measurands, or estimates for two different, but related sets of measurands, it can be shown that the uncertainty expressed by \mathbf{U} is larger than that expressed by \mathbf{U}' if the standard uncertainty of any derived quantity obtained by linear combination of estimates is larger with \mathbf{U} than with \mathbf{U}' , or if

$$\sum_{i=1}^N \sum_{k=1}^N p_i (U_{ik} - U'_{ik}) p_k \geq 0 \quad \forall (p_1, \dots, p_N). \tag{32}$$

This is the case if and only if $(\mathbf{U} - \mathbf{U}')$ is a positive semi-definite matrix. A symmetric square matrix with real entries is positive semi-definite if and only if all its eigenvalues are nonnegative. An efficient way to test this is based on the Cholesky decomposition [28].

When the level of derived uncertainties is concerned, the derived uncertainty $u(p)$ has to be normalised with respect to the level of the coefficients p_i for comparison. One may obviously introduce the Rayleigh quotient (e.g., see [29])

$$\Phi(\mathbf{U}, \mathbf{p}) = \frac{u^2(p)}{\sum_{i=1}^N p_i^2} = \frac{\sum_{i=1}^N \sum_{k=1}^N p_i U_{ik} p_k}{\sum_{i=1}^N p_i^2} = \frac{\mathbf{p}^T \cdot \mathbf{U} \cdot \mathbf{p}}{\mathbf{p}^T \cdot \mathbf{p}} \quad (33)$$

as a measure. The minimum value of the Rayleigh quotient (over all vectors \mathbf{p}) is given by the lowest eigenvalue of the matrix, and the maximum is given by the largest eigenvalue. The level of uncertainty of a variance–covariance matrix \mathbf{U} (i.e., the range of the uncertainty of standardised linear combinations), is determined by the range of eigenvalues $\lambda(\mathbf{U})$ of the uncertainty matrix, $[\lambda_{min}(\mathbf{U}), \lambda_{max}(\mathbf{U})]$, according to $\lambda_{min}(\mathbf{U}) \leq (\mathbf{U}, \mathbf{p}) \leq \lambda_{max}(\mathbf{U})$.

The above general principles can be extended to (i) consideration of only nonnegative coefficients p_i , and (ii) the use of specific features of the matrix as for example, block diagonal, nonpositive, or nonnegative off-diagonal elements, zero row sums, and other. The application allows for an easy handling assessment and comparison of variance–covariances matrices for multiple measurands. Guidance on these specific topics and worked-out examples showing the impact of correlations can be found in [30]. Reference [31] deals with, and exemplifies, correlations between dilution series in standard addition, an experimental technique frequently used in analytical chemistry for estimation of possible matrix influences.

Functional relationships

The multiple measurands dealt with so far were interrelated (i.e., correlated), but still standalone properties of an artefact or sample. Many tasks in testing refer to properties which are related by a – known or unknown – functional relationship governed by physical principles. Irradiance curves or fluorescence efficiency/yield (in the wavelength domain), the uptake of sorbent (pressure domain), or implantation profiles (space domain) may serve as examples.

Determination of these relationships is, however, pointwise, and the assessment is based upon raw data which are replicated runs from different sources, being different instruments, methods, or laboratories, and are presented as nonequidistantly spaced data tables of different length. The relationship is measured as a set or table of paired values $\mathbf{y}|\mathbf{x} = \{y_i, x_i\}_1^m = \{y_1, x_1; y_2, x_2; y_3, x_3; \dots; y_m, x_m\}$ in a single run of an instrument or a set of instruments. Normally, this measurement is repeated q times with the same instrument yielding to q tables $\mathbf{y}|\mathbf{x}_k$. As a rule, even with the same instrument the number of data pairs (y_i, x_i) in the tables will differ; that is, each table has its own number m_k of data pairs (with k ranging from 1 to q). For more than one data source, this measuring scheme is repeated by a set of p instruments,

methods, or laboratories each providing q_j tables $\mathbf{y}|\mathbf{x}_{k,j}$. Although one will normally aim at keeping both the number m of data pairs in the table and the number q of repetitions (runs) constant, it cannot be avoided that both $m_{k,j}$ and q_j differ from source to source. As a result, the data assessment has to deal with a total number of $Q = \sum q_j$ tables containing altogether $M = \sum \sum m_{k,j}$ data pairs (y_i, x_i) . In practical situations, $\min(m_{k,j})$ and $\max(m_{k,j})$ may differ by a factor much larger than unity. These sampling tables have to be combined appropriately, and a common or reference function has to be determined.

As long as the functional relationship is known, nonlinear regression is the method of choice. Numerous approaches and algorithms have been developed and are well established (see, e.g., [28, 29]). A large class of problems of practical interest are governed by a minimum of three, mostly four functional parameters for description, namely maximum position, maximum height, and (left and right) FWHM, and may be approximated and handled by using a bell-shaped, symmetric or asymmetric model function. The advantage of regression is that it can easily cope with different numbers of points in the sampling tables $\mathbf{y}|\mathbf{x}_k$ as described above. Less attention has been paid thus far to the uncertainties which are best described by the variance–covariance matrix of the function parameters. For assessment and comparison, techniques as described in earlier sections of this chapter may be applied.

It is often felt that even thoroughly selected model functions are too stiff to interpret fine details in a complex relationship, and preference is given to interpolation techniques. Again, an enormous amount of work has been done, and many algorithms developed in this field, with spline interpolation being probably the most prominent one. Uncertainty estimation for a splined relationship meets some problems, such that frequently simpler approaches are used, including the following.

(a) Reference curve determination using grid projection

A sequence of equally or log-equally spaced sampling points (the grid) is spanned over the range covered by the sampling tables $\mathbf{y}|\mathbf{x}_{k,j}$. The distance between sampling points is chosen such that the number of intervals equals or exceeds $\max(m_{k,j})$. Such a choice avoids co-ordination problems occurring when an interval contains, for example, only one datapoint of a certain table (k_1, j_1) , but more than one datapoint of another (longer) table (k_2, j_2) . According to their x_i -value, the y_i -values are now projected onto the sampling points. The projection rule is linear interpolation. The projections are subsequently treated as individual measurement points, and means and standard deviations are calculated for them.

(b) Reference curve determination using kernel estimation

This is a piecewise fit of an approximating function to the data. Normally, one assumes polynomial approximating functions. The level of localisation

is governed by the introduction of a kernel, normally a function symmetric with respect to the origin, diminishing outside an interval h (the bandwidth). The squared deviations with respect to the approximating function are now weighted by the kernel, and a global minimisation according to

$$\sum_{i=1}^N (y_i - \sum_{j=1}^p \beta_j \cdot (x - x_i)^j)^2 \cdot K_h(x - x_i) = \min \quad (34)$$

is carried out for the parameter vector β . This provides local solutions within each band width, that is, in the vicinity of x , where an estimate for the unknown functional relationship is required. From the residuals, confidence and prediction intervals can be calculated. The main advantage of this technique consists in the fact that no specific assumptions concerning the kind of the functional relationship is needed. Because the local environment of each point to be approximated is taken into account, the method takes up and conserves more of the (hidden) functional information than a grid projection as described above. Uncertainty estimates may be derived from the overall residual SSD.

(c) Reference curve estimation using 2D averaging

This technique is motivated by the following fact. All of the above approaches consider the independent variable (e.g., a scale) error-free. This is most often not the case in practice, because scales of different sources (instruments, methods) are also subject to uncertainty and a general systematic bias. This idea has been taken up in the development of advanced regression techniques for calibration (e.g., [32, 33]). The technique is most often called generalised least squares regression (although other brand names are in use), and is based upon weighted regression in both the dependent and the independent variables. In 2001, the technique was standardised for the field of gas analysis [34], its application is, however, generic and not field-specific.

2D averaging serves the same needs. It does not use weighted regression (which is able to cope with heteroscedastic datasets) but attributes common biases to both variables. In this sense, it is similar to warping, an emerging technique used for improvement of reproducibility of chromatographic measurements [35]. 2D averaging comprises the following steps.

Step 1: Each sampling table included in the optimization procedure is made ‘continuous’ by an interpolation rule which retains the experimental sampling points in space, and allows continuous interpolation between them. Rules may be as sophisticated as needed (polynomials, splines, rationales); frequency polygons (i.e., straight-line interpolations between sampling points) are often sufficient.

Step 2: The original data are adjusted according to appropriate bias models which should be as simple as possible, and as complicated as needed.

An appropriate model could be scaling for the responses (y -axis) which takes up biases in sensitivity of the different instruments or methods, and shifting the x axis, which accounts for additive biases between scales.

$$\begin{aligned} \varphi(x_{jk}) &= x_{jk} + \delta_k & : & \quad x_{jk} \rightarrow x_{jk} + \delta_k \\ y_{jk} &\rightarrow f_k \cdot y_{jk} \end{aligned} \quad (35)$$

Step 3: At any of the adjustment steps, each of the single measurement results from the sampling tables is attributed to, and grouped around the closest actual position in the independent variable. The estimated point on the reference function RF is calculated as the corresponding mean in both the x - and the y -direction according to

$$\overline{f_k \cdot y_{jk}} = RF_j(\overline{x_{jk} + \delta_k}) \quad (36)$$

and the RF represented by a frequency polygon interpolating all points calculated in this simple way.

Step 4 : The individual deviations of each point in the sampling table from the reference value according to Equation (36) at the actual position, and the total SSD (i.e. the sum over all these deviations) are estimated.

Step 5 : The total SSD is minimized by adjusting the parameters f_k and δ_k using an appropriate iteration procedure. Steps 2 to 4 are repeated until convergence is reached.

When convergence is reached, the best-fit estimates for f_k and δ_k are found. Now the joint confidence region (JCR) for each of the points which make up the RF (and were calculated in accordance to Equation (36)) is determined. Upper and lower CI of the points on the RF are then estimated as the points where the bisecting line of the RF frequency polygon passes through the JCR. Thus, the approach allows a joint estimation of both the best estimate for the RF and the biases which can reasonably be attributed to the data sources.

3 Assessment and combination of qualitative results

Although calculation, assessment, and statement of measurement uncertainty in *quantitative* testing are required by international standards [36], supervised by accreditation bodies, and guided by basic documents (e.g., [20–22]), the problem of assigning uncertainty-like statements to *qualitative* (also called *nonquantitative* to avoid confusion with ‘quality’) test results remains widely unsolved. The *ILAC Guide* [37] states that for qualitative testing, consideration is still being given as to how uncertainty of measurement applies in such cases. To overcome this situation, the European laboratory association EURACHEM formed a working group for drafting corresponding guidelines, and recently also ISO/REMCO initiated a study into this issue. This is motivated

by the fact that the volume of not purely quantitative testing and analysis is immense, and so is the economic impact of decisions taken on the basis of qualitative results.

3.1 Approaches and tools

Very much in general, basic approaches such as method repeatability and method reproducibility also remain valid for semi-qualitative test procedures. Any testing laboratory may create an estimate of its internal method dispersion simply by assessing the results of an appropriate number of repetitions. Such experiments can be carried out under either repeatability (controlled, constant environmental and other influential parameters, same operator) or reproducibility conditions (maximum variation of all influential parameters within the limits of method specification, and different operators).

Laboratory intercomparisons may be designed and executed in the very same way. Average lab precision and average *between*laboratories bias may be calculated from the intercomparison results using established algorithms.

Testing with measurements involved

Nowadays a very large group of qualitative testing procedures are not as qualitative as it may seem at a first glance. Rare are the situations where testing procedures are purely descriptive, or solely depend on human senses. Most frequently, one will have a situation where a complex characteristic (e.g., of a material) discontinuously changes in dependence on one or more recordable, continuous quantitative parameters. Two examples are material breakdown or detonation of explosive mixtures. Whether the characteristic has changed at a certain set of parameter values is assessed by judgement (e.g., pattern recognition, colour change) which may be instrument-assisted. Such procedures are semi-qualitative in the sense that the judgement is qualitative, but the result of the test is nevertheless a fully quantified (set of) parameter(s). A certain fuzziness of judgement induces an uncertainty in quantification. The *change* (or the status) of the characteristic to be tested is modelled as a Bernoulli random variable depending on a set of fully quantified parameters

$$\zeta = f(k_1, k_2, \dots, k_n). \quad (37)$$

The variable may take only two values: 0 (fail = status has not changed) or 1 (success = status has changed). During the test, the status of ζ in dependence on \mathbf{k} is assessed by judgement, and \mathbf{k}_c determined as the value at which the status of ζ changes. Normally, there will be a region of certainty where the outcome of the experiment is undoubtedly 0 or 1, and a region in the vicinity of the threshold \mathbf{k}_c where the outcome is uncertain, mainly due to the fuzziness of judgement.

Within this region, replicate measurements at each set of parameter values are taken, and the parameter values are increased (or decreased) in a stepwise

manner according to the actual test requirements. With n replicates, ζ at each sampling point should follow a $B(n, p)$ binomial distribution with known mean and variance. The (changing) probability parameter p is estimated as the frequency of occurrence (or nonoccurrence) of the event as observed, and the corresponding variances $\text{var}(p)$ are calculated accordingly.

A unit step function (with the step at \mathbf{k}_c) is fitted to the frequency-of-occurrence data. The variable parameter for the fit is \mathbf{k}_c . Different minimisation criteria are feasible, namely the following.

1. The normal SSD = min criterion known from regressions. Due to the discrete spacing of the sampling and the infinitive first derivative of the Heaviside function at \mathbf{k}_c , this criterion provides multiple solutions in an interval between two sampling points.
2. The criterion of equilibrated (i.e., equal) cumulated probabilities (quantiles) on both sides of the step. This criterion provides a unique solution for \mathbf{k}_c , which will fall in the interval according to (1) and, thus, also satisfies the minimum residual SSD requirement of an ordinary least squares regression (OLS) fit.

The criterion according to (2) should be preferred. Because at both sides of the step estimates of an average occurrence probability are only available at the discrete sampling points, quantiles may only be calculated with an additional assumption on an appropriate envelope function. Detailed information on the envelope is rather unavailable, but pragmatic approaches can be used including simple frequency polygons which connect the measured datapoints in a segmentwise manner by straight lines. A disadvantage of this pragmatism is that the areas beneath and above the frequency polygon will normally be different (i.e., the strict $p = 1 - q$ rule is violated), but because one will use in practice only one type of the areas (the cumulated p) on both sides of the step, the method will provide a sensible estimate for the equality point.

It may also seem reasonable to fit the quantile of the normal distribution to the data obtained for p . One aims at finding estimates for the variability (scatter) of one (or a couple of) value(s) \mathbf{k}_c , which is/are influenced by a considerable number of influential factors such that an overall distribution close to the normal distribution can be assumed. On the other hand, this coincides with the philosophy of the GUM [1], which also recommends the transformation of different, most often only assumed distributions to Gauss distributions, and the corresponding measures of scatter to standard uncertainties.

For the given set of datapoints $\{p_i, \mathbf{k}_i\}$ one has to solve the minimisation problem

$$\sum_{i=1}^N (p_i - Q(\mathbf{k}_i, \mathbf{k}_c, s))^2 = \min \tag{38}$$

(with Q being the quantile of the normal distribution) with respect to the parameters \mathbf{k}_c and s . This automatically generates the standard uncertainty; the confidence interval is calculated from its multiplication by the corresponding

(one-sided) t factor. The above-described approach may be further refined into a weighted (in y) regression using the inverse variances of the p_i as weights.

With or without such weighing in y , one is still not completely done because in a last step, the quantifiable uncertainties of the k_i should be taken into account by uncertainty propagation [34]. Following this propagation scheme, each k_i is separately altered by $\pm 1/2 u(k_i)$, and the minimisation problem solved for \mathbf{k}_c with a fixed s . The contribution of the uncertainty in \mathbf{k}_i to the total uncertainty is obtained as

$$u(\mathbf{k}_c : k_i) = \mathbf{k}_c(k_i + 1/2u(k_i)) - \mathbf{k}_c(k_i - 1/2u(k_i)) \quad (39)$$

and will accordingly be summed up for all sampling points.

Identification

In a generic sense, identification and recognition (of something) are not quite the same but similar terms. This subsection deals with identification in a more specific sense, namely of a chemical substance or compound. For analytical identification of a substance, measurements are also carried out, most often using spectroscopic techniques such as UV/Vis, NIR, IR, MS, or NMR. The results of the measurement are spectra, and identification basically reduces to a comparison of known and unknown spectral pattern.

Figure 3 illustrates that identification problems (which are usually considered qualitative) have much in common with quantitative analysis if the subsequent decision process is taken on board. In quantitative analysis, decisions on the compliance of the data with certain specifications are made on the basis of the value and the uncertainty, the latter taken as a measure for the performance of the method used. This is one of the reasons why this chapter deals with assessment strategies providing reliable values and uncertainties. From the outcome of the decision process, statistics can be generated which provide estimates for the probabilities that

- A complying object is recognised as such (true positive).
- A noncomplying object is recognised as such (true negative).
- A noncomplying object is admitted as complying (false positive).
- A complying object is rejected as non-complying (false negative).

Although measurements are involved, the results of an identification (pattern recognition) experiment are treated differently. They are compared with certain specifications regardless of any performance criteria. Where one finds measurement uncertainty in quantitative analysis, there is a question mark here.

Instead, the statistics on the true positive, true negative, false positive, and false negative outcomes are normally considered as the uncertainty estimates. This widely leaves out the performance characteristics (reproducibility, biases) of the concrete equipment used for doing the test, at the concrete time the

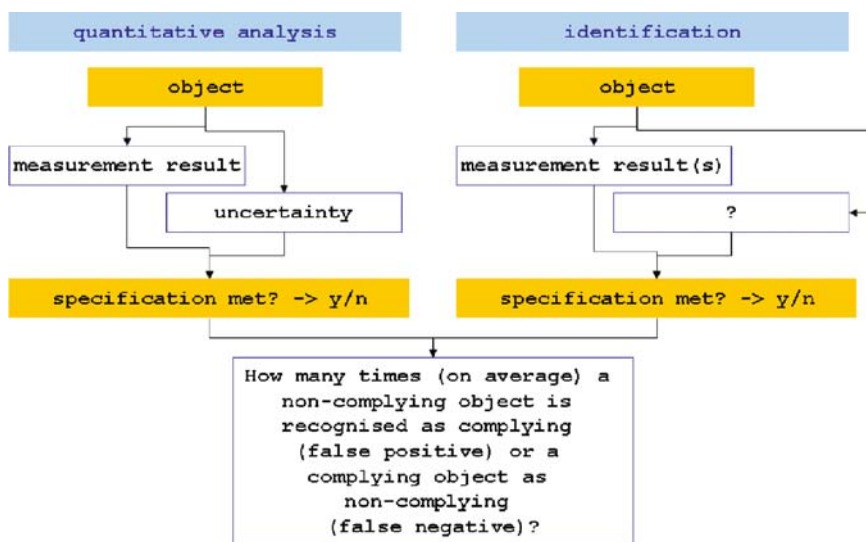


Fig. 3. Analogies between quantitative testing (including decision-making on the outcome) and substance identification as an example of qualitative testing.

test is carried out. An obvious candidate for filling the box with the question mark would be principal component (PCA) or discriminant analysis (DA), the latter if maximum distinction between patterns is the aim. PCA has successfully been used for discrimination of composition patterns of a large variety of food products (wines, high-alcoholic beverages, olive oils, honey, etc.). Characteristic patterns were then attributed to the origin of the product or the raw materials from which these products were made, for example, [38–40].

The result of a (spectroscopic) measurement for substance identification is a data table $\mathbf{y}|\mathbf{x}_k$ containing signals \mathbf{y} (peak heights or areas, extinctions, emission densities, abundances, etc.) at well-separated or even quasi-continuous sampling positions \mathbf{x} (wavelengths, retention times, masses, ratios, etc.). The length of the table may vary widely depending on the method. Measurements are made in replicate for at least the reference substance and the unknown. Both tables are normalised, preferably with respect to the sums of all signals involved. This normalisation removes the influence of sensitivity and/or injection variations of the equipment. Normalised data for the reference substance are subjected to PCA, and reference loadings are obtained. These are applied to the normalised data for the unknown to give scores which are displayed together with the scores for the reference. Joint confidence regions for both score sets are calculated and serve as the basis for assessing the overlap and the difference between the unknown and the reference. This is done in the very same manner as scalar uncertainties are used for deciding about nonsignificant or significant differences between scalar values. The JCR for the reference substance is centred in the origin of the score plot. The

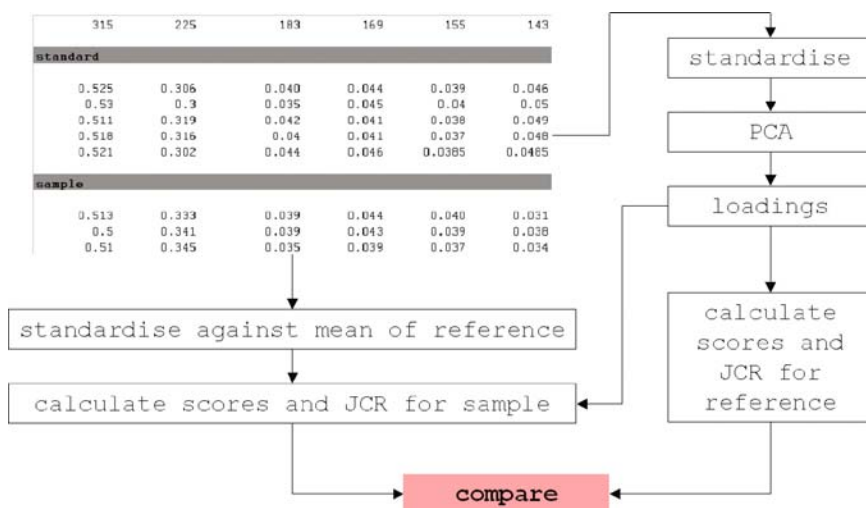


Fig. 4. Data treatment for comparison of an unknown with a reference substance. Starting point is a data table with replicate measurements for the reference and the sample; here abundances of spectral mass peaks are determined by MS/MS at 6 peak positions.

minimum identity requirement is that the JCR for the scores of the unknown also contain the origin of the plot. More sophisticated models which assess the degree of overlap between both JCRs are possible. The described data treatment is visualised in Figure 4.

Inspection

Inspection, either by visual or sensual judgment or supported by auxiliary measurement, follows the principles of attribute sampling. Quite often, mutually excluding nominal properties such as intact/broken, tight/leaking are considered, and modelled using the formal approach introduced at the beginning of this section. Occurrence of the ‘event’ follows a binomial or Poisson distribution. Using these distributions, a number of problems of practical interest can be solved, such as the determination of optimum sample size given a certain acceptance or rejection criterion. Closely linked to production process control, available tools are well developed, and laid down in several international standards [41–43]. Note that, for example, ISO 2859 (reference [41]) has several parts, with part 1 dealing with lot-by-lot inspection, part 2 with isolated lot inspection, and parts 3 to 5 with derived techniques including skip-lot and sequential sampling.

But even if the testing procedure involves supporting measurement, equipment uncertainty (implementation effects, reproducibility, biases) is normally disregarded (see Figure 3). A certain fuzziness of the equipment is, however,

recognised and taken into account as a probability of not detecting the property if present, and detecting it in absence. Thus method performance is characterised by the positive/negative decision rates described above. Bayes' theorem is applied to the conditional probabilities in order to derive risk estimates for acceptance or rejection given the feasible sample sizes and sampling conditions.

Characteristics derived from the basic probabilities are often used in non-destructive testing. Nondestructive testing aims primarily at the detection of defects and faults in components and assemblies. The diagnostic signals may be quantitative (e.g., the peak height of an ultrasound echo) or qualitative (e.g., the blackening pattern on an X-ray film). The quality of the testing procedure is assessed from a receiver operating characteristic (ROC) diagram in which the probability of detection (POD) of a defect is plotted against the probability of false alarm (PFA). Blind testing of the equipment under realistic conditions in the industrial environment allows an integral assessment of the system capabilities given the application conditions and the human factor. A comprehensive description and application to land mine detection and de-mining is given in [44].

3.2 Uncertainty estimation

In addition to the above-mentioned approaches to the description of equipment capabilities, general principles of stating and reporting uncertainty in qualitative testing have been developed in [45].

The use of uncertainty information in compliance assessment is comprehensively described in [46]. Metrological requirements to test equipment stipulated in the new EU Measurement Instrument Directive caused attempts of incorporating measurement uncertainty in the decision-making process in conformity assessment when sampling is limited. Risks of incorrect decisions where test results lie in the vicinity of a specification limit are assessed in terms of percentage probability and the costs of measurement and environmental consequences [47], leading to an optimised uncertainty methodology based on attribute sampling.

4 Some final remarks

Any data assessment and consolidation strategy should start from thorough data exploration aiming at the detection of possible bias. Plenty of well-elaborated tools and approaches exist; some of them have been described here. Which tools are appropriate depends on the suspected sources of bias. Uncertainty budgets which are as complete as possible should be the basis for exploration. Bare repeatabilities are normally insufficient, but in some cases the only estimates attainable so far.

The outcome of data exploration suggests the appropriate way of combining results. Again, plenty of approaches exist, and the most appropriate should be selected. Most often, for specific problems both a tailored analysis and a tailored combination strategy will be needed.

A data consolidation strategy is successful if individual data are compatible with the final result within the stated uncertainties (assessed by, e.g., the E_n criterion, the chi-square technique, or the Birge ratio).

Multianalyte analyses and functionally related properties require multivariate data exploration tools. Efforts should be made to enable existing multivariate tools to take up, and account for, type B uncertainty estimates. This would make them fully GUM-compliant and bridge another gap between quantitative and qualitative testing.

References

1. BIPM, IEC, IFCC, ISO, IUPAC, IUPAP and OIML *Guide to the Expression of Uncertainty in Measurement*, ISO, Geneva, 1st edition (reprinted with corrections) 1995
2. H. W. Coleman, W. G. Steele: *Experimentation and Uncertainty Analysis for Engineers*, John Wiley & Sons., Inc., New York, 1989
3. M. Neuilly: *Modelling and Estimation of Measurement Errors*. Intercept Ltd., Hampshire, UK, 1999
4. L.J. Lyons, *Analyst*, 2000, **125**, 2020–2025
5. M.S. Levenson, D.L. Banks, K.R. Eberhart, L.M. Gill, W.F. Guthrie, H.K. Liu, M.G. Vangel, J.H. Yen, N.F. Zhang, *J. Res. Natl. Inst. Stand. Technol.*, 2000, **105**, 571–579
6. ISO Guide 35, *Reference materials—General and statistical principles for certification* (3rd edition 2006)
7. W.G. Cochran, *Biometrics*, 1954, **10**, 101–129
8. R.J. Douglas and A.G. Steele, *Metrologia*, 2006, **43**, 89–97
9. C.R. Paule and J. Mandel, *J Res Natl Bur Stand*, 1982, **87**, 377–385
10. S.B. Schiller and K.R. Eberhardt, *Spectrochimica Acta*, 1991, **46B**, 1607–1613
11. A.L. Rukhin and M.G. Vangel, *J Am Stat Assoc*, 1998, **93**, 303–308
12. W. Hässelbarth, W. Bremser and R. Pradel, *Fresenius J Anal Chem*, 1998, **360**, 317–321
13. R. Kacker, *Metrologia*, 2004, **41**, 132–136
14. R.T. Birge, *Phys Rev*, 1932, **40**, 207–227
15. W. Weise and W. Wöger, *Meas Sci Technol*, 2000, **11**, 1649–1658
16. R. Kacker, R. Datla and A. Parr, *Metrologia*, 2002, **39**, 279–293
17. J.W. Müller, *J Res NIST*, 2000, **105**, 551–555
18. W. Hässelbarth, and W. Bremser: *Combination of results from several measurements – An everlasting problem*. in: *Combining and Reporting Analytical Results*. The Royal Society of Chemistry 2007. RSC Publishing, Cambridge, UK.
19. M.G. Cox: A discussion of approaches for determining a reference value in the analysis of key-comparison data. NPL Report CISE 42/99

20. EURACHEM/CITAC Guide: *Quantifying Uncertainty in Analytical Measurement*, 2nd edition, 2000
21. ISO 5725-parts 1, 2, 6:1994, *Accuracy (Trueness and Precision) of Measurement Methods and Results*. ISO, Geneva
22. V.J. Barwick, S.L.R. Ellison, *Analyst*, 1999, **124**, 981–990
23. Nordtest TR 537(2003): *Handbook for Calculation of Measurement Uncertainty in Environmental Laboratories*
24. Codex alimentarius Commission: *Guidance on Measurement Uncertainty* (March 2007)
25. Eurolab Technical Report 1/2007: *Measurement Uncertainty Revisited: Alternative Approaches to Uncertainty Evaluation*
26. H.Scholze, *Der Einbau von Wasser in Gläser, Glastechnische Berichte*, 1959, **32**, 81–88
27. A.S. Brown, M.J.T. Milton, C.J. Cowper, G.D. Squire, W. Bremser, R.W. Branch, *J Chromatog A*, 1040, 2004, 215–225
28. Press WH, Teukolsky SA, Vetterling WT, Flannery BP, *Numerical recipes in Fortran*, Cambridge University Press, 2nd edition, 1994
29. G. Engeln-Müllges, F. Reuter, *Numerik-Algorithmen mit Fortran77 Programmen*, Wissenschaftsverlag, Mannheim-Leipzig-Wien-Zürich, 1993
30. W. Hässelbarth, W. Bremser, *Metrologia* **44** 2007 128–145
31. W. Bremser, W. Hässelbarth, *Accred Qual Assur* 1998 **3** 106–110
32. W.E. Deming, *Statistical Adjustment of Data*, John Wiley & Sons, New York, 1946
33. B.D. Ripley, M. Thompson, *Analyst*, 1987, **112**, 377
34. ISO 6143:2001, *Gas analysis—Comparison methods for determining and checking the composition of calibration gas mixtures*, ISO, Geneva
35. G. Tomasi, F. van den Berg, and C. Andersson, *J. Chemometrics*, 2004, **18**, 231–241
36. ISO/IEC 17025:1999, *General Requirements for the Competence of Testing and Calibration Laboratories*. ISO, Geneva
37. ILAC-G17:2002, *Introducing the Concept of Uncertainty of Measurement in Testing in Association with the Application of the Standard ISO/IEC 17025*
38. W.-O. Kwan, B.R. Kowalski, and R.K. Skogerboe, *J Agricult Food Chem*, 1979, **27**, 1321–1326
39. V.F. Taylor, H.P. Longrich, and J.D. Greenough, *J Agricult Food Chem*, 2003, **51**, 856–860
40. X. Capron, J. Smeyers-Verbeke, and D.L. Massart, *Food Chem*, 2007, **101**, 1585–1597
41. ISO 2859-1:2001, *Sampling Procedures for Inspection by Attributes —Part 1: Sampling Schemes Indexed by Acceptance Quality Limit (AQL) for Lot-by-Lot Inspection*, ISO, Geneva
42. ISO 3951:2005, *Sampling Procedures for Inspection by Variables, Parts 1–5*, ISO, Geneva
43. ISO 14253-1:1998, *Geometrical Product Specifications (GPS) —Inspection by measurement of workpieces and measuring equipment —Part 1: Decision rules for proving conformance or non-conformance with specifications*, ISO, Geneva
44. C. Mueller, M. Gaal, M. Scharmach, U. Ewert, A. Lewis, T. Bloodworth, P.-T. Wilrich, and D. Guelle, *Reliability Model for Test and Evaluation of Metal Detectors*. Final report of the ITEP project 2.1.1.2, Sep 2004. Available at <http://www.itep.ws/>

45. S.L.R. Ellison, *Accred Qual Assur*, 2000, **5**, 346–348
46. EURACHEM/CITAC *Guide: Use of Uncertainty Information in Compliance Assessment*, 1st edition, 2007
47. L.R. Pendrill, H. Källgren, *Accred Qual Assur*, 2006, **11**, 496–505

Monte Carlo Modeling of Randomness

Alan G. Steele, Robert J. Douglas

Institute for National Measurement
Standards, National Research Council of Canada, Ottawa, Canada K1A 0R6
alan.steele@nrc-cnrc.gc.ca, rob.douglas@nrc-cnrc.gc.ca

Summary. The evaluation of expressions of measurement uncertainty can be enhanced by Monte Carlo simulation techniques. For any claimed uncertainties that have been expressed (either explicitly or implicitly) as probability density functions, Monte Carlo techniques can propagate uncertainties through any measurement equation or algorithm – including statistical aggregates. These techniques are introduced in this chapter, and applied to examples drawn from measurement comparisons. Examples illustrate the power of the techniques in extending uncertainty analysis to cases where the applicability of the usual methods is in doubt, such as consistency testing using chi-squared-like statistical aggregates.

Key words: Measurement, uncertainty, Monte Carlo, simulation, comparison, measurement equation, measurand, random, GUM

1 Introduction

In the context of uncertainty, Monte Carlo modeling is based on an extremely simple concept. Experience shows that measurements do not repeat exactly, but have a random component, as is illustrated in Figure 1.

A Monte Carlo computer model can replicate the randomness in the measurement process, creating pseudo-measurements that incorporate the modeled randomness, as shown in Figure 2. These pseudo-measurements may be used instead of real measurements in support of conventional uncertainty analyses [GUM95]. This chapter illustrates how modeled randomness can emulate uncertainties evaluated not only by statistics (Type A methods) but also by other means (Type B methods).

Monte Carlo modeling of randomness provides an easy means for combining uncertainties [GUM-S1 07] that is particularly useful when the adequacy of the usual methods are in doubt.

It is very simple to combine uncertainties using Monte Carlo simulation:

- Model the uncertainties of the input quantities with pseudo-random number generators.




Fig. 1. Repeated measurement of a 25 ohm platinum resistance thermometer in an ice bath over a time interval of about five hours. To the right is a histogram of the measurements.




Fig. 2. Monte Carlo simulation of the measurements of Figure 1. To the right is a histogram of the simulation (gray points) compared to a histogram of the measurements (black points).

- Model the calculation of the output quantity from the input quantities using the measurement equation.
- In a loop, generate a set of pseudo-random input quantities, calculate its output quantity, increment a histogram bin—repeat loop many times (commonly $\geq 10^5$).

Given a means of generating the appropriate pseudo-random numbers, notice that there is nothing complicated here. Each iteration through the loop represents a single Monte Carlo event, where the instance values of the input quantities have been sampled from their attributed distributions. A picture of the resulting randomness associated with the output quantity is built, one event at a time, by performing the calculation many times and generating a histogram of the simulated events.

Even the choice of modeling strategy is only rarely a complicating issue. For any scalar input quantity, subtleties arise only when its uncertainty is nonsymmetric about that quantity's reported value [DSWH05], so usually no distinction needs to be made between the distribution of

measurements (or *gedanken* measurements) of Figures 1 and 2 and the distribution of the measurand (to which formal uncertainty distributions refer). The formal justification for this is facilitated by the standard practice by metrologists of using the *same value* to be the best representation both of the (fully corrected) measurement, and of the measurand. This “fiducial value” simplifies, and in our view strengthens, the fiducial argument [F35].

The utilitarian potential of a measurement can be realized directly when it is compared to one or more other measurements. For this chapter we have chosen measurement comparisons as examples to illustrate the power of Monte Carlo simulation in the chains of metrological inference that provide the practical benefits of measurement science.

1.1 Strengths of Monte Carlo methods

With the Monte Carlo simulation method of uncertainty analysis, there is no restriction on the measurement equation. Neither linearization nor derivative-evaluation need to be done. The Monte Carlo method will work every bit as easily for a measurement algorithm as it does for a measurement equation. Thus, unlike the conventional methods that *require* a measurement equation, Monte Carlo simulation can be applied to uncertainty analysis of measurement processes that incorporate decisions or branching, or incorporate order statistics that depend on sorting a set of values into nondecreasing order.

For any case where there is some doubt about the adequacy of the approximations inherent in the classical GUM analysis, Monte Carlo simulation is a useful tool to investigate, and usually to confirm, that the classical methods for combining uncertainties are indeed adequate. For all the power of Monte Carlo simulation, it is also important to recognize its limitations: although it can easily and quickly give answers to obscure questions about the model, the simulation in and of itself does not address questions about reality.

Monte Carlo simulation is applicable to uncertainty claims that have been derived by *any* method and expressed as a standard uncertainty, or as a standard uncertainty with a stated degrees of freedom, or as an explicitly given probability density function. In this chapter, we restrict the discussion to real, scalar input, and output quantities, but address all other intricacies.

The power of the Monte Carlo method is further revealed when applied to aggregates of extensive sets of experimental measurements to judge the mutual consistency of such results relative to the uncertainty claims. It can analyze the predictions implied by these claims however they have been prepared, whether by Type A methods, frequentist methods, Type B methods, Bayesian methods, Monte Carlo simulation, or quantum mechanics. By construing a prediction from the uncertainty claim associated

with each measured value, it is possible to invoke the classical scientific method of comparing prediction with experiment to create and to convey confidence for the methods of uncertainty analysis used in a particular application.

2 Distributions for uncertainty claims

x	An input quantity
$f_G(x)$	Gaussian probability density function (PDF) for x
a_G	Normalizing constant for f_G
u	Standard uncertainty of input quantity x
x_0	Mean, or reported value, of x
$f_S(x)$	Scaled and shifted Student- t PDF
a_S	Normalizing constant for Student- t PDF
ν	Degrees of freedom for Student- t PDF
a	Semi-width of a rectangular distribution

The recommended [GUM95] form of an uncertainty claim is as a standard uncertainty, taken as the standard deviation of a Gaussian distribution, or as the *sample* standard deviation (see Section 2.2) of a t -Student distribution if a finite degrees of freedom is given, or as the distribution standard deviation if there is any other significant departure from a Gaussian, with its departure described explicitly. This is sometimes done using moments, but it is usually better to do this using tabulated distributions. Any asymmetry will have to be correctly oriented, using an appreciation of whether the distribution describes the measurand, or is a more direct prediction of measurement results.

2.1 Gaussian or Normal distributions

The most common function used to provide a basis for uncertainty statements is the Gaussian or normal probability density function (PDF). A random variable x is used to represent the possible value of the measurand, described by its mean value x_0 and its standard deviation u . If only x_0 and u are given to describe a measurand, a Gaussian PDF is to be construed for x :

$$f_G(x) = a_G \times e^{-(x-x_0)^2/(2u^2)}, \quad (1)$$

where a_G is a normalizing constant, equal to $(u\sqrt{2\pi})^{-1}$, chosen so that the integral of f_G over all x , the overall probability, will be equal to 1.

2.2 Student- t distributions

Another function that is explicitly recognized as the basis for expressing uncertainty [GUM95] is the scaled and shifted Student- t distribution. This has an additional parameter, the *degrees of freedom*, ν , and in the limit as $\nu \rightarrow \infty$ the Student distribution converges to the Gaussian distribution discussed above. The scaled and shifted Student- t PDF is given by

$$f_S(x) = a_S \times [1 + (x - x_0)^2/(\nu u^2)]^{-(\nu+1)/2}, \quad (2)$$

where a_S is a normalizing constant, equal to $\Gamma([\nu + 1]/2)/[\Gamma(\nu/2) u\sqrt{\pi\nu}]$, chosen so that the integral of f_S over all x , the overall probability, will be equal to 1. Although tabulated most commonly for ν as a positive integer, ν does not have to be an integer when expressing uncertainty. This distribution is centered at x_0 and has a standard uncertainty u , the *sample* standard deviation, which for all $\nu > 2$ is $\sqrt{(\nu - 2)/\nu}$ times the distribution's standard deviation (the square root of the distribution's variance about x_0).

The *sample* standard deviation is the parameter u in Equation (2), and is to be used as the standard uncertainty [GUM95] whether it is derived as the standard deviation away from the mean from a sample of $\nu + 1$ independent measurements drawn from a Gaussian, or if it is derived in any other way (including a consideration of how well the uncertainty is known). If an uncertainty budget has been expressed using the guidance of the ISO-GUM, and a degrees of freedom has been quoted, then the standard uncertainty *must* be construed as the parameter u in Equation (2) to accord with the intent of the authors of any uncertainty budget who have followed the ISO-GUM's development of expanded uncertainty (Annex G of [GUM95]) and its use of the Welch-Satterthwaite approximation. Note that this required usage is contrary to the usage of the standard uncertainty as the *distribution* standard deviation for a general PDF, and a subordinate suggestion [GUM-S1 07] for scaled and shifted Student- t distributions. Our perspective is that supplements to the ISO-GUM should not be interpreted as modifying the ISO-GUM, and the way that is most obviously faithful for construing an ISO-GUM standard uncertainty associated with a degrees of freedom is to use Equation (2) with u as the quoted standard uncertainty and with ν as the quoted degrees of freedom.

There are different reasons why those preparing an uncertainty statement might give a finite degrees of freedom. Each reason can be appropriate in its own context, but the reasoning need not concern us here. Whenever a standard uncertainty claim is accompanied by a finite degrees of freedom, in the absence of other information users must construe a scaled and shifted Student- t PDF in order to be consistent with the accepted standard [GUM95] and its explicit calculations of expanded uncertainty.

2.3 Rectangular distribution

The rectangular or uniform distribution is constant over a range $(x_0 - a, x_0 + a)$, and has zero probability density outside this range. It is commonly used for uncertainties where only the bounds of an input quantity can be specified. It also has a more rigorously justified use in accounting for the quantization uncertainty between a finite-resolution measurement and an analog measurand.

2.4 U-shaped distribution

For an input quantity that is sampled at random in a situation there is a sinusoidally varying fluctuation between two fixed limits, a particular U-shaped PDF—the arcsine distribution—may be useful. It may be appropriate where an on-off temperature controller is cycling on and off between two fixed limits in a system that has a time constant much longer than the controller cycle time. We discuss this distribution in more detail below, as a worked example in the section on algorithms.

3 General distributions

The building block PDFs discussed above do not cover all possible cases of interest for uncertainty analysis. It is possible to list other forms, but eventually circumstances arise that cannot be covered by any list, however long the list may be. For example, for an input quantity needed in your measurement equation, you may be handed an uncertainty claim for an input quantity that includes a description of a wholly new PDF.

Although this is not now a common occurrence, it is not fanciful. It will likely become more common as we improve our ability to describe the uncertainty properly and so meet the requirements of the international standard [GUM95]. If, in the opinion of the person preparing an uncertainty statement, there are significant departures of the PDF from the Gaussian distribution, or from the Student- t distributions, such departures are to be described in stating the uncertainty. There are several ways in which this may be done.

3.1 General distributions by moments

One traditional way of describing a more general PDF is to use its moments. Indeed, we have already encountered the first moment (the mean value, x_0) and the second moment about the mean (the variance, u^2 , or for the special case of the Student- t distribution $u^2\nu/(\nu - 2)$). For higher moments, the n th moment M_n of a PDF $f(x)$, about its mean x_0 is

$$M_n = \int_{-\infty}^{\infty} (x - x_0)^n f(x) dx. \quad (3)$$

The second moment is the variance (usually u^2), and the third moment may be expressed by a coefficient of skewness, usually $M_3/M_2^{3/2}$. If a PDF has a more pronounced tail to the right, its skewness is positive, and if the tail is more pronounced to the left the skewness is negative. A distribution that is symmetric about x_0 has zero skewness, and M_n is zero for any n that is odd.

For any uncertainty that is not symmetric about x_0 , it is important to appreciate and to be able to discuss the sense of the PDF and whether it “tails left” or “tails right”. The skewness parameter is a convenient mechanism for keeping track of the orientation of the distribution in such discussions.

Even a Gaussian PDF has a fourth moment, or kurtosis, which is equal to 3. We often encounter the “excess kurtosis” parameter $(M_4/M_2^2) - 3$, which is zero for a Gaussian and which thus describes a significant departure in peakedness from a Gaussian PDF.

We believe that describing general distributions by their moments is of only limited use for generalizing Monte Carlo methods to handle arbitrary probability distributions. However the moments are reported, the uncertainty statement is only complete when a scheme for expanding the moments back into the PDF is also specified by the persons preparing the uncertainty statement, otherwise their uncertainty specification is not unique. With the expansion specified, it would be straightforward to calculate and to use a tabulated approximation of the PDF, as discussed next.

3.2 General distributions by tables

Another traditional method for describing a function is to tabulate x and $f(x)$ for a reasonable number of points and to specify an interpolation scheme for calculating $f(x)$ for any value of x in between the tabulated values. Every time a PDF is graphed, this is the basic method that used, and this method can be adapted to describe the PDFs of any claimed distribution in a form that may be used easily in Monte Carlo simulation of the uncertainty claims.

The general form of tabulated function and interpolation scheme can be quite simple. Consider a table of N values of x , $\{x_i\}$, $i = 1, \dots, N$, sorted in ascending order so that $x_i < x_{i+1}$. For each x_i there is a corresponding $f(x_i)$, with $f(x_1) = 0$ and $f(x_N) = 0$. For any value of x outside the range (x_1, x_N) , $f(x) = 0$. For any value of x between x_i and x_{i+1} , the value of $f(x)$ can be calculated by linear interpolation:

$$f(x) = f(x_i) + [f(x_{i+1}) - f(x_i)] \times [x - x_i]/[x_{i+1} - x_i]. \quad (4)$$

Other, higher-order, interpolation schemes are also possible. The points do not need to be equally spaced, so it is feasible to approximate a discontinuous function very efficiently. For example, a rectangular distribution could be approximated well by a table of length $N = 4$, having a flat top and arbitrarily steep sides. Note that $f(x)$ is the probability *density*, so that when the $\{x_i\}$

are unequally spaced, the tabulated $f(x_i)$ are not all simply proportional to the probability (they can all have different proportionality constants).

Monte Carlo simulation can treat any uncertainty expressed as a general function that can be tabulated and interpolated in this way.

3.3 Skewed general distributions

In the authors' opinion, one of the benefits of Monte Carlo simulation for analysis of uncertainties is that there can be a direct mapping of intermediate results of the simulation onto possible experiments where repeated real measurements and repeated simulations could be compared. There is a minor difficulty to be addressed: the experimental data are expressed in the form of a histogram of replicated measurements, and the simulated uncertainty distributions should be in the form of a histogram of replicated simulated values for the measurand if the uncertainty is being expressed in the standard way [GUM95].

For the scalar measurands being discussed here, there is a simple and rather robust mechanism for converting one sort of histogram into the other: by reflecting the histogram of simulated values about x_0 [DSWH05]. For symmetric distributions, this is a null operation and the minor difficulty vanishes. For claimed uncertainty distributions that are not symmetric, care is needed to construe the meaning of the claimed distribution and to propagate its skewness appropriately through any changes from the measurand perspective to the measurement perspective and back again. Skewness reversal can also arise from the measurement equation (e.g., in subtracting), but this will be looked after automatically by the Monte Carlo simulation process.

4 Multivariate distributions

Although we have discussed the use of a general PDF of a real variable by tabulating and interpolating the PDF of any combined uncertainty, there are sometimes advantages in backing up one (or more) layers in order to use the constituent uncertainty components of a combined uncertainty statement as the mechanism for creating a Monte Carlo simulation of a new output quantity. In essentially redoing an uncertainty calculation authored by someone else—that is, in recombining their uncertainty components within your own uncertainty model, there can be no objections provided that there are no significant differences in the calculations.

The disassembly of a combined PDF into more elemental PDFs can sometimes really simplify the treatment of covariant uncertainties, and avoid assumptions and approximations that are often invoked when a multivariate PDF is used as the basis of handling covariances.

4.1 Covariance: Simplifications – Type A

x_1, x_2	Two input quantities, possibly covariant
u_1, u_2	Standard uncertainties of x_1, x_2
$r_{1,2}$	Correlation coefficient of x_1 with x_2
y_1, y_2, y_3	Three independent random variables to model uncertainty of x_1 and x_2

A covariance that has been evaluated by a purely Type A method (statistical evaluation of correlations between repeated *paired* measurements) can be handled by Cholesky decomposition to obtain an efficient means for a Monte Carlo resampling of the multivariate probability density function that has the desired variance and covariance characteristics.

We can obtain additional clarity by sacrificing a bit of efficiency, using more than the minimum number of pseudo-random variables in order to describe the multivariate Gaussian distribution. For example, the general bivariate Gaussian distribution of two (zero mean) input quantities x_1 and x_2 would usually be characterized by their standard uncertainties u_1 and u_2 , and their correlation coefficient $r_{1,2}$ (or covariance $r_{1,2} u_1 u_2$).

An equivalent description can be created using *three* independent zero-mean random variables, $y_1, y_2,$ and $y_3,$ drawn from Gaussian distributions having standard deviations $v_1 = \sqrt{u_1^2 - r_{1,2} u_1 u_2}, v_2 = \sqrt{u_2^2 - r_{1,2} u_1 u_2},$ and $v_3 = \sqrt{r_{1,2} u_1 u_2}.$ It is clearer to consider v_1 as the standard uncertainty of the independent perturbations of $x_1,$ and v_2 for $x_2,$ with v_3 representing the standard uncertainty of the perturbations shared between x_1 and $x_2.$ In the context of Monte Carlo simulation, it is a simple matter to generate a set of three values for $y_1, y_2,$ and $y_3,$ drawn from PDFs characterized by $v_1, v_2,$ and $v_3,$ and to calculate $x_1 = y_1 + y_3$ and $x_2 = y_2 + y_3$ for the event, where the *same* value of y_3 is used for calculating both x_1 and $x_2.$

For Type A evaluations of a covariant uncertainty component, generally no cause (or causes) are identified for this “shared” perturbation, $y_3.$ The Monte Carlo method can gracefully handle covariances evaluated by Type A methods, but does not help us to focus on obtaining clearer insights. In our experience with high-level metrology, it is rare to rely only on Type A methods for evaluating covariances in uncertainty budgets.

4.2 Covariance: Simplifications – Type B

For covariances evaluated by Type B methods, Monte Carlo methods can focus on the physical and practical meanings of our variables in evaluating covariances implicit in the input quantities used in our measurement equation.

In practice, two or more variables in a measurement equation can be examined for possible covariances by asking a few simple questions about their values as used in the measurement equation:

- Are the values the same? This is a particularly simple but rather weak test. Particularly for interlaboratory comparisons, identical corrections *may* indicate a common origin. The test is weak in two ways: it may be just a chance occurrence that the values are the same in this instance, and that there is no underlying correlation. Furthermore, even if corrections by different laboratories are not numerically identical, they may still be correlated.

- Are the values *procedurally constrained* to be the same? If two uncertainty budgets each import a value for one of their input quantities from the same source, the associated uncertainty component is completely correlated (with correlation coefficient $r = 1$). Although this system might be modeled as two correlated variables, simplicity dictates that the Monte Carlo model should use only one pseudo-random variable for this circumstance.

- Are the values *physically constrained* to be the same? For example, in a measurement equation, many parts of the associated measurement system may be affected by temperature, and although each part may be at a slightly different temperature, the temperatures of all parts are expected to track the laboratory temperature pretty closely. For this situation, a single pseudo-random variable (for laboratory temperature) is the simple starting point.

These questions are relatively simple for measurement specialists to answer in their own fields. The more complete question, “Are the values *procedurally or physically constrained* to be *partially* covariant?,” is a much more difficult question to answer. It may be possible to develop insights about these questions by reformulating the measurement equation in terms of independent and shared perturbations (similar to the y_1, y_2, y_3 basis discussed above).

4.3 Covariances in building measurement equations

In creating a Monte Carlo model, a sound objective is to use as few pseudo-random variables as possible, *but no fewer*. If simplicity can be enhanced by increasing the number of pseudo-random variables, we believe that it is appropriate to do so. In our opinion, Monte Carlo simulation is best exploited when one goes beyond considering a standardized measurement equation, and considers anew all of the assumptions that had to be made to arrive at a measurement equation that was tractable using standard methods. Because Monte Carlo simulation for combining uncertainty components can easily cope with many more intricacies than the traditional approach, it is a natural way to explore extensions to a measurement equation without incurring the overhead of re-evaluating derivatives and covariances.

4.4 Example: Covariance matrix to perturbation vector

α	Temperature coefficient
T_i	Temperature at which measurement i was taken
T_0	Reference temperature for reporting measurement
$u(\alpha)$	Standard uncertainty in α

To illustrate how covariances are automatically handled by the direct approach inherent in Monte Carlo model building, consider a Type B evaluation of covariance due to the temperature coefficient α of a measurand, with dominating uncertainty $u(\alpha)$. One example of this might be the thermal expansion of a gauge block, needing a correction of the measured length at a temperature T , to the specified reference temperature T_0 , of $-\alpha[T - T_0]$. To first-order, the correction's uncertainty has two components, each with its own standard uncertainty: $\alpha u(T - T_0)$ and $u(\alpha)[T - T_0]$. It is this last standard uncertainty, $u(\alpha)[T - T_0]$, that can give rise to covariant uncertainties if there is a shared value of the thermal expansion coefficient that is used for applying the correction to different measurements.

We require the full covariance matrix of a set of N' measurements made at known temperatures T_i , rather than at the specified reference temperature T_0 , and assume $u(T_i) \times \alpha \ll \alpha[T_i - T_0]$. With elements $u^2(\alpha)[T_i - T_0][T_j - T_0]$, this covariance matrix has a determinant of zero, so the usual methods for Cholesky factorization fail.

Considered from the approach of Monte Carlo simulation, however, the uncertainty in the i th result is just $u(\alpha)[T_i - T_0]$. A single random variable can simulate the uncertainty in α , scaled by $[T_i - T_0]$ to the N' simulated errors of the corrected measurand. Determining the scope of each value for α is part of the uncertainty modeling, determined by the physical constraints on the temperature coefficient and described by the PDF associated with $u(\alpha)$. Here, the temperature coefficient is to be represented by the same value for each set of simulated measurements, resampled to reflect its uncertainty $u(\alpha)$.

In the Monte Carlo approach, it is clear that the random variable for α can be drawn from *any* shape of distribution, in contrast to the usual multivariate distribution approach where it would have to be approximated by a Gaussian.

5 Pseudo-random number generators

The ascendance of Monte Carlo simulation is based on the easy availability of programmable computers that can quickly perform deterministic calculations involving millions—or millions of millions—of arithmetic operations. A particular class of deterministic program has been developed that can emulate randomness, producing a sequence of (floating point) values that are difficult to distinguish from ideally random values. Programs that do this are termed *pseudo-random number generators*. The “pseudo” is not referring to readily discernible departures from randomness in the generated sequence (many are practically impossible to distinguish from ideal randomness) but rather to the fact that the sequence is deterministic: it can be repeated at will, and so the usual validation procedures for deterministic software can be used. Even if true random number generators (i.e., nondeterministic) were available everywhere, the preference would be for *pseudo*-random number generators as a tool for

uncertainty analysis and measurement science and in any other application where software validation is an everyday concern.

For ordinary computers, a wide variety of pseudo-random number generators is available, and although a few of the earlier methods have serious flaws for uncertainty analysis, most are quite suitable for use with the Monte Carlo method of combining uncertainty components. For a few applications such as “six-nines” calculations, or for other chi-squared-like tests comparing probabilities at 10^{-8} or less, additional care is needed when selecting your generator. Also, it has happened that a broadly trusted application has been “updated” to break a formerly working pseudo-random number generator. Clearly, in the context of metrology and ISO 17025 quality systems, these tools and their validation must be properly documented and maintained.

In the context of uncertainty analysis, the computation speed of different pseudo-random number generators is not normally an issue: any typical computer can be used with any generator to calculate millions of points for an output quantity in a very reasonable amount of time, usually without concern about the computer language used for the implementation.

5.1 Underlying pseudo-random number generators

Virtually all pseudo-random number generators are based on some algorithm that creates a sequence of floating-point numbers, with each number restricted to the interval $(0, 1)$ or $[0, 1)$. The generators aim to approximate a sequence of samples of a random variable with a uniform probability density function on the interval $(0, 1)$. The sequence is not really random because the sequence can be restarted so that, for each restarted sequence, exactly the same sequence of floating point numbers will result. Good pseudo-random number generators have no discernible patterns to their sequence of numbers. Although the sequence has a finite length and would eventually repeat, many generators have a repeat length so long that accidental repetition is simply not a problem.

5.2 Compact algorithms

In metrology one aspect of floating-point random numbers is occasionally important: the dynamic range of the smallest steps near 0 and 1 to the interval size of 1. The Hill–Wichmann [HW82] generator has a dynamic range of about 2^{15} (30,000), which is adequate for many but not all applications in metrology. For some applications, the period of this generator, 7×10^{12} numbers, would also be considered to be short. The same authors have published an improved version [WH06] that is still quite compact.

The main advantage of a compact algorithm is found in the greater ease of checking the source code of the computer program. The other advantages of compactness can be translated into an expectation for a greater speed of execution. If proper software validation is to be performed on the executable code from any source code (however compact it might be) the apparent advantage of source code compactness will largely disappear.

5.3 Demonstrable randomness

One particularly good implementation of a pseudo-random number generator is known as RANLUX [J94], [L94]. This algorithm adds an interesting theoretical underpinning to the usual tests for assessing the statistical quality of randomness. Generally these tests only try to find some departure from ideal randomness, but ideal randomness in all its manifestations cannot be proven. For the RANLUX generator, the chaotic nature of the individual bits of binary representation is given a theoretical justification for portions of the underlying sequence [L94]. By discarding some, or all, of the potentially nonchaotic values, a family of pseudo-random number generators is created characterized by a specified “luxury level”. A luxury level of 4 means all potential departures from bit-by-bit chaos have been removed. The most commonly used member of this family is for a luxury level of 3, which also passes all of the usual empirical statistical tests for randomness. The period of this generator is 10^{171} numbers.

The RANLUX generator provides 24-bit pseudo-random numbers, and a standard extension is RANLUX64 that is used to concatenate two adjacent 24-bit pseudo-random numbers from the RANLUX sequence into a 48-bit mantissa for a double-precision floating-point number on $[0,1)$. It is easy to see that the bit-chaotic character is preserved by this concatenation, and so the theoretical basis for the randomness also applies to RANLUX64.

A version of RANLUX (and of RANLUX64) is available from the authors which has a symmetric approach to generating numbers close to 0 and to 1. Unlike the original implementation, no extra bits are concatenated for pseudo-random numbers near zero. For simulating uncertainty distributions, this has the effect of symmetrizing the wings of the resampled distributions. Asymmetry in the “extreme outliers” of measurements—even of simulated measurements—is likely to capture the attention of good metrologists, so it is worthwhile to avoid artificially introducing asymmetry via an underlying algorithm in a fashion that may be difficult to explain to other metrologists.

5.4 Testing with pseudo-random number generators

In the context of uncertainty analysis, testing with pseudo-random number generators is needed to support the validity of conclusions made using Monte Carlo methods. Overall, “testing” in uncertainty analysis should address a very broad range of issues, including:

- The implementation of the agreed basis for uncertainty analysis (random variables)
- The implementation of resampling of specific distributions that have been claimed to represent the uncertainties of the input quantities
- The adequacy of the measurement equations used for representing real measurements

- The adequacy of uncertainty claims to represent a set of measurements on quantities construed as invariant by a broad consensus
- The adequacy of the statistical methods themselves, tested by using the classical scientific method.

The Monte Carlo approach is powerful enough to support testing in all of these ways, we believe. The above scope of testing may seem daunting—offering too many years (or perhaps even centuries) of testing and improvement—until one examines what might reasonably be wanted by a skeptical client, colleague, or quality system assessor. In our view, it is really only the first two statements that need to be addressed at this time in order to justify a particular implementation of Monte Carlo methods as a means of delivering confidence in a specific uncertainty analysis. In its role of supporting the standard methods [GUM95], and justifying their approximations, only rather simple testing of the Monte Carlo tools must be performed and documented.

5.5 Tests of pseudo-random number generators

The primary test for a pseudo-random number generator (PRNG) aims at discovering departures from randomness in the underlying code that is to generate uniform random numbers on the interval $(0,1)$. When an algorithm for generating pseudo-random numbers is first published, some comment about the quality of its randomness is normally included. Exceptionally, a theoretical justification may be provided [L94], but more usually only results from one or more empirical test suites [M85], [LS02], [B07] are given. Generally, if one is relying on empirical tests, the more recent tests can provide stronger testing, but the requirements for randomness in some important applications (such as the characterization of cryptographic security) can be very different than the real requirements for modeling uncertainty.

For modeling uncertainty claims, the real requirements on randomness can be quite lax in comparison. In high-level metrology, uncertainty budget authors would only rarely be adamant in defending their value for the standard uncertainty in discussions about the possibility of a 10% change to the standard uncertainty. However, only a modest effort is required in documenting the randomness of your PRNG to a *much* higher level than this, with the great benefit of eliminating potentially distracting arguments about PRNG quality. This will involve some understanding, and perhaps tracking, of the requirements *perceived* by your target audience.

5.6 Tests on your computer

At first, it might seem important to repeat the tests, discussed above, of the underlying PRNG that is to generate uniform random numbers on the interval $(0,1)$. Careful checking of the source code, and rerunning one or more of

the empirical test suites might seem to be required. Because these empirical tests should ideally be rerun and checked whenever there is a change to your computer, this could burden users who automatically update their operating system and other software monthly (or even more frequently).

A practical and much shorter test can be almost as convincing. The PRNG is restarted from a standard state, and after generating a specified number of pseudo-random numbers, the next pseudo-random number is recorded and compared to the output by the standard implementation that has been used to run the empirical test suites. Almost all possible deterministic errors should be revealed by this process. For a variety of PRNG methods, the GNU Scientific Library and associated testing programs [C02] give a standard initialization and a standard random number reference value (for the value after 10^4 random numbers have been generated). This provides a simple, easily repeated, and easily documented test that can reliably associate the PRNG *in the present state of your computer* with earlier results of the empirical test suites.

Because neither theory nor empirical tests can provide generic assurance about randomness that is rigorously applicable to all circumstances that you may wish to examine in your Monte Carlo simulations, another particularly simple test is very helpful: to run your simulation with PRNGs of different qualities. Regions perturbed by PRNG artefacts will usually exhibit differences between results obtained using different PRNG methods. If the answers to questions about the model are the same for simulations performed with different PRNGs, this provides a convincing yet simple argument for there being no PRNG artefacts in the answer about the model.

For uncertainty analysis, there is another simple demonstration that is capable of convincing a wide audience of measurement experts: showing that the simulated randomness does indeed regenerate the PDFs claimed by the authors of the uncertainty budgets in question. Collecting and histogramming the results is quite tractable (see Figure 3), and can provide graphical reassurance to those whose primary concern rests in the position, width, and shape of these distributions. Most measurement experts are quite used to accepting the assumption of independence for random variables that represent perturbations which in reality have unknown correlations, correlations that are likely far greater departures from ideal randomness than would ever be found in a modern pseudo-random number generator.

5.7 Tests of dynamic range

There is a simple visual variant to the last approach that we have just discussed. Using a logarithmic vertical scale instead of a linear one for comparing the histogram of the simulated results with the form of the claimed PDF, which is usually expressed analytically, changes the character of the comparison from one of *probability* accuracy to one of *probability ratio* accuracy, allowing one to appreciate the *fractional* accuracy of the simulation relative to the claim. Confidence about the fractional accuracy is important when we are interested in statements about the chance occurrence of outliers.




Fig. 3. Histogram comparison of the limited dynamic range of the Hill–Wichmann PRNG versus RANLUX64, for simulating two uncertainty distributions (with 2×10^{10} events for each) by the table method: a Gaussian and a Student- t with 6.3 degrees of freedom. The two thin white lines are the analytic targets, which show the failings of the Hill–Wichmann PRNG for target values less than 10^{-4} of the peak; many histogram bins have 0 or 1 counts where 10s or 100s are expected.

This type of logarithmically scaled graph is a very simple means of establishing the validity of the dynamic range of the simulation: it can provide a convincing demonstration of the dynamic range of the underlying PRNG in the context of each specific algorithm for transforming to the target PDF. There are circumstances where users of Monte Carlo results really need to know that the implementation of the Monte Carlo resampling has sufficient dynamic range in the probability density to justify using the simulation to formulate conclusions about outliers of the modeled claims. Figure 3 shows the difficulties that the Hill–Wichmann PRNG can create when used outside its somewhat restrictive dynamic range. Also shown is the improvement that can accrue with a PRNG having a larger dynamic range, the RANLUX64 generator (with luxury level = 3). The table lookup method, used for the histograms in Figure 3 targeting a Gaussian, and a Student- t distribution with degrees of freedom = 6.3, uses but a single value at a time from the underlying uniform PRNG, and places the most stringent demands on dynamic range. Even so, the worst of the distributions in Figure 3 demonstrate that they are quite adequate for simulations requiring a dynamic range of less than 10^4 to 10^5 .

Of course, with large dynamic range comes a particular problem in measurement science. The low-probability tails of the targeted uncertainty claims

are usually difficult to evaluate experimentally, and the role of confidence about our Monte Carlo implementation is to direct attention towards improving the claims for these tails, rather than to have others dwell on possible deficiencies in the mechanics of the simulations. The claimed tails themselves may be of interest in making decisions [SWD05], and the Monte Carlo simulation can incorporate any specific decisions to record the statistical consequences about this complex kind of process.

6 Algorithms for generating specific distributions

To model the claimed distributions properly, we will need tools that can convert the uniform (pseudo-) random variate into a (pseudo-) random variate with the distribution claimed by the author of the uncertainty statement. The underlying uniform random number generator creates pseudo-random floating-point numbers from the interval (0,1), and all other needed distributions will be created by transforming samples from this sequence of numbers into other floating-point numbers that have the desired distribution. Each output number is based on one or more samples from the underlying uniform PRNG.

A multiplicity of input quantities needs to be simulated in most Monte Carlo simulations, and usually these all draw in turn on numbers from a single sequence from a single underlying uniform pseudo-random number generator. Below, we use a preceding subscript to denote the multiplicity of pseudo-random variates that usually needs to be created: $i = 1, 2, 3, \dots$ Gaussians; $j = 1, 2, 3, \dots$, Student- t variates, $k = 1, 2, 3, \dots$ uniform variates, ... etc.

6.1 Gaussian

i	Preceding subscript for the i th Gaussian PDF
${}_i x$	i th input quantity described by a Gaussian PDF
${}_i x_0$	The mean or reported value of the input quantity ${}_i x$
${}_i \sigma$	The standard uncertainty of ${}_i x$, its Gaussian standard deviation
${}_i y$	Standard-normal variate used to resample ${}_i x$

The Box–Muller method [BM58] is a standard method for transforming samples from a standard uniform PRNG on $[0, 1)$ into samples from a Gaussian (or “standard normal”) distribution, with mean value = 0 and standard deviation = 1. To use a Box–Muller output variable ${}_i y$ for uncertainty analysis, we convert to the quantity ${}_i x$ (which is an input quantity of our measurement equation) by using the simple shifting and scaling transformation

$${}_i x = {}_i \sigma \cdot {}_i y + {}_i x_0 \quad (5)$$

so the resulting variable ${}_i x$ has a Gaussian distribution, with a standard deviation equal to the claimed standard uncertainty, ${}_i \sigma$, and centered on the reported value ${}_i x_0$.

Although it may be the most common method, this is not the only way that a Gaussian might be approximated: table lookup and interpolation can be applied here just as readily as in the case of more complicated distributions; a reasonable approximation is to use the departure from 0.5 of the average of $N > 5$ samples from the underlying uniform PRNG; the Gaussian needed by one of our input quantities might originate in a GUM-compliant combination of other uncertainty components, and if the full uncertainty budget for this input quantity is available, it could be simulated in all its constituent components rather than as its combined standard uncertainty that is represented by a Gaussian.

The important observation in this discussion is that the way in which we use the modeled randomness is always independent of the mechanism used in the computation, and a specified PDF can always be regenerated in whatever way is clearest or is most convenient for us, even when the distribution in question is this most common basis PDF found in uncertainty budgets. Only if reanalysis of an uncertainty budget shows significant discrepancies from its author’s combined uncertainty calculation does one need to distinguish between the author’s combined uncertainty claim, and our claim of the combined uncertainty from the author’s uncertainty budget. This distinction applies to any method that we use for the recalculation, and is not an additional burden peculiar to Monte Carlo uncertainty analysis.

6.2 Student- t

j	Preceding subscript for j th input quantity described by a Student- t
${}_j x$	The j th input quantity described by a Student- t PDF
${}_j x_0$	The mean or reported value of the j th input quantity ${}_j x$
${}_j \sigma$	The standard uncertainty of ${}_j x$ (also the “sample standard deviation”)
${}_j \nu$	The degrees of freedom of ${}_j x$
${}_j y$	A Student- t variate used to resample ${}_j x$

The Kinderman–Monahan–Ramage method [KMR77] is a standard method for transforming samples from a standard uniform PRNG on $[0, 1)$ into samples from a Student- t distribution with a specified degrees of freedom ${}_j \nu$ (note that ${}_j \nu$ is not constrained to be an integer). It produces a zero-mean distribution having a *sample* standard deviation equal to 1. For ${}_j \nu > 2$ the

distribution standard deviation is $\sqrt{{}_j\nu/({}_j\nu - 2)}$. As published, the method fails if the underlying PRNG returns exactly 0.25, but it is a simple matter to avoid 0.25.

To use a Kinderman–Monahan–Ramage output variable ${}_jy$ for uncertainty analysis, we convert to the input quantity ${}_jx$ of our measurement equation by using the simple shifting and scaling transformation

$${}_jx = {}_j\sigma \cdot {}_jy + {}_jx_0 \tag{6}$$

so the resulting variable ${}_jx$ has a scaled and shifted Student *t*-distribution, with a standard deviation equal to the claimed standard uncertainty, ${}_j\sigma$ (recall that this is always to be interpreted as the *sample* standard deviation [GUM95]), and centered on the reported value ${}_jx_0$.

As was the case for the Gaussian distribution, there are other computation methods at our disposal, and for integer ν we can generate the Student-*t* variate by combining results from $\nu + 1$ samples from a Gaussian distribution. This approach can directly simulate the finite sampling uncertainty and its broadening of a Gaussian PDF's wings into those of a Student-*t* PDF.

6.3 Uniform distribution

k	Index for the k th input quantity described by a uniform distribution
${}_kx$	The k th input quantity described by a uniform distribution
${}_kx_0$	The mean or reported value of the k th input quantity ${}_kx$
${}_ka$	The semi-width of the uniform distribution for ${}_kx$
${}_ky$	A uniform variate on (0,1) used to resample ${}_kx$

It is particularly simple to use the underlying uniform PRNG to simulate any quantity described by a uniform rectangular distribution. Using a value ${}_ky$ from our underlying PRNG, for uncertainty analysis we convert to the input quantity ${}_kx$ of our measurement equation by using

$${}_kx = {}_ka({}_ky - 0.5) + {}_kx_0 \tag{7}$$

so the resulting variable ${}_kx$ has a scaled and shifted distribution, with a full rectangular width that is equal to a and centered on the reported value ${}_kx_0$. Recall [GUM95] that $a/\sqrt{12} = (a/2)/\sqrt{3}$ is the distribution standard deviation and is to be the claimed standard uncertainty.

Once again, the numerical method can be chosen from many alternatives: in this case the choice can be made from the many candidates available for the underlying uniform pseudo-random number generators.

6.4 U-distribution

l	Index for the l th input quantity described by an arcsine-U distribution
${}_l x$	The l th input quantity described by an arcsine-U distribution PDF
${}_l x_0$	The reported value of ${}_l x$, and the mean of its arcsine-U distribution
${}_l a$	For ${}_l x$, the semi-width of its arcsine-U distribution
${}_l y$	Uniform variate on (0,1) used to resample the arcsine-U distribution for ${}_l x$
v	A random variate used for integrating the arcsine-U distribution
δx	A small increment in the input quantity x
θ	A uniformly distributed random variate on $(0, \pi)$
$\delta\theta$	A small increment in theta

The U-distribution, or arcsine-U distribution is chosen here as our last class of classical analytic uncertainty distributions. It is neither the next most common nor necessarily the next most useful distribution for uncertainty analysis. It does offer a particularly clear physical basis for the transformation from one uncertainty distribution to another, and is a good example of the simplification that can be had in uncertainty modeling by choosing the input quantities carefully. It is also a convenient distribution to use for demonstrating the general methods that we need when no such physical model is available.

This distribution is used to describe random sampling from a measurand subject to a sinusoidal perturbation. For example, consider the temperature of a system with a long time constant, controlled by a temperature controller between two set limits. When the cycling between the two limits is faster than the natural time constant, the main perturbation can be sinusoidal. If measurement times are uncorrelated with the temperature controller cycling, then the probability density can be modeled by emulating the sampling process of the cyclic variation: we draw a pseudo-random number ${}_l y$ from a distribution uniform on $(0, 1)$, multiply by 2π to emulate a random phase at which the sinusoid is sampled, scale it to the appropriate amplitude ${}_l a$ (half the peak-to-peak amplitude), and shift it to the desired mean value ${}_l x_0$,

$${}_l x = -{}_l a \cdot \cos(\pi {}_l y) + {}_l x_0, \quad (8)$$

a process that is physically transparent in situations where we are considering a random-phase sample of a sinusoidally varying input quantity. The standard uncertainty is the root mean square of the departures of ${}_l x$ from ${}_l x_0$, ${}_l a/\sqrt{2}$. Here we are considering one branch: as ${}_l y$ runs from 0 to 1, ${}_l x$ runs from ${}_l x_0 - {}_l a$ to ${}_l x_0 + {}_l a$.

However, this is not the only route to resampling. We can also evaluate the PDF for this resampling process. The density of samples from the sinusoidal waveform giving a value between x and $x + \delta x$ will be proportional to the fraction of the time that the waveform spends at that value, $\delta\theta/2\pi$, derived

from the inverse of slope of the sinusoid at $x = \sin(\theta)$, $(\delta x/\delta\theta) = \cos(\theta)$; that is, $1/\cos(\theta) = 1/\sqrt{1 - \sin^2 \theta} = 1/\sqrt{1 - x^2}$. Thus the PDF of x is proportional to $1/\sqrt{1 - x^2}$.

6.5 General PRNG methods – for the U-distribution

We now wish to consider a general procedure for converting a specified shape of PDF into a usable means for converting a value from a uniform PRNG into a pseudo-random sample from the specified shape.

The first step is to create the *cumulative distribution function*, $CDF(x)$ from our specific probability density function $PDF(x)$ on the interval $[-1,1]$ and zero outside $[1,1]$. The general definition of the CDF is

$$CDF(x) = \int_{-\infty}^x PDF(v)dv \tag{9}$$

which, in the above example of the U-distribution, can be evaluated explicitly

$$CDF(x) = \frac{\int_{-1}^x (1 - v^2)^{-1/2} dv}{\int_{-1}^1 (1 - v^2)^{-1/2} dv} \tag{10}$$

to arrive at

$$CDF(x) = 0.5 + \frac{\arcsin(x)}{\pi} \tag{11}$$

which is the origin of the usual name for this kind of random variate, an arcsine-U variate. When dealing with PDFs that are described by un-normalized shapes, as here, it is necessary to divide by the normalizing area under the shape as shown in Equation (10). Note that in this example the shape of $PDF(x)$ is symmetric about $x = 0$, and so $CDF(0) = 0.5$ exactly.

In general, the cumulative distribution function will provide the path for transforming values ${}_m y$ from our underlying PRNGs, uniform on $(0, 1)$, into samples ${}_m x$ of a random variate that match the CDF and its own defining PDF shape. Given a value ${}_m y$ from our underlying PRNG, the transformation involves solving for ${}_m x$ in the equation

$$CDF({}_m x) = {}_m y \tag{12}$$

which in general would need to be solved numerically, but in this case also has an analytic solution, ${}_m x = \sin(\pi {}_m y - \pi/2) = -\cos(\pi {}_m y)$, which after shifting and scaling gives Equation (8).

The use of the arcsine-U distribution is not restricted to cases for which there is a neat physical justification. It also is substituted for a uniform distribution when the author of an uncertainty budget wants to explore the consequences of a “less central” distribution than the uniform (rectangular) distribution. Defined as here, spanning $[-1, 1]$, the PDF is $\pi^{-1}[1 - x^2]^{-1/2}$ and has a variance of 0.5, giving a standard uncertainty of $1/\sqrt{2}$ (of the half-width which is equal to 1 in this unscaled version of the arcsine-U PDF).

6.6 General PRNG Methods for Tabulated Distributions

(x_i, x_{i+1})	Interval containing a specific value x with PDF tabulated as $f(x_i)$
Δ	The bin width used in a histogram
${}_m x$	A random variate on $(-\infty, +\infty)$ possessing a specific CDF
${}_m y$	A uniform random variate on $(0,1)$
ι, κ, η	Indices to N tabulated values of an input quantity and its PDF
x_ι	ι th tabulated value of input quantity x , with a PDF value $f(x_\iota)$
f_ι	$f(x_\iota)$, the tabulated value of the PDF at a tabulated value x_ι
CDF	Cumulative distribution function
$F(x)$	CDF of x , integrating linear interpolation of PDF points f_ι at x_ι
F_κ	The tabulated value of the CDF, $F(x_\kappa)$
RV	Random variate
${}_n x$	RV representing some n th input quantity
${}_n x_r$	The reporting value for the input quantity ${}_n x$: a “fiducial value”
${}_n x'$	Random variate representing measurand–measurement conjugate to ${}_n x$
${}_n y$	RV uniform on $(0,1)$, to make RV ${}_n x$ from its tabulated CDF, $\text{CDF}({}_n x)$

In uncertainty analysis, we can simulate the most general randomness using only a tabulated PDF. We do not need access to the physical model that gave rise to the tabulated PDF. The tabulated PDF can be interpreted as a continuous probability density function $f(x)$ given in Equation (4). For any x between x_ι and $x_{\iota+1}$, and for any ι denoting $f(x_\iota)$ as f_ι we can define the appropriate cumulative distribution function

$$F(x) = \frac{[\sum_{\eta=1}^{\iota} (x_{\eta+1} - x_\eta) \frac{f_{\eta+1} + f_\eta}{2}] + (x - x_\iota) f_\iota + \frac{(x - x_\iota)^2}{(x_{\iota+1} - x_\iota)} (f_{\iota+1} - f_\iota)}{\sum_{\eta=1}^N (x_{\eta+1} - x_\eta) \frac{f_{\eta+1} + f_\eta}{2}} \quad (13)$$

and use this to transform a value ${}_n y$ from a uniform pseudo-random number generator into a sample ${}_n x$ of a random variate with a PDF = $f(x)$. Again, we solve the equation

$$F({}_n x) = {}_n y \quad (14)$$

which can be done in two steps. Denoting $F(x_\kappa)$ as F_κ , in the first step we find an index k such that $F_\kappa < {}_n y < F_{\kappa+1}$. In the second step we solve Equation (14) for x , knowing that $x_\kappa < {}_n x < x_{\kappa+1}$. A linear interpolation of $F(x_\eta)$ will only simulate a staircase approximation to $f(x)$, but by solving Equation (14) using the full Equation (13) we can recover our original $f(x)$ as defined in Equation (4). If $f_k \neq f_{\kappa+1}$, then the solution for Equation (14) is a solution to a quadratic equation

$${}_n x = \frac{-f_\kappa + \sqrt{f_\kappa^2 + 2(ny - F_\kappa)(f_{\kappa+1} - f_\kappa)/(x_{\kappa+1} - x_\kappa)}}{(f_{\kappa+1} - f_\kappa)/(x_{\kappa+1} - x_\kappa)} \tag{15}$$

and if $f_\kappa = f_{\kappa+1}$ the solution is

$${}_n x = \frac{{}_n y - F_\kappa}{f_\kappa} \tag{16}$$

producing a value for ${}_n x$ which may be used directly, if the mean and standard deviation are exactly what is wanted in the simulation. Otherwise, it may be scaled and shifted as we have done for the distributions discussed above.

For any general distribution that is not symmetric about its mean, it is significantly simpler to simulate a zero-mean process. In using such a distribution rigorously, one is forced to choose whether one is simulating a distribution of samples of the measurement, or is simulating samples of the measurand. Usually, it is the mean value of an input quantity that is the reported value, and then a zero-mean simulation needs only a simple change of sign to transform between the interpretation of the simulated values ${}_n x$ as samples of the measurement process and the interpretation of the simulated values as samples of possible values ${}_n x'$ of the measurand. Otherwise, if ${}_n x_r$ is the reported experimental (or fiducial) value of an input quantity, the transformation is

$${}_n x' = {}_n x_r - ({}_n x - {}_n x_r) \tag{17}$$

or the neater but more obscure ${}_n x' = (2 \cdot {}_n x_r) - {}_n x$. Clearly, there is no significant simplification of the arithmetic. However, by simulating a zero-mean process we can keep this question of measurement versus measurand picture within the context of uncertainty analysis, whereas the use of Equation (17) evokes the additional scrutiny demanded by any source of potential bias.

6.7 Choosing the number of simulated events

N_e	Number (#) of fully-resampled simulated events in a Monte Carlo simulation
n	# of simulated events with an output quantity between specified limits
n_l	# in the histogram bin (width Δ) with $p_l N_e$ events to the left of this bin
n_r	# in the histogram bin (width Δ) with $p_r N_e$ events to the right of this bin
p_l	The coverage probability to the left of a coverage interval
p_r	The coverage probability to the right of a coverage interval

The output quantities of interest will generally be proportional to the fraction n/N_e of the number of simulated events N_e that create n measurement

equation values between two limits. The accuracy with which the fraction n/N_e can be ascertained can be evaluated easily in several different ways.

- For $n \ll N_e$, or for $N_e - n \ll N_e$, the resampling uncertainty (the standard deviation) in n can be estimated as \sqrt{n} , and the relative accuracy of n/N_e as $(n/N_e)(\sqrt{n}/n) = \sqrt{n}/N_e$.

- As n gets larger, a better estimate may be useful that takes into account the randomness, not only in n , but also the anticorrelated randomness in $N_e - n$ (there being no randomness in N_e). The resampling uncertainty (the standard deviation) in n for this case is to be estimated as $\sqrt{n(1 - n/N_e)}$.

- If the simulation involves complex decision trees, for which the two estimators above are in doubt (they are based on properties of the Poisson and binomial distributions), then by repeating the simulation m times, without restarting the PRNG, and by calculating the standard deviation of the resulting m values for n , and dividing by \sqrt{m} , one would usually obtain a good estimate of the standard deviation of the mean of n . The PRNG must be random across each subset n of the N_e events, but this is expected from the PRNG if N_e is less than the PRNG sequence length.

If there is some particular level of accuracy wanted in n/N_e , a reconnaissance run can give an estimated fraction n/N_e and allow a simple calculation of a value for N_e that will deliver the appropriate accuracy in n/N_e by using one of the above techniques.

For some quantities, such as coverage intervals, the estimate is only slightly more involved. Consider a coverage interval leaving a fraction p_l events to the left of the coverage interval, and p_r to the right. By treating the two ends as independent in a simulation with N_e events overall, the uncertainty in the coverage interval can be estimated from the bin width Δ , the number of events n_l in the bin closest to having $p_l N_e$ events to the left of it, and the number of events n_r in the bin closest to having $p_r N_e$ events to the right of it. Due to the finite number of Monte Carlo resampled events, there is an additional uncertainty component for breadth of the coverage interval having a standard uncertainty

$$\Delta \sqrt{\frac{p_l(1 - p_l)N_e}{n_l^2} + \frac{p_r(1 - p_r)N_e}{n_r^2}}.$$

Commonly this will be for a probabilistically symmetric interval ($p_l = p_r$), with a coverage probability of $95\% = 100\%(1 - p_l - p_r)$, so $p_l = p_r = 0.025$ and the simulation-induced standard uncertainty in the coverage interval is $0.1561\Delta\sqrt{N_e(n_l^{-2} + n_r^{-2})}$. The choice of an appropriate bin width Δ , and the number N_e of events to be used in the simulation, can be easily obtained from one or more reconnaissance simulations using skills that are highly developed in measurement scientists by extensive laboratory experience with real experiments.

7 Examples of Monte Carlo simulation

i, k	Index to indicate Laboratory i or k in an M -lab comparison example
${}_i x_0$	Mean or “fiducial” value reported by lab i
${}_i u$	The uncertainty reported by lab i
${}_i \sigma$	Gaussian’s standard deviation representing uncertainty of Lab i , usually ${}_i u$
${}_i x_j$	The j th event’s simulated value for lab i : ${}_1 x_j$ for lab 1, etc.
\bar{x}_j	Simple mean of the M values simulated for the M laboratories in event j
\tilde{x}_j	The inverse-variance weighted mean of the M simulated values in event j
M	Number of labs in a comparison and its simulated comparison

7.1 Gaussian PDFs

To illustrate the broad scope of Monte Carlo analysis of uncertainties in measurement, we have constructed an example based on interlaboratory measurement comparisons that circulate a stable artefact to be measured by a group of laboratories, aiming to provide confidence in the claimed measurement capabilities by demonstrating consistency. Table 1 and Figure 4 show what might be a typical comparison involving $M = 6$ laboratories: each has reported a value and a standard uncertainty that are conventionally plotted as shown in the left panel (although sometimes as error bars an expanded uncertainty might be used instead of, or in addition to, the standard uncertainty).

Recall that in the absence of explicit information to the contrary, the primary interpretation of the standard uncertainty [GUM95] is as the standard deviation of a Gaussian probability density function. If a degrees of freedom is explicitly given, it is a Student- t distribution that is to be construed, with the standard uncertainty to be used as its *sample* standard deviation. Allowance is made for any other departures from Gaussians, provided that the departures are described. In this example, we start by simulating Gaussian PDFs and work towards incorporating non-Gaussian PDFs.

Table 1. An artificial set of comparison values ${}_i x$ with standard uncertainties ${}_i u$, used as the example in all the analyses that follow.

Lab i	${}_i x_0$	${}_i u$
Lab 1	0.335	0.335
Lab 2	-0.251	0.251
Lab 3	-0.262	0.393
Lab 4	-0.271	0.575
Lab 5	0.640	0.640
Lab 6	-0.471	0.270




Fig. 4. Left: the artificial set of comparison values from Table 1, with error bars showing the claimed standard uncertainties. Right: the same data with axes interchanged to also show Gaussian PDFs, having standard deviations equal to the claimed standard uncertainties, and means equal to the reported values.

The six PDFs that account for the randomness in the measurand are reinterpreted as describing resampled measurements, [DSWH05], resampling the measurement as all influence quantities are imagined to vary randomly over their natural ranges. As an alternative, with some minor exceptions it would be possible to run Monte Carlo simulations of uncertainties within the perspective of resampling the measurand; but with impaired opportunities, we believe, for practical guidance from laboratory experience and experiments. The resampling is not, of course, from reality but is from the claimed model of uncertainty and randomness. The role of the Monte Carlo simulation is restricted to answering questions relating to the model and the adequacy of the model to describe the available measurements of reality. Note that this reinterpretation is what would commonly be done when a measurement (and its uncertainty) is actually used with another measurement to make a decision: “. . .the reported value is . . . but the reported value *could* also have been distributed like this . . .”, and that this reinterpretation is a null operation except for asymmetric distributions.

The quality of the six resampled variables from the six Gaussians of Figure 4 (derived from Table 1) is illustrated graphically in Figure 5, where the histograms of the six Monte Carlo resampled variables have been recorded for 10^5 events, each drawn from the Hill-Wichmann uniform PRNG [HW82] and the Box-Muller [BM58] transformation to a standard normal distribution which is scaled and shifted (as per Equation (5)) to have the claimed mean ${}_i x_0$ and standard uncertainty ${}_i u$ equal to the standard deviation of the Gaussian, ${}_i \sigma$. Superimposed on each histogram is its analytic Gaussian distribution. This kind of graph can convincingly show that the claimed means and randomness are being properly resampled. Note that here it is inappropriate to impose any external constraint (such as a common-mean constraint) to permit us to simulate the consequences of the as-claimed values and uncertainties. The




Fig. 5. The Gaussian approximation to the comparison of Figure 4. The histograms of 10^5 simulated points are shown for the six labs, with the analytic Gaussians superimposed in white.

data for Figure 5 can be calculated in a few seconds, and it is not difficult to prepare a graph where there is no visible simulation noise left in the counts in the histogram bins. In our experience, graphs with visible noise are very often more effective than smooth ones for introducing the idea that the resampling is being properly done.

For our first example of a measurement equation, let us consider the simple mean as the output quantity. For the j th set of six pseudo-measurements, $\{i x_j, i = 1, \dots, 6\}$ it is completely straightforward to calculate the simple mean \bar{x}_j of the set as

$$\bar{x}_j = \frac{1}{M} \sum_{i=1}^M i x_j. \tag{18}$$

Of course, any other measurement equation involving the $\{i x\}$ could be substituted for the simple mean, and for physically meaningful measurement equations we would expect no difficulties. The resulting output quantity values, \bar{x}_j , are histogrammed, as shown in Figure 6 for the same 10^5 sets used for the six histograms of Figure 5. Also shown in Figure 6 is the analytic Gaussian distribution that would usually be attributed to the simple mean: centered at the simple mean of the six experimental values, and having a standard deviation equal to $\sqrt{(1/M)/\sum_{i=1}^M i \sigma^2}$. As we should expect, the




Fig. 6. The Monte Carlo simulation of 10^5 samples of the simple mean of the set of six simulated measurements of Figure 5, each drawn from a Gaussian distribution. Superimposed in white is the analytic answer.

agreement between the Monte Carlo simulation and the analytic methods of conventional statistics is excellent in this example where the claimed uncertainties are Gaussian and can be handled by the techniques of conventional statistics.

However, we do not have control over the claims of the six participating laboratories. Their uncertainty budgets might not be Gaussians. They may have degrees of freedom specified, or they have been prepared by combining uncertainties using Monte Carlo methods and specified as a Monte Carlo histogram. We now turn to a more interesting simulation for which the applicability of conventional statistical tools could be in doubt.

7.2 Non-Gaussian PDFs and measurement equations

To make the point that arbitrary claimed PDFs can be handled by Monte Carlo simulation, we now simulate the same example as above (Table 1), but for a set of fanciful PDFs that all have marked departures from Gaussian distributions. The sources of the six PDFs are shown in Figure 7. The PDFs, shown in Figure 8, are not all symmetric, so we also need to clarify that we are taking these PDFs as specifying the distribution for measurements that the model predicts when all influence quantities are varying randomly over their natural ranges. The distribution standard deviations ${}_i\sigma$ are taken as the standard uncertainties ${}_i u$ from Table 1, and the distribution means are taken as the ${}_i x_0$ from Table 1, as illustrated in Figure 9. Repeated comparisons



Fig. 7. A demonstration of coping with arbitrary uncertainty distributions. Six familiar shapes are used to provide the basis for PDFs in a fanciful comparison of six NMIs.




Fig. 8. Tabulated (but unnormalized) PDFs from Figure 7. These PDFs are used, as described in the text, to construct six pseudo-random variates, one sampled from each tabulated PDF, and representing resampled measurements.

are each simulated from six consecutive pseudo-random numbers, uniform on $(0,1)$, from the underlying PRNG, transformed to the six values that constitute a simulated comparison. As discussed in Section 3.3, tabulated PDF distributions were used (Figure 10A) to calculate the appropriate CDF (Figure 10B) that is used to transform samples from a uniform PRNG into samples




Fig. 9. The comparison data of Table 1 and Figure 2, with a fanciful PDF raised to be beside each point. Each PDF's mean and standard deviation are matched to the value and standard uncertainty of Table 1, also plotted as the point and error bar.

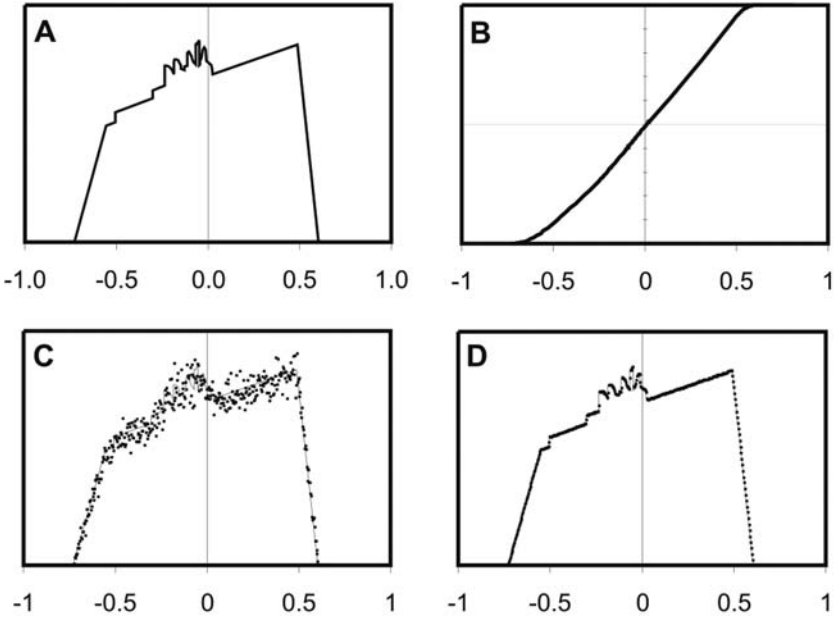


Fig. 10. (A) Tabulated input PDF for lab 1 of Figure 9. (B) CDF for A, probability is vertical axis, from 0 to 1. (C) Monte Carlo histogram of 10^5 events compared to tabulated input. (D) As for (C), but with 10^8 events. Note the standard deviation in a histogram bin containing n events is \sqrt{n} , or a fractional accuracy of $1/\sqrt{n}$.

from the tabulated PDF (Figure 10C for a histogram of 10^5 samples and Figure 10D for a histogram of 10^8 samples). From these simulated values, any measurement equation, or indeed any measurement algorithm, can be used to calculate an output value for which simulated values can be accumulated in a histogram.

Measurement equation: simple mean

Despite the markedly non-Gaussian nature of the individual PDFs for the six input quantities, when we calculate the simple mean as a candidate reference value (using Equation (18) as our measurement equation), we find that the resulting distribution of the output quantity \bar{x} differs only very slightly from the all-Gaussian approach discussed above. For the simple mean, no departures are evident on the linear graph of the left panel of Figure 11, where the analytic Gaussian PDF for the simple mean is shown on top of the Monte Carlo histogram for the simple mean obtained by using the tabulated-PDF method of Section 3.3 to model the PDFs shown in Figures 7–10.

For applications exhibiting no significant differences between results from the simple Gaussian approximation and from the full Monte Carlo simulation using the tabulated PDFs, a simple linear PDF comparison graph, such as those in Figure 11, may be a sufficiently quantitative characterization of the differences over the scope or scopes of interest. For specific applications, it may be appropriate to use other means to describe the adequacy of the simple Gaussian methods. If the low-probability tails are of particular interest, the adequacy of the simple Gaussian methods may be better appreciated if the PDF axis is graphed using a logarithmic scale. If one specific parameter is of interest, the accuracy of the simple Gaussian method may be simply characterized as an accuracy in percent for that parameter. However it is reported, this description should aim for simplicity in the eyes of the end users.




Fig. 11. Monte Carlo histograms (black) using 10^8 events to simulate three methods for calculating a reference value for measurements from Figure 9: the simple mean (left), the inverse-variance weighted mean (middle), and the median (right). The superimposed gray lines are based on GUM-derived uncertainty methods (see text).

Measurement equation: inverse-variance weighted mean

The inverse-variance weighted mean is another candidate reference value that has an uncertainty which we can model by Monte Carlo simulation. This weighted mean is defined by the measurement equation

$$\bar{x}_j = \frac{1}{\sum_{i=1}^M i\sigma^{-2}} \sum_{i=1}^M \frac{ix_j}{i\sigma^2} \quad (19)$$

which is still just a linear combination of the six simulated lab measurements that constitute an “event set” in the Monte Carlo simulation. In the Gaussian approximation, the expected uncertainty in \bar{x} is

$$u(\bar{x}) = 1/\sqrt{\sum_{i=1}^M i\sigma^{-2}},$$

shown in the center panel of Figure 11 as the light gray line superimposed on the black line of the histogram of weighted means from a Monte Carlo simulation of the fanciful PDFs of Figures 7–10. The Gaussian gives a good approximation to the full simulation’s distribution of the weighted mean’s within-method randomness to be expected from the six uncertainty claims shown in Figure 9.

Measurement algorithm: median

Monte Carlo simulation can analyze uncertainty for any automatable algorithm. As an example of this, we now consider the median of the fanciful comparison as the output quantity of interest. As applied to a set of M measurements, the usual definition of the median is the value of the measurement having as many measurements above as it has below, if M is odd. If M is even, the median is taken as the simple mean of the adjacent pair of measurements that have $(M - 1)/2$ measurements below them in value, and the same number above.

It is straightforward to automate the procedure by sorting the measurements as to value, and taking—for odd M —the $[(M + 1)/2]$ th value as the median. If M is even, the average of the $[M/2]$ th and the $[(M/2) + 1]$ th sorted values is the median value of the measurements. The same algorithm can be applied to each set of pseudo-measurements in the process of resampling the median to obtain a histogram of possible median values that the stated measurement values and the claimed uncertainty distributions can produce. This is the commonest form [SWD06] of a broader family of order statistics that rely on the relative positions of measurements, and sometimes also use the claimed uncertainties, that seek to de-emphasize “outliers”. It is not generally

feasible to analyze uncertainty propagation for these algorithms by using the standard methods presented in the ISO-GUM [GUM95], yet the resampling technique has no difficulty in producing resampled histograms showing how the claimed values and their randomness produce pseudo-measurements and how these sets of pseudo-measurements react to the analysis algorithm, which in this case is the median algorithm.

Note that the median algorithm has been clearly defined for measurements, and both intuitions and the possibility of direct experimental tests lead us to prefer the measurement picture here, speaking of the “ M measurements” rather than the “ M values reported for the measurand.” Thus we reinterpret measurand-perspective uncertainty statements about our input quantities as statements about the randomness to be expected in the measurements if all components of the uncertainty budget were allowed to vary over their natural ranges. Furthermore, we may want to reconvert the histogram of output quantity (expressed in the perspective of measurements) back into the measurand perspective [DSWH05]. For the usual definition of the median, there is another detail to be mastered in these translations of perspective: the median of the reported values is not necessarily equal to the mean of the median distribution. If we are scrupulous about reflecting in the reported value for the median [DSWH05] it is possible to analyze the median algorithm rigorously and exactly as it is normally used. In the right panel of Figure 11, the black line shows the resampling histogram of median values (still expressed as a distribution of measurements). On the same graph, the gray line shows a Gaussian approximation with a standard deviation derived from the square root of the noncentral second moment of the distribution of resampled medians about the experimental median.

The power of the Monte Carlo method allows us to address other questions about the median as a measurement algorithm. In the context of comparisons of laboratory results, the question, “How often is each lab expected to be the median?” is something that can be tracked within a Monte Carlo simulation. For an odd number M of values, it is possible to track which one is the median value and to count the number of times that each one contributes the median value. For even M , the median value is contributed by a pair of labs, and each lab’s count would be incremented by one-half. In our fanciful example of Figure 9, from top to bottom the “median weights” for the resampled median are 13%, 30%, 22%, 16%, 7%, and 12%. Note that in the experiment, the median is the average of the second and third values. These percentages can provide useful insight for explaining how the median algorithm handles the randomness claimed for a particular comparison.

Measurement equation: Departure from the weighted mean

In CIPM key comparisons, it is common practice to report the measurements as departures from a “key comparison reference value” (KCRV) and to report

the expanded uncertainty of this departure for a coverage probability of 95%. Usually the KCRV is algorithmically derived from the reported measurements and uncertainties. Often the KCRV is calculated as the simple mean, or as the inverse-variance weighted mean, or as the median. As shown in Figure 11, Monte Carlo methods can easily resample any of these methods. Sometimes outlier rejection schemes are invoked that are based on the observed values and the claimed uncertainties, and if the method for doing so has been fully described, these can also be incorporated into the KCRV algorithm that is simulated.

Here, we illustrate the simulation of a comparison (Figure 9) to report for one lab (lab 1, as shown in Figure 9 and in Figure 10) departure from a reference value that is the inverse-variance weighted mean, with no outlier rejection. This example illustrates how the Monte Carlo simulation simplifies the handling of correlations due to the reuse of values in a measurement equation. The measurement equation is

$${}_1x_j - \bar{x}_j = {}_1x_j - \frac{1}{\sum_{i=1}^M i\sigma^{-2}} \sum_{i=1}^M \frac{i x_j}{i\sigma^2} \tag{20}$$

and the reuse of ${}_1x$ as $i x$ for $i = 1$ is obvious and is automatically incorporated in the simulation with no additional effort. For traditional methods, there is an unavoidable covariance to be accounted for. The ISO-GUM uncertainty for Equation (20) is described by a Gaussian with standard deviation $u({}_1x_j - \bar{x})$ given by

$$u({}_1x_j - \bar{x}) = \sqrt{u^2({}_1x) + u^2(\bar{x}) - 2 \text{cov}({}_1x, \bar{x})}, \tag{21}$$

where both $u^2(\bar{x})$ and $\text{cov}({}_1x, \bar{x})$ are equal to $(\sum_{i=1}^M i\sigma^{-2})^{-1}$, so that Equation (21) becomes $u({}_1x_j - \bar{x}) = \sqrt{u^2({}_1x) - u^2(\bar{x})}$. The Monte Carlo simulation’s histogram of ${}_1x_j - \bar{x}$ is compared to the ISO-GUM’s Gaussian approximation (Equation (21)) in Figure 12. The expanded uncertainty for a coverage probability of 95% is also shown for each method in Figure 12. The ISO-GUM expanded uncertainty is for a coverage factor $k = 1.96$ appropriate for a Gaussian distribution. It is plotted as the gray horizontal error bar in Figure 12. The probabilistically symmetric coverage interval of the Monte Carlo histogram is shown as the black horizontal error bar: 2.5% of the histogram is to the left of this interval and 2.5% is to the right. The 95% coverage interval is slightly smaller than the Gaussian, and is slightly asymmetric: $0.50 + 0.53, -0.57$ (for measurements). In the measurand perspective, this coverage interval should be reported as $0.50 + 0.57, -0.53$.

This example illustrates that even with strikingly large departures from Gaussian distributions, the Gaussian approximation can yield results that may be satisfactory for many purposes.




Fig. 12. Monte Carlo resampled comparison of Figure 9 claims as repeated pseudo-measurements: histogram of “lab 1 minus the weighted mean” (black) compared to the GUM Gaussian (gray). The error bars show 95% coverage intervals. The Monte Carlo interval (black error bar) is the probabilistically symmetric interval.

8 Monte Carlo simulation for hypothesis testing

χ^2	Reduced chi-square of departures from inverse-variance weighted mean
z	A particular value of a chi-square
E_n	Normalized error: departure divided by its standard uncertainty
DoF	Degrees of freedom, describing a χ^2 distribution
χ^2_{APD}	χ^2 -like aggregate over all pair differences for nominally equal pairs

All of the simulations above have used *reported* values and *claimed* uncertainties and predict where measurements might be found on the basis of these claims. However, Monte Carlo simulation can also use these as claims about reality to test for compatibility with an external hypothesis.

To test an external hypothesis in this way, we run a simulation where *we assume that this hypothesis is true*. There is no escaping making this assumption temporarily, for the duration of the test; and the test may only provide a means for quantifying the compelling evidence revealed for rejecting the hypothesis.

Unlike any of the simulations in earlier sections, hypothesis-testing Monte Carlo simulations examine the probabilistic consequences of the modeled randomness (i.e., the claimed uncertainties in the context of a particular measurement equation) in a fictitious universe where the hypothesis under test is forced to be true. This is a profoundly different kind of simulation, and will only touch reality in statements of how likely it would be for a real experimental value of a statistical aggregate to be exceeded in the modeled fictitious

universe where the hypothesis has been forced to be true. Finally, this is interpreted as quantifying the strength of the evidence for rejecting the hypothesis under test, paralleling the process of conventional theorem-based χ^2 -testing.

Although we made the assumption for the purposes of testing the hypothesis, it is crucial to convey to all users of our test results that the *absence* of compelling evidence to reject our hypothesis is not particularly good evidence to accept the hypothesis. Specifically, if the test “passes” a criterion that limits false rejection to 5%, all users of the test result must be protected from the disastrous logical mistake of associating this in any way with a 95% confidence in the hypothesis being true.

8.1 Chi-squared testing of the common mean hypothesis

One of the triumphs of theorem-based statistics is the clarity of analysis that can be brought to bear in testing the external hypothesis of the common mean for claimed Gaussian uncertainty distributions that are independent (which may be a second external assumption). Another external hypothesis may be needed—that using the inverse-variance weighted mean \bar{x} really is an acceptable method for determining a reference value—with no need to impose weight limits as would often be considered in high-precision comparisons.

The idea is that the square of the departures of the M constituent measurements from the weighted mean can be aggregated to yield an analytic reduced-chi-squared distribution if the hypotheses discussed above are true.

$$\chi^2 = \frac{1}{M-1} \sum_{i=1}^M \frac{(ix - \bar{x})^2}{i\sigma^2}. \quad (22)$$

The hypotheses are usually considered as a single composite hypothesis termed the “null hypothesis” in chi-square testing. The experimentally observed value of this aggregated statistic can be associated with a probability for falsely rejecting the hypotheses when they are true, and this can be used to help evaluate how compelling the evidence is for actually rejecting the combined hypothesis.

The reduced chi-square that is appropriate is one with a degrees of freedom parameter equal to $M - 1$. The reduced chi-squared PDF is defined on the interval $(0, \infty)$, as is its cumulative distribution function (CDF). The probability that M measurements, with a common mean and fulfilling all the other assumptions we have just discussed, has a χ^2 aggregate that exceeds a value z is given by $(1 - \text{CDF}(z))$, as shown by the dotted gray lines in Figure 13, for a degrees of freedom = 5, which is appropriate to the Gaussian model of six measurements shown in Figure 4.

From the values and uncertainties of the comparison in Figure 4, for this imagined experiment the value of the reduced chi-squared statistic is 1.060, and if the common mean hypothesis is forced to be true then the probability of the reduced chi-squared exceeding 1.060 is 38.0%, taken from the analytic

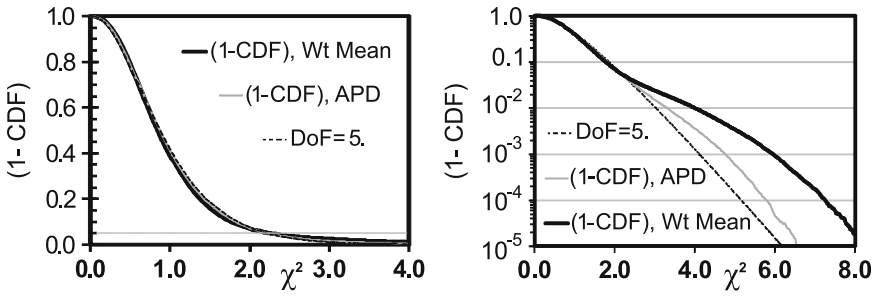


Fig. 13. Probability of three test statistics exceeding a value of χ^2 , for a six-lab comparison, as linear (left) and logarithmic (right) graphs. The thin dashed line is the analytic reduced chi-squared for a degrees of freedom (DoF) = 5. The other two χ^2 -like statistics need Monte Carlo simulation with a null hypothesis of “agreement” to evaluate the randomness of Figure 9: the black line is for agreement with the weighted mean reference value (Wt Mean), and the gray line is for pairwise agreement of all labs, using the all pair differences (APD) mean-square aggregate of E_n , χ_{APD}^2 .

reduced chi-squared distribution for a degrees of freedom = 5 shown as the dotted line in Figure 13. There is no compelling evidence in the experiment’s dispersion to reject the common mean hypothesis in the context of the claimed uncertainties. Note that “not rejecting” is not the same as “accepting” or “validating”.

The main limitations to the method are related to the method’s inapplicability when non-Gaussian uncertainties are claimed, or a candidate reference value is a variant on the strict inverse-variance weighted mean.

8.2 Chi-squared-like testing of the common mean hypothesis

Fortunately, by using Monte Carlo simulation, it is relatively easy to extend full rigor to the chi-squared style of analysis for these cases where the strict conditions for chi-squared analysis cannot be met. The common-mean hypothesis, with the subsidiary hypothesis that the candidate reference value algorithm can represent the unknown common mean, can be tested with any specified randomness, such as that of Figure 9.

By forcing the simulated distributions to have the same mean, we can emulate the PDF of the χ^2 , calculated as per Equation (22), and thence calculate the function $(1-CDF(z))$, that would be expected if the null hypothesis were to be true. Note that this does not force us to accept the null hypothesis, but we must make this assumption for the purposes of evaluating the consequences that accrue if the null hypothesis is true.

This curve is plotted in both panels of Figure 13 as the solid black line. The departures of the statistic from an analytic chi-square (with degrees of freedom = 5) are relatively minor for $\chi^2 < 3$ but get rather serious for larger

values of the reduced chi-squared. At $\chi^2 = 6$, the probability of exceeding a $\chi^2 = 6$ is some 100 times greater than would be estimated from the analytic chi-squared function.

Thus for some purposes, the established analytic chi-squared method can be shown to give results that are good enough, whereas in other cases it will only be appropriate to use the hypothesis-constrained Monte Carlo simulations' results to establish the (1-CDF(z)). Practically, the constraint about the shared mean is particularly easy to arrange when there are zero-mean random variables and setting the common mean to zero is a convenient way for applying the common-mean constraint.

Notice that if the common-mean constraint is not imposed, then Equation (22) is simply the measurement equation for an output quantity and the PDF that has been left unconstrained by the null hypothesis is simply the uncertainty distribution for χ^2 . This may even be of some use in deciding on the solidity of conclusions based on values of χ^2 .

Of course, in the hypothesis-constrained Monte Carlo simulation, *any* algorithmically defined reference value can be used for the chi-square-like testing. The Monte Carlo method can test the simple mean, the inverse-variance weighted mean with weight limits, the median, or any other reference value.

Similarly, we are not constrained by the form of the statistic. However, because the purpose of the testing is to quantify *and to communicate* confidence in the adequacy of the model, familiarity and simplicity from the perspective of the end-user will dictate a preference for chi-squared-like statistics for the foreseeable future.

If the data in Figure 4 are reinterpreted as having non-Gaussian uncertainty distributions as shown in Figure 9 the Monte Carlo calculation of the reduced chi-squared (Equation 22), shown as the black line in Figure 13, can be used to test the common-mean hypothesis, represented as the inverse-variance weighted mean, against the detailed non-Gaussian (tabulated) uncertainty distributions. If the common mean hypothesis is forced to be true, the chance of exceeding the actual reduced chi-squared of 1.060 is 35.4%. This happens to be but little different than the analytic chi-squared result, but the simulation has provided the easiest and most convincing mechanism for establishing this near equivalence. For different experimental values, giving a larger value of the reduced chi-squared, the probabilities could have been very different. In this case, there is no compelling evidence in the experiment's dispersion to reject the common mean hypothesis in the context of the claimed uncertainties. Again, note that "not rejecting" is not the same as "accepting".

8.3 Chi-squared-like testing of the agreement hypothesis

The above model is testing mediated agreement, mediated by a real value for the measurand. If this value is to be experimentally accessible and will be used as the referent for future measurements, all is well.

However, when the mediating statistic is under debate, or the reference value has not been specified or will not be used, an unmediated approach may have significant advantages. In this kind of testing, the unmediated agreement of Labs i and k , in the normalized error sense of $E_n = (i x - k x) / (u(i x - k x))$ can be used instead, as can its mean-square averaged over all distinct measurement pairs. This all-pair-difference chi-squared-like statistic is

$$\chi_{\text{APD}}^2 = \frac{1}{M(M-1)} \sum_{i=1}^M \sum_{k=1}^M \left[\frac{i x - k x}{u(i x - k x)} \right]^2. \quad (23)$$

It is sometimes an exact reduced-chi-squared having degrees of freedom equal to $M - 1$, and at other times it is only moderately close. (See the gray line in Figure 13). The point that we are finishing on here is that Monte Carlo simulation can provide a solid foundation for extending the scope of quite familiar-looking tests, allowing them to be applied before choosing a reference value or a reference value method [SHD02] [SD06a] [DS06] [SD06b].

Applying this chi-squared-like statistic, and imposing the null hypothesis of pair agreement, the probability of exceeding the experimental value of 1.017 is 39.9%. In this case, there is no compelling evidence in the experiment's dispersion to reject the pair agreement hypothesis in the context of the claimed uncertainties. Yet again, note that "not rejecting" should not be taken to have the same meaning as "accepting" or "validating".

9 Final remarks

Monte Carlo simulation has been presented here as a means of modeling the randomness of measurements, using reported values and their associated uncertainty claims as the basis for predicting the randomness of the result of a measurement equation (or algorithm). With a measurement equation, only arithmetic needs to be done. It does need to be done many, many times; it may take an everyday computer a second or so to perform enough iterations.

The random number generators can be conveniently packaged as a toolkit that minimizes the amount of reprogramming required. On the accompanying DVD, we provide and document the full toolkit that was used for the Monte Carlo simulations presented here. It is provided as a set of Excel macros that use Excel's Visual Basic for Applications programming and debugging environment. We have found these Excel workbooks to be a simple way for a beginner to adapt a Monte Carlo simulation program to a new measurement equation. An introductory tutorial is available online [NRC05] with many examples. Sharing these workbooks can also be a very effective way to communicate with colleagues, communicating both results and methods with less effort than other mechanisms we have tried. For Monte Carlo simulation, we prefer to link directly with experiments (real and gedanken) by thinking about measurements throughout the simulation process. By considering the

histograms that we can collect as *predictions* that might be compared with real experiments, we facilitate the application of the classical scientific method to measurement science and metrology: prediction, followed by experiment and comparison. We use a simple means for importing and exporting uncertainty claims expressed instead in terms of a distribution for scalar measurands. In simple circumstances Monte Carlo simulation might be done equally well from the perspective of resampling measurands, but when asymmetry is present in the distributions, special care is needed [DSWH05].

Monte Carlo simulation can only reinterpret our model (measurement equation, measurement values, and uncertainty information), and does not of itself create any really new knowledge, although it may reveal the consequences of the model in new and clearer ways. New knowledge can be obtained from the enhanced abilities to consider different models and different assumptions. We have indicated some of these directions in this chapter, but we believe that the full capabilities of Monte Carlo methods have yet to be explored in dealing with the randomness attributed to measurements and measurands.

References

- [BM58] Box, G.E.P., Muller, M.: A Note on the Generation of Random Normal Deviates. *Ann. Math. Statist.*, **29**, 610–611, (1958)
- [B07] Brown, R.G.: *The DieHarder Battery of Tests of Randomness*. (GNU Public Licence, version 2b). Available at <http://www.phy.duke.edu/~rgb/General/dieharder.php> (2007)
- [C02] Christensen, C.H.: *PRNG test*. Available at http://cholm.home.cern.ch/cholm/misc/gslmm/random_2generator-test_8cc-example.html —(2002)
- [DSWH05] Douglas, R.J., Steele, A.G., Wood, B.M., Hill, K.D.: A useful reflection. *Metrologia* **42**, L35–L39, (2005)
- [DS06] Douglas, R.J., Steele, A.G.: Pair-difference chi-squared statistics for key comparisons, *Metrologia* **43**, 89–97, (2006)
- [Ex03] <http://support.microsoft.com/kb/833855>, Issue 834520. *A discussion of the error, introduced into Excel's random number generator, by early releases of Office XP and Excel* (2003)
- [F35] Fisher, R.A.: The fiducial argument in statistical inference. *Ann. Eugenics*. **6**, 381–398, (1935)
- [GUM95] *Guide to the Expression of Uncertainty in Measurement*, International Organization for Standardization, 1995
- [GUM-S1 07] *Supplement One to the Guide to the Expression of Uncertainty in Measurement*, BIPM, 2007
- [HW82] Hill, I.D., Wichmann, B.A.: An efficient and portable pseudo-random number generator, *Appl. Statist.* **31**, 188–190, (1982)
- [J94] James, F.: Ranlux: a Fortran implementation of the high-quality pseudorandom number generator of Lüscher. *Comput. Phys. Commun.* **79**, 111–114, (1994)

- [KMR77] Kinderman, A., Monahan, J., Ramage, J.: Computer Methods for Sampling from Students t-distribution. *Math. Comput.*, **31**, 1018–1019, (1977)
- [L94] Lüscher, M.: A portable high-quality random number generator for lattice field theory simulations. *Comput. Phys. Commun.* **79**, 100–110, (1994)
- [LS02] L'Écuyer, P. Simard, R.: *TestU01: A Software Library in ANSI C for Empirical Testing of Random Number Generators*. Version 1.2.1, dated April 23, 2007. (Non-commercial redistribution licence). Available at <http://www.iro.umontreal.ca/~simardr/testu01/tu01.html>, (2002)
- [M85] Marsaglia, G.: *The Diehard Battery of Tests of Randomness*, On the Marsaglia Random Number CDROM, from Florida State University. Available at <http://www.cs.hku.hk/~diehard>, (1985)
- [NRC05] <http://inms-ienm.nrc-cnrc.gc.ca/qde/montecarlo/> (2004)
- [SD06a] Steele, A.G., Douglas, R.J.: Extending chi-squared statistics for key comparisons in metrology, *J Comput. Appl. Math.* **192**, 51–58, (2006)
- [SD06b] Steele, A.G., Douglas, R.J.: Extending E_n for measurement science, *Metrologia* **43**, S235–S243, (2006)
- [SHD02] Steele, A.G., Hill, K.D., Douglas, R.J.: Data pooling and key comparison reference values. *Metrologia* **39**, 269–277, (2002)
- [SWD05] Steele, A.G., Wood, B.M., Douglas, R.J.: Outlier rejection for the weighted-mean KCRV. *Metrologia* **42**, 32–38, (2005)
- [SWD06] Steele, A.G., Wood, B.M., Douglas, R.J.: The middle ground in key comparison analysis: revisiting the median, In *Advanced mathematical and Computational Tools in Metrology VII*, P. Ciarlini *et al.* (eds): ISBN 9812566740, pp. 245–252, (2006)
- [WH06] Wichmann, B.A., Hill, I.D.: Generating good pseudo-random numbers. *Computa. Statist. Data Anal.* **51**, 1614–1622, (2006)

Software Validation and Preventive Software Quality Assurance for Metrology

Norbert Greif, Dieter Richter

Physikalisch-Technische Bundesanstalt, Abbestrasse 2-12, 10587 Berlin, Germany
Norbert.Greif@ptb.de, Dieter.Richter@ptb.de

Summary. As everywhere in technological development, IT has become an indispensable part of metrology. Modern measuring instruments rely on embedded software. Data obtained through measurement are transferred via communication networks and processed further by software systems. This development has led to increasingly complex and distributed metrological functions. This complexity has different facets: it can not only enhance the functionality of a measuring system itself including self-checking facilities, but also the distribution and transmission of measurement data and the methods used for data analysis. One can say that IT opens up a new world of concepts that essentially contribute to the development of metrology. In this, different fields of information technology have to be taken into account, extending from software engineering, computational methods, databases, communication technology, and security techniques to knowledge-based systems. This chapter of the book gives a survey of the software validation and software engineering methods in metrology. For a better illustration, the accompanying DVD presents two case studies of validation procedures performed at PTB. Finally, an outlook on future issues is given.

Key words: Measurement software, software validation, software quality assurance

1 Introduction: The risks of software in metrology

Today, a metrological system cannot do without software technology, as this is an indispensable base. However, the improved software-based functionality and the complexity of the metrological systems not only bring about advantages, they also raise new problems and questions to be answered. By no means should we ignore the risks that are involved with software. Software of insufficient quality may lead to malfunctions with unlimited consequences. Moreover, software must always be considered to be a target for accidental or even intentional corruption. It is therefore evident that software quality assurance, including the security of the software and the measurement data, plays an important role in metrology today. The public needs to be able to trust

the software-supported functionality associated with high-quality metrological service for science, technology, medicine, consumer protection, and environmental protection.

Looking at metrology, we cannot say yet that software technology can already be regarded as a mature and commonly understood discipline of metrology. Of course, we do not ignore the many successes that have already been achieved in the area of metrological software. However, a common understanding still remains to be developed. This first section illustrates the special character of software and is particularly addressed to people who are not software professionals.

1.1 What is software

A fundamental part of a common understanding is the clarification of what software is. This question might be irritating. Everyone interested in technology naturally has an idea about software: computer programs which can do many different things such as serve the customer at the banking terminal, control the activation of an airbag, process texts, control measuring instruments, process measurement data, or check other computer programs. So we formulate our question more precisely and ask what software encompasses: the bit string on a storage unit (on a chip, etc.), the instruction sequence written in a programming language (C++, Java, etc.), the logic design for the program flow, the mathematical calculation rule, the structure of the data, the architecture of the subunits, the design of the user interface, the information necessary for operation, the documentation needed for the description, or the input and output data. Answering this question requires some thought, and the individual answers reflect a lot of different opinions.

From the point of view of software technology, the answer is clear as far as the principles are concerned: all elements listed above are part of the software. The question about the particular design must, however, be answered in concrete terms, and the answers then follow individual subjective assessments. This subjective character complicates the answers – which are difficult anyway – to questions of the following kind. How can it be ascertained whether a piece of software can be trusted? Is software A better than software B? Is a certain software package suitable for use in area X? What must be done when software is accepted to ascertain whether the required quality is complied with? How is software validated? Can software be calibrated? How can the uncertainty associated with software be determined?

It is clear that the issues associated with such questions depend on the context and it is important to pose the questions precisely and in keeping with the context. This often is the key to success. At the same time, however, this is commonly the most difficult part of problem solution because the questions are not of a purely software-technological nature but have an interdisciplinary character and may additionally be superposed by economic,

social, and other issues. Software problems can best be mastered if they can be reduced to clear software-technological questions.

1.2 Software is different

Metrology is largely founded on the physical sciences. Software has, however, introduced new types of risks, which are not known in the technical world that is shaped by physics. It is necessary to understand that quality ideas cannot be transferred blindly from one area to another. In the following, fundamental differences between physical and software engineering are presented taking up the theses of Boris Beizer [Be00], which are adjusted or further developed to fit in with the backgrounds of this chapter (see also [Ri02]).

What are the prominent differences?

Fundamentals

First, and as a principle, software is not structured in accordance with physical laws. The existing mathematical bases of information science do not imply concrete program structures but offer a wide range of structures which can generally be used by the software engineer within the scope of his or her freedom in design. This freedom in design – following well-considered principles of software engineering, but in many cases, to the chagrin of software engineers – is increasingly restricted by conventions and regulations in the form of standards, guidelines, or procedural models. These do not, however, have their roots in mathematics or natural sciences but follow comprehensive demands such as rules from quality management systems, customer's specifications, or restrictive regulations for reasons of business economics. As a consequence, there is no possibility of tracing back quality-defining software characteristics on the basis of natural science or mathematics.

Ageing

Software does not age in physical terms, whereas physical components, in the course of time, undergo changes in their properties as a result of environmental effects or wear. This is not so for software whose characteristics do not change with time. What can vary over time are its carriers or the storage units, or they can fail, which would indirectly lead to a change in the software features. Also, the environment of the software can vary, resulting in a different assessment of the utility value. Furthermore, it can be ascertained only in the course of time whether a software product is corrupt or otherwise does not meet the expectations, which also changes the assessment of the utility value. In all these cases, the software attributes originally realized remain, however, the same.

Spatial and temporal aspects

Another difference stems from the fact that physical effects are related to space and time. Experience has shown that in the physical world, the effects of an action bear a temporal and spatial relation to the action. This cannot be expected from software. As a result of the use of global data, for example, the location of the effect of a fault can be completely different from the source inducing the fault. Likewise, effects may occur at a time which cannot be foreseen. This is true, for example, for program abortion due to exceptional arithmetic phenomena.

Proportionality of cause and effect

There is also a difference as regards the relation of cause and effect. In the physical world, it is in general true that an effect is approximately proportional to the force expended or that damage is roughly proportional to the destroying force. Excepted from this are cases in which boundary values are exceeded thus leading to rupture or in nonlinear systems in which small changes in elements of the system can lead to abrupt or catastrophic changes in behavior. For software, it is not possible to delimit an area in which, for example, the effects of a software failure can be related to the kind or number of the faults. It is not even confirmed that identical software faults lead to identical effects, not even in the case of identical faults of the same software. This strongly affects the measurement uncertainty concept. Metrology depends on being able to assess the cause-and-effect relations and to estimate the residual uncertainty [GUM95] but this is not valid for the uncertainty inherent in software. The effects of unrecognized software faults can be estimated neither in quantitative nor in qualitative terms.

Negative synergy effects due to complexity

The complexity of software and the negative synergy effects related to it are phenomena deserving great attention. It is typical for software that failures are produced by more than one fault; software failures that arise from a single cause are usually detected and eliminated at the development phase. What as a rule is not discovered in the development phase are the failures induced by the concurrence of several different causes. With increasing complexity of the software systems, the relation between causes and an error state is increasingly hard to discover. So-called negative synergy effects occur. This situation affects the structure of software systems and the interaction between subsystems and touches on the management of the software development which must ensure the coordination necessary for the development. As it is difficult to fix bugs retrospectively, constructive (i.e., preventive) quality assurance gains in importance.

The increasing complexity of software systems is the result of an increasing variety of functions as well as of ever more comfortable user interfaces.

As a price for user friendliness, that is, simple interfaces guiding the user and recognizing operating errors, the size of the programs increases.

Security limits and tolerances

In software engineering the principles of security limits and tolerances are unknown. Classical engineering products such as bridge structures have a calculated and a designed loading limit. For use a safety factor is added so that the allowed load is clearly below the calculated loading limit. In production technology, tolerances are stated within which the dimensions of manufactured workpieces are allowed to lie. In software engineering analogies are not known.

Measurability

The bases for measurability and quantifiability are also entirely different. The essential difference from the physical measurement is that in software engineering the measurands are not defined on the basis of the natural sciences. What quantities are of interest and what system is defined by measurands is rather a question of subjective assessment and of reasonableness in a particular context. Software metrics can very well contribute to objectifying the software assessment and to providing comparability, provided the parties involved have clearly stated the definition, the interpretation, and the context of the application. In software engineering, different metrics systems for the quantification of software characteristics are known (cf., e.g., [Zu97]). Software metrics serve, for example, to quantify the complexity of software by measured values. The metrics systems for software are in a state of flux. Although efforts to order them have been made for a long time (cf., e.g., [KHL01]), international standards for fundamental units of software measurement are not in sight. Although software measurement indirectly affects the reliability of software, the feasibility of a standard system of fundamental units for software measurement is basically questioned [Be00].

What conclusions can be drawn so far?

- (1) Software quality assurance measures cannot start from recognized effects (i.e., observed failures), but must be planned independently and separately. The reason is that causes of failure often cannot be linked to the location and the time of failure observation. This is also true for the planning of quality assurance measures. With the increasing complexity of software systems, the detectability of cause-and-effect relationships diminishes.
- (2) The degree of the deviation from a design rule or specification (as far as it can reasonably be determined) is not indicative of the severity of the consequences associated with a failure caused by the deviation. Error-avoiding quality assurance measures thus cannot be founded on a taxonomy of the compliance with design rules and specifications. They must be based alone on the analysis of the consequences of potential failures.

- (3) The appreciation of the suitability of metrics is dependent on the context and based on experience. This is also valid for the assessment of software development and testing methods. Any change in context is likely to require a reassessment of the software.
- (4) Software failures in metrology do not fit into the existing measurement uncertainty systematics. They need their own fundament. They require a different assessment methodology.

1.3 Tackling software issues in metrology

From the organisational point of view, there are indeed various ways to handle software issues in metrology. Here we briefly describe the approach PTB has taken.

PTB has established special working groups for IT support in metrology. One of them has its focus on software testing and software process related quality assurance in metrology. Another one is especially devoted to software issues in legal metrology. In order to raise customer confidence, the first unit has been further developed to a software test centre that received in 2001 the accreditation as a software test laboratory according to ISO/IEC 17025. This accreditation confirms the competence of the software test centre, the observance of widely accepted requirements specified in international standards for test laboratories, and the comparability of the laboratory's test results with those of other software test centres throughout the world. The scope of accreditation ranges from the testing of software for functionality, reliability, security, and usability (ISO/IEC 25051/12119, ISO/IEC 9126) to the detailed testing by ergonomic criteria (ISO 9241). Test objects are measuring device software, software in measuring facilities and test assemblies, software in calibration devices, software for the processing of measurement data, and software in other devices to be approved by PTB or other state institutes (e.g., voting machines). If appropriate, software development processes can be subjected to testing.

The software test centre offers its services not only to PTB laboratories and other state institutes but cooperates with the measuring instruments industry and metrological services. The guidelines and requirements catalogues developed are made publicly available [GSS07].

2 Standards and guidelines related to software in metrology

In metrology, the amount of software used steadily grows. However, its quality differs. It is necessary, for principal reasons in metrology, to have a comparability of software quality. Unlike most metrology quality concepts, the software quality cannot be traced back to any measurement standard. That role must be taken over by software standards or guidelines [GR00].

Requirements play a fundamental role for the comparability. They have a twofold importance: they are to be observed by the manufacturer, and they are a quality criterion for the tester or user. Standards and guidelines are a supporting means for deriving particular requirements. The parties concerned can find generic as well as concrete requirements in these documents.

Unfortunately, the situation of standardisation in the field of software quality assurance is rather confusing. The situation is caused by different aspects.

- Numerous projects for drawing-up standards or guidelines have been launched by various bodies (e.g. ISO, IEC, IEEE, EA, EUROLAB).
- The lines of action followed within the bodies have not been harmonised (overlapping with respect to contents, differing terminology).
- The results of the bodies' work have reached different stages (e.g., multistage drafts of standards/guidelines, tentative standards, standards adopted).
- The scopes of the regulations differ (they are valid for different fields of software application).
- The purposes for which the standards and guidelines are intended differ (definition of terms, specification of product and/or process requirements).

In the following, a survey of the present state of software standards with a particular concern on quality issues is given. Particular attention is laid on such regulations which formulate special software requirements for testing and calibration laboratories. Drafts of standards and guidelines already published will be taken into account to some extent. In view of the problems outlined above and due to the limited space available, this section cannot claim to represent a complete overview.

For an overview of guiding documents concerning IT issues in legal metrology refer to [TG06].

2.1 Approaches to software quality and systematics of standards

The quality of software is determined by the quality of the software product, including its intermediate products, but also by the quality of the processes that belong to the software lifecycle, in particular, the design processes. Product quality is evaluated on the basis of internal product characteristics (e.g., results of static analyses of intermediate products), on the basis of external product characteristics (e.g., results of dynamic functional tests), and on the basis of special quality in use characteristics (cf. also Section 2.3.1).

Product quality is influenced by the quality of the processes involved in the manufacture of the product. Because the characteristics of processes basically differ from product characteristics, these classes of characteristics are usually treated in separate standards and guidelines. Figure 1 shows a systematics of

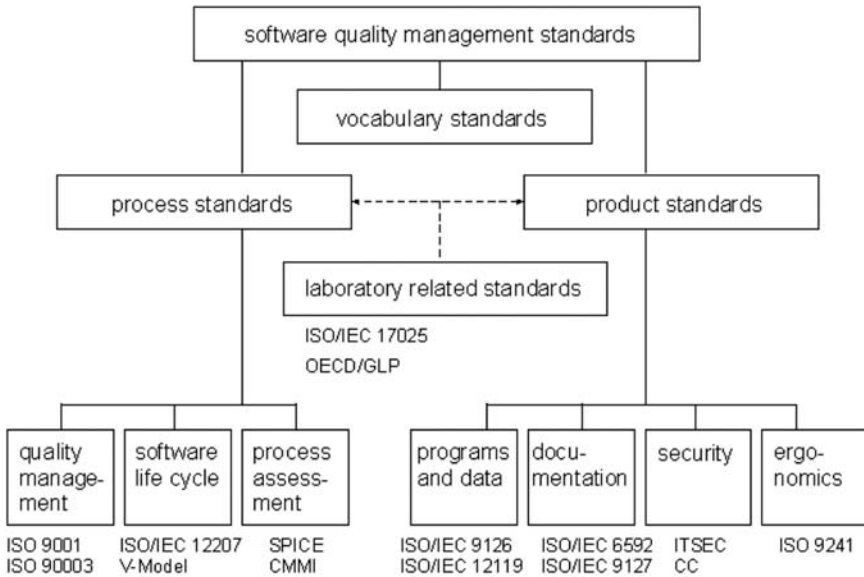


Fig. 1. Systematics of software quality management standards.

the standards concerning the field of software quality. As regards the purpose of the standard, a distinction is made above all among vocabulary standards, product standards, and process standards:

- Vocabulary standards serve the unification of terminology by providing definitions of terms.
- Product standards formulate requirements to be met by the product classes under consideration (final or intermediate products). The manufacturing process (here: software design process) is left out of consideration.
- Process standards formulate requirements for the implementation of processes (e.g., design, production, quality management processes). They describe the line of action to be followed in the processes. The spectrum of the process standards ranges from the description of framework processes (e.g., quality management, model of how to proceed in software development) to detailed rules of procedure (e.g., concrete testing instructions, auditing methods). No statements concerning the quality of the final product are made.

The separation between the definition of product and process requirements often is not consistent. Some overlap in the contents can be found, in particular, in regulations specific to a particular technical field or laboratory (cf. Figure 1). In the following, standards and guidelines concerning software quality management are dealt with according to the systematics shown in Figure 1.

2.2 Standards and guidelines concerning process quality

This section indicates and briefly explains process standards which are significant for the complexes of quality management, software development, and for the development of safety-related systems.

2.2.1 *Quality management*

The series of standards ISO 9000 ff. provides basic standards for quality management. It assists the customer in gaining confidence in the capabilities of a producer or supplier. The supplier wins the customer's confidence by introducing a quality management system and furnishing proof of its existence. Minimum requirements to be fulfilled by such a quality system are described in ISO 9001. This standard is, however, valid for a wide range of products, from material products to services. ISO 90003 has been published to take the specific features of software production into consideration. However, it is not intended to be used as assessment criteria in quality system certification. ISO 90003 only sets out guidelines to facilitate the application of ISO 9001.

2.2.2 *Software development*

The standard ISO/IEC 12207 (Information Technology - Software Life Cycle Processes) defines a well-accepted software life cycle and describes in detail all processes which play a part in the life cycle. For so-called primary processes, such as development or maintenance, but also for supporting processes, such as configuration management, or quality assurance, the hierarchy of the activities to be performed is presented in detail. The Software Lifecycle Process Model of the German Development Standard [DS97] describes all activities to be carried out and all products to be drawn up in the course of software development. With the aid of this process model, software projects can be carried out in compliance with the standard ISO 9001. The process model also covers quality assurance, configuration management, and project management. In addition to these process standards which are universally applicable there are a large number of standards which describe special software processes. Some examples are given in the following.

- ISO/IEC 15910, Information Technology – Software User Documentation Process
- ISO/IEC 18019, Software and System Engineering – Guidelines for the Design and Preparation of User Documentation for Application Software
- ISO/IEC TR 9294, Information Technology – Guidelines for the Management of Software Documentation
- ISO/IEC TR 15846, Information Technology – Software Life Cycle Processes – Configuration Management for Software
- ISO/IEC 14764, Information Technology – Software Maintenance
- ISO/IEC TR 14471, Information Technology – Software Engineering - Guidelines for the Adoption of CASE Tools

- BS 7925-1/2, Software Testing – Vocabulary/Software Component Testing
- IEEE 1008, Software Unit Testing
- IEEE 829, Software Test Documentation

Additionally, there are some regulations which deal specifically with the assessment and improvement of software processes:

- ISO/IEC 15504, Information Technology – Software Process Assessment (SPICE)
- Capability Maturity Model Integration (CMMI) [CMM02]

Due to the fact that software processes are an integral part of the overall system processes, the well-known system engineering standard ISO/IEC 15288 is often taken into account.

2.2.3 Standard for the implementation of safety-related systems: IEC 61508

The international standard IEC 61508 functional safety of electrical/electronic/programmable electronic safety-related systems is a generic standard for building up safety-critical systems. The standard deals with the total system life cycle of safety-related systems, for example, from initial concept, through design, implementation, operation, and maintenance to decommissioning. The standard comprises amongst others detailed subsets of process-oriented requirements:

- General requirements for systems
- Requirements for safety-related systems
- Software requirements
- Examples of methods for the determination of safety integrity levels
- Overview of techniques and measures to fulfil the requirements

A requirements catalogue for software processes is provided in Part 3 of the standard. In general, the standard adopts a risk-based approach. The determination of well-defined system-based safety integrity levels (SIL) and software safety integrity levels is an important element of the required risk assessment process and a base for the final selection of risk-minimising measures and validation procedures.

2.3 Standards and guidelines concerning product quality

This section deals with standards and guidelines which specify requirements for software products.

2.3.1 ISO/IEC 25000 – The SQuaRE series of standards

The ISO/IEC 25000 SQuaRE (Software Product Quality Requirements and Evaluation) series of standards covers both software product requirements

and corresponding evaluation procedures. Recommendations are given for developers, purchasers, and evaluators. Within SQaRE, a fundamental quality model for software products is described in the part ISO/IEC 25010. This model is based on the well-accepted quality model of the standard ISO/IEC 9126-1. The current version of this standard defines quality characteristics with corresponding subcharacteristics (see Figure 3) and a specific quality model for quality in use. The following aspects are taken into consideration.

- Internal product quality (via internal product characteristics such as results of static analyses)
- External product quality (via external product characteristics such as results of dynamic functional tests)
- Quality in use (via special quality in use characteristics)

Six quality characteristics for internal and external product quality (e.g., functionality, reliability, usability, maintainability), including subcharacteristics, which are also normative, are explicitly defined in the standard. Four special characteristics for assessing the quality in use are formulated in addition. These characteristics (effectiveness, productivity, safety, satisfaction) determine the degree of the overall efficiency of the software product for the user.

The quality characteristics defined in the standard are to be used for the following activities.

- Definition of software requirements (user requirements, software design objectives)
- Evaluation of software products (definition of quality criteria and testing objectives)
- Definition of user acceptance criteria for a completed software product
- Validation of the completeness of a requirements definition

The evaluation section of SQaRE is based on ISO/IEC 14598. Here, an overall process model for software evaluation is described and refined for different roles such as developers, acquirers, and evaluators.

2.3.2 The standard ISO/IEC 25051

The standard ISO/IEC 25051 (Software Engineering – Software Product Quality Requirements and Evaluation (SQaRE) – Requirements for Quality of Commercial Off-The-Shelf (COTS) Software Products and Instructions for Testing) has been drawn up to replace the older well-accepted standard ISO/IEC 12119. As a specific part of the SQaRE series of standards (see Section 2.3.1), it defines requirements to be met by COTS software packages and lays down regulations for checking compliance with these requirements. Here, software packages are software products as offered, supplied, and used, that is to say not as an intermediate or changeable product. Database programs, text processing programs, but also programs for technical and scientific functions are given as examples. Metrology software thus also comes within the

scope of this standard. The standard is intended, for example, for software manufacturers, software users, testing laboratories, and certification bodies endeavouring to stipulate requirements for software products or to evaluate software products.

The standard's software requirements are divided into the categories of documentation and programs/data. For example, requirements to be met by the documentation are:

- Each software package shall have documentation (product description and user documentation).
- General requirements for the documentation are: understandability, completeness, consistency, correctness, and ease of overview.

Requirements to be met by programs and data correspond to the content of ISO/IEC 9126-1 (see Section 2.3.1). Among them are functionality (installability, availability of functions, correctness, consistency), reliability (disaster recovery, data security, plausibility), and usability (understandability, clarity, operability).

2.3.3 Standards concerning documentation, security, and ergonomics

Although some software standards, for example, ISO/IEC 25051 or ISO/IEC 12119, also formulate requirements for documentation, security, and ergonomics, specific standards have been drawn up which concern exclusively software documentation, security, or ergonomics. Some of these standards are listed in the following. Examples of documentation standards are:

- ISO/IEC 6592, Information Technology – Guidelines for the Documentation of Computer-Based Application Systems
- ISO/IEC 9127, Software Engineering – User Documentation and Cover Information for Consumer Software Packages
- DIN 66270, Information Technology – Software Document Evaluation – Quality Characteristics

Based on these international standards, a guideline for software documentation was elaborated at PTB [GS06]. The guideline serves as a basis for quality assurance in software development, and for the evaluation of software documentation. It is also a basis for drafting and controlling of contract documents when software development commissions are placed with third parties. The proposed process and contents of software documentation are illustrated by an example from metrology.

The security of software as part of the security of information technology is checked and evaluated on the basis of special security requirements laid down in international standards. For example, security evaluation criteria are specified in the document ITSEC (Information Technology Security Evaluation Criteria) [ITS91] and in the international standard ISO/IEC 15408 (Common Criteria for Information Technology Security Evaluation). Additionally,

guidelines exist which define a uniform method for conducting evaluations, which apply security requirements and which thus contribute to achieving the mutual recognition of security evaluations. In correspondence to the requirements catalogues mentioned above, these are the documents ITSEM (Information Technology Security Evaluation Manual) [ITS92] and CEM (Common Evaluation Methodology for Information Technology Security) [CC04].

In order to ensure efficient interaction between computer users and application software, ergonomic design criteria should be taken into account. General ergonomic requirements, such as requirements for fitness for use, user friendliness, realisation of dialogues, presentation of information, and expectation conformity are defined, for example, in the standards ISO 9241 (Ergonomic Requirements for Office Work with Visual Display Terminals) and ISO 13407 (Human-centred Design Processes for Interactive Systems).

2.4 Laboratory and metrology related standards and guidelines

Standards and guidelines dealt with in the following formulate requirements for computers and in particular for software products used in testing and calibration laboratories, and specifically in metrological applications. In addition to the well-known accreditation standard ISO/IEC 17025 which contains only a few explicit software requirements, some specific guidelines are considered in closer detail. Besides a basic document of the Organisation for Economic Co-Operation and Development (OECD) [GLP95], metrology-oriented guidelines of NPL [WPB07] and NORDTEST [NOR03], and the FDA guideline [FDA02] are discussed. A guideline for the validation of metrological software according to ISO/IEC 17025 was elaborated at PTB [Gr06b] and is outlined in detail in Section 4.

2.4.1 *The standard ISO/IEC 17025*

The standard ISO/IEC 17025 (General requirements for the competence of testing and calibration laboratories) lays down the requirements a laboratory has to fulfil when it wants to be recognised as being competent for the performance of tests and calibrations. The requirements must be met for both the accreditation of a laboratory in accordance with ISO/IEC 17025 and for a conformity declaration according to ISO/IEC 17050. Both conformity assessment procedures confirm the correctness and reliability of the testing or calibration activities in a laboratory.

If a laboratory uses computers or other software-controlled systems for its testing or calibration activities, the requirements outlined in ISO/IEC 17025 for the management of the laboratory and the test methods and test equipment used also apply to the handling of computers and software and/or the software products themselves. In addition, the standard defines a number of requirements explicitly for computers, software products, and software processes. The term “software processes” covers thereby the whole range of the

software life cycle, that is, all activities, starting from contract design to purchase, development, testing, and maintenance.

In contrast to the OECD document [GLP95], the standard does not specifically refer to the use of computers and software in laboratories. It rather describes management requirements and general technical requirements (e.g., for staff competence, accommodation, and environmental conditions, testing and calibration methods). Only a few of the requirements can be applied to software as a means of measurement. They essentially concern the quality characteristics of functionality and security as well as the software documentation. For example, the laboratory shall ensure that software is documented in sufficient detail (user manual, installation guide) and suitably validated or otherwise checked as being adequate for use. Protection against unauthorised access to data and protection of data integrity in case of input, storage, transmission, and processing of data are requirements of the category data and program security. Calculations and data transfer shall be checked in an expedient and systematic way. Original data shall not be lost or changed. Documents generated by software (surveys, reports, charts, etc.) shall be uniquely identified. For further details see Section 4.

2.4.2 The OECD consensus document

On the basis of the OECD principles laid down in the document OECD Series on Principles of Good Laboratory Practice (GLP) and Compliance Monitoring, the so-called GLP consensus document ‘The Application of the Principles of GLP to Computerized Systems’ has been drawn up [GLP95]. The document formulates requirements which computers, connected hardware, and installed software must meet when they are used in laboratories within the framework of tests, approvals, or the like. The document says literally: ‘All computerised systems used for the generation, measurement or assessment of data intended for regulatory submission should be developed, validated, operated and maintained in ways which are compliant with the GLP principles.’

With reference to the main sections of the GLP principles, namely responsibilities, training, facilities and equipment, maintenance and disaster recovery, data, security, validation of computerised systems, documentation, and archives, the requirements for both the software development processes and software products are formulated.

Process requirements

The application of acknowledged rules of science and technology, acknowledged technical and quality standards (e.g., ISO, IEC, IEEE), and documented procedures is explicitly required, especially for processes covering the introduction of computerised systems, software development (design, programming, testing, documenting), and for processes covering the validation, operation,

and maintenance of software. The software should be developed in compliance with acknowledged software development standards (e.g., the Software Lifecycle Process Model of [DS97]) and documented procedures (e.g., guidelines for software documentation or programming directives). The application of acknowledged procedures allows in particular the software manufacturer's reliability and thus, indirectly, the product quality to be assessed.

Product requirements

The requirements for software products specified in the chapters of the consensus document are compiled here and assigned to the requirement categories of security, identification, functionality, reliability, and documentation.

Requirements with regard to the security of data and programs generally rank high in the OECD consensus document. In the section about data and program security, documented procedures for the protection of data and programs against unauthorised (intentional or unintentional) access (confidentiality), and unauthorised modification, falsification, and loss (data integrity) are required. The following means/actions are suggested: unambiguous user identification and multistage protection by password(s), 'electronic' signature with date/time and reason for modification so that all modifications of data can be traced back to persons, file verification routines (use of checksum-protected binary data files), and plausibility checks (recording of nonplausible values). Archiving procedures must be uniformly applied to all data types. The unique identification of programs, subparts, and versions is required.

The laboratory has to ensure that only validated program versions are used. The availability of adequate documentation is required, including program documentation and program development documentation. A test documentation (test plan, test procedures, test data, test results) is especially required.

In the scope of functionality, the traceability of the measurement results and of all data modifications as well as the archiving of all raw data and supporting information according to the expected life is required. To achieve reliability, documented procedures for disaster recovery to be taken in the event of partial or total failure (backup copies of data and programs) and documented procedures for preventive maintenance and fault repair are required.

2.4.3 NPL Best Practice Guide No. 1

The NPL guidance document validation of software in measurement systems [WPB07] addresses validation of metrology-related software both from the perspective of the user and the supplier. Current best practices in software quality assurance are surveyed and applied to measurement systems. Especially, validation aspects within safety-related systems by reference to the generic safety standard IEC 61508 are considered. The contents of the guide comprise a management overview and a specific technical part. The

part covering management aspects defines a framework for the overall quality assurance approach based on software risk assessment. The applied risk analysis corresponds to the standard IEC 61508. The technical part of the guide provides recommendations for specific validation techniques. Based on risk assessments, well-defined software risk factors are determined. As a function of these risk factors, a so-called “measurement software index” is computed which serves as a base to determine appropriate validation techniques.

2.4.4 NORDTEST Software Validation Method

The NORDTEST software validation guideline [NOR03] provides assistance in validation of small and medium-sized software used in accredited and other testing and calibration laboratories. The proposed validation method is implemented as a software tool to be used for systematic computer-aided software validation. The text processing tool recommends a working strategy based on a common software life cycle model. The software requirements used in the validation process are mainly stated in the standard ISO/IEC 17025. Furthermore, the guideline ISO 90003 is taken into consideration for requirements specification. The proposed method can be used to validate the following.

- Purchased software products that are not standard or configurable software packages
- Self-developed or purchased software products where the source code is available
- Software products being developed in control of the laboratory

A template to elaborate the validation report summarising all validation activities is suggested.

2.4.5 FDA Guide: General Principles of Software Validation

The FDA guidance document [FDA02] describes the general context for software validation (terminology, quality assurance approaches) and the fundamental principles of validation in the scope of medical device software. The validation principles shall provide helpful measures to improve the design, development, and manufacturing processes of medical devices. System and software requirements are taken from the FDA medical device quality system regulation. The guidance addresses the complete software life cycle and all software components that influence the final software-controlled medical device. Furthermore, risk management activities are recommended to qualify the validation process. The proposed guidance applies to the following kinds of software.

- Software used as a component, part, or accessory of a medical device
- Software that is itself a medical device (e.g. blood establishment software)
- Software used in the production of a device (e.g., programmable logic controllers in manufacturing equipment)

- Software used in implementation of the device manufacturer's quality system (e.g., software that records and maintains the device history record)

For software developers, the guide proposes to establish a complete software life cycle model that is appropriate for their products. Software validation is seen as an important part within the applied life cycle model. Especially, repeated validations in case of software changes are emphasised.

2.5 Future work

The standards and guidelines presented define requirements to be met by software products and software processes. From the viewpoint of metrology, the formulation of the majority of the requirements is, however, a very general one. Application of the requirements to metrological software requires concretisation. One important objective is the definition of special metrological software quality characteristics. This would allow the relevant requirement catalogues to be derived more easily for different kinds of software tests (type approval procedures, testing of measuring assemblies) and for different measures of preventive quality assurance.

Quality criteria which are easily applicable are also advantageous in the development and purchase of metrological software. Reference to metrological aspects is concretely made in the standard ISO/IEC 17025, in the NPL, PTB (see Section 4), and NORDTEST guidelines, and in the GLP principles. In contrast to the general standards such as ISO/IEC 9126-1, these documents take the special conditions in testing and calibration laboratories into account. In view of the historical development and bearing in mind the respective purpose, the standards and guidelines referred to above emphasize different features. The emphasis on security aspects – for example, protection against unauthorised access and data tampering – is striking in the documents, in particular in the GLP document. Requirements for the quality of the processes are also taken into account to an ever-increasing extent.

In view of the great variety of standards concerning software quality assurance and the necessity of defining special metrological quality characteristics, the development of a framework document for the application of software quality standards in the field of metrology would be a forceful way towards quality criteria which are easily applicable. By serving as a guide to the derivation and application of requirements for metrological software, such a document might bridge the gap between software specialists and metrologists.

3 Analytical and preventive software quality assurance

For two reasons, the use of software-controlled systems will always entail a factor of risk: implementation errors can never be completely ruled out [Be90], and on the other hand, software solutions will always provide a target for intentional or unintentional manipulation. To maintain consumer confidence in

the correct and reliable functioning of the software in measuring systems, the existing risk has to be minimised. Today this is mainly achieved by different kinds of conformity assessment techniques:

- *Product testing*: The final software product (or an intermediate product) is thoroughly tested by an independent test centre. In analogy to the general inspection carried out for motor vehicles, the independent body certifies – if the test result is positive – that a certain software product complies with the relevant quality standards, and that it is suited for its intended use.
- *Process assessment and improvement*: The software development process is regularly accompanied by preventive quality assurance measures (audits) already at the producer's. If the result of the audits is positive, the independent auditor certifies that the producer's software development process, including all its subprocesses, complies with the relevant quality standards. If the outcome of the audit is negative, suitable suggestions for improvement are derived from the audit results.
- *Producer's declaration*: The software producer himself declares compliance of his software products or his software development processes with the technical and quality standards applicable.

3.1 Product testing versus preventive quality assurance

Software quality assurance consists of (a) the analytical testing of software products and (b) the assessment and improvement of the software development processes as preventive action to be taken as early as possible in the development phase. These two components are inseparable, and they supplement each other. The quality of software is determined by the quality of the final software product and its intermediate products, and is also a result of the quality of the underlying processes of the software life cycle, especially of the development processes (see Figure 2). Not only is the quality of software evaluated by its internal characteristics (e.g., the results of static analyses), its external characteristics (e.g., the results of dynamic function tests), and its special quality in use characteristics (see also Section 3.3), but also it is directly influenced by the quality of the different processes contributing to its development. For the assessment and permanent improvement of development processes by means of special process characteristics, models such as ISO/IEC 15504 (SPICE) or CMMI [CMM02] are available and already in use worldwide (see Section 3.4). In the long run, high software quality can be achieved only if the underlying software development processes are also of high quality. The logical consequence of this is that analytical and preventive aspects of quality assurance have to be linked and directed towards the common objective of software quality [Gr06a]. The relationship between testing and preventive quality assurance also becomes clear when we look at the application of quality requirements (see Section 3.2). The requirements which will finally be applied by the test engineer as quality criteria for the software test must be defined as

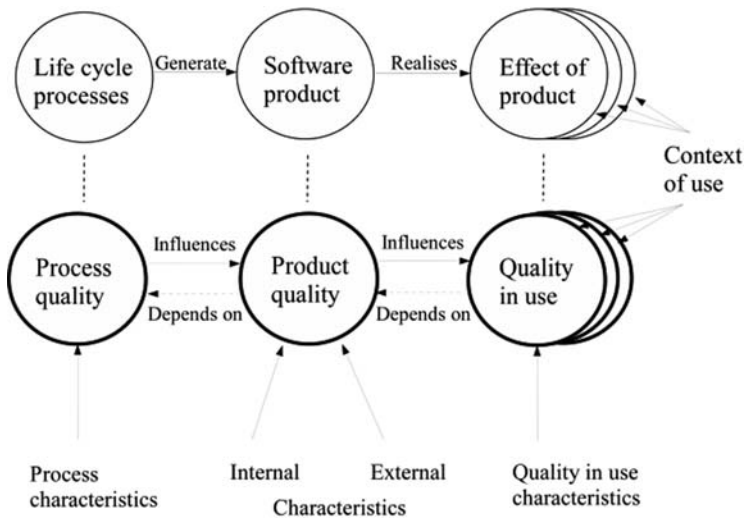


Fig. 2. Product- and process-oriented aspects of software quality.

quality objectives already in the design and development phases as it is not possible to ‘test quality into a product’ once it is completed.

The requirements applying to a software product and its development processes should be known not only to the test engineer but also to the device manufacturer and to the software producer. Efforts should therefore be made to provide them with the necessary information, ideally in the form of checklists, guidelines, and process models. If these guidelines are applied at an early stage, it will have a positive impact both on the manufacturing process and on the quality of the software development process, and finally lead to an improvement of the software product itself. The aim of this approach is to strengthen customer confidence in the producer, to enhance product testability, to facilitate explicit software testing and render it more cost-effective, and to replace final product testing in part or completely by analysing and assessing the software development process right from the beginning by audits.

3.2 Definition of testable requirements

It is a precondition for assessing the quality of software used in measuring systems that testable requirements are defined for the software products and their development processes [SS97, Gr04, GS00]. These requirements, which must cover both metrological and software-related aspects, will serve the software producer as target functions in the development of the product and the test engineer or software user as quality criteria. Unfortunately, at the beginning of a test cycle, testable requirements are hardly ever in place and the producer’s ideas about the objective of the test usually rather vague. All the producer generally wants is the product to be certified as being ‘of high

quality,' or as "functioning correctly." For the test engineer, however, these formulations are much too imprecise. He needs detailed test specifications and defined targets. Each time a product is tested, concrete test objectives are therefore determined beforehand in close cooperation with the producer, and a catalogue of specific technical requirements is compiled.

For the definition of requirements, existing standards and guidelines can be used as a basis. It is, however, to be noted that the state of standardisation in the field of software quality assessment is not very clear.

A survey of the situation is given in Section 2 with particular reference to standards and regulations containing special software requirements for testing and calibration laboratories.

As an example, Figure 3 shows the quality characteristics for software products according to ISO/IEC 9126-1. For the specific problems encountered in metrology, most of the software requirements specified in standards and guidelines are formulated in a much too general way. Before they can be applied to metrological software, they have to be tailored for their intended use. Only in a few regulations, some special needs of metrology are taken into account, for example, in the GLP Principles of the OECD [GLP95] and in the NPL and PTB guidelines (see Section 2.4). Unlike general standards such as ISO/IEC 9126, the standard ISO/IEC 17025 and the PTB guideline (see Section 4) also account for the special conditions in testing and calibration laboratories. In certain cases, these documents can therefore be directly used for the testing of measuring software. Special features of

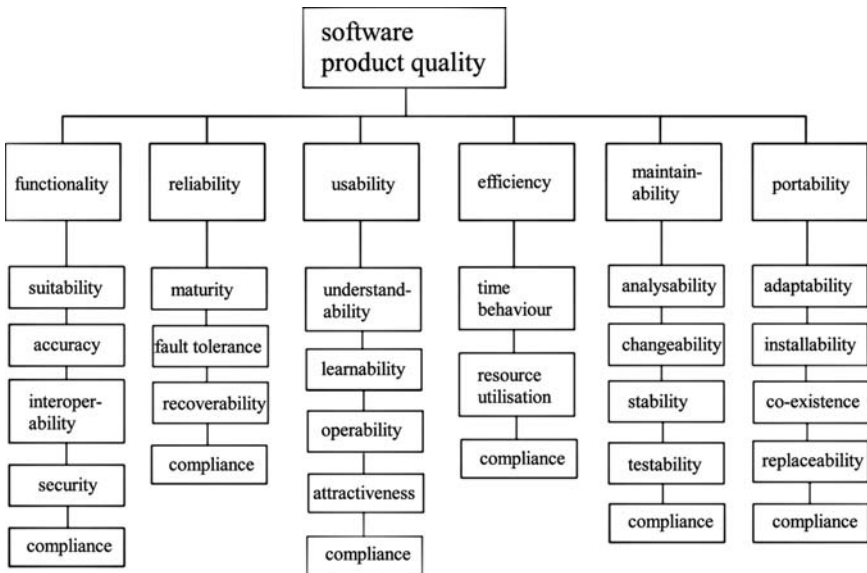


Fig. 3. Software quality characteristics according to ISO/IEC 9126-1.

legal metrology are considered in the WELMEC-Guide 7.2 [WEL05]. Based on the MID [MID04], concrete requirements and validation guidance are provided for measuring systems subject to legal verification.

Usually, however, the requirements must be separately drawn up and defined for the different classes of measuring instruments based on the standards and guidelines to which they refer (see [GR00]). The time and work spent on drawing up the requirements catalogues would take up a large amount of the total test work. To reduce this amount, the PTB Software Test Centre has prepared a series of catalogues covering different types of requirements and allowing a software test to be carried out from different aspects. The following catalogues are available.

- General nonfunctional software requirements
- Requirements for software used as testing equipment
- Ergonomic software requirements
- Requirements for the use of software-controlled systems in testing and calibration laboratories
- Checklist for auditing a software producer

The catalogues are permanently updated. Also, new catalogues for further aspects are under preparation. Detailed requirements lists of the various catalogues have been compiled in [GS00].

From the requirements catalogues and the applicable standards and guidelines, concrete requirements for special device classes can be derived. For a weighing instrument, for example, the requirement ‘functionality’ means that the mass of a weight must be correctly indicated (within the limits of the declared measurement uncertainty). For the software of this weighing instrument, this implies that, for example, the zero position of the instrument (the offset of the weighed value) is determined in a self-check and stored in a suitably dimensioned program variable.

For example, the Software Test Centre of the PTB has worked out catalogues of concrete requirements for the following fields of application: calculation of measurement uncertainties [GSR06], calibration of gauge blocks (length measuring technique) [GSR99], and calculation of the radiation dose for flight attendants. One objective of the future work will be the definition of specific metrological quality criteria for further classes of test applications. Such criteria would allow concrete requirements for different types of software tests and for different instrument classes to be derived more easily. It is also to be noted here that reasonable software requirements are also needed for preventive quality assurance and for the development or purchase of measurement software.

3.3 Software product testing

The objective of a software product test can vary strongly from case to case, with the test method varying accordingly. As a consequence different methods

and procedures are applied in the test. Software tests can be distinguished according to the following levels.

- With regard to the different software components, a distinction can be made among the testing of the executable program (including the user interface); the testing of the source code, the testing of the documentation, and, where appropriate, the testing of data (see Figure 4).
- With regard to the instantaneous position of a software product within the software life cycle, we can distinguish between the testing of the final product and the testing of intermediate products. Apart from the testing of the final product (which is described below in more detail), the scope of testing also includes the testing of products in the early stages in the software life cycle, for example, reviews of requirement and design documents. The test of a final software product against the original user requirements is called software validation.
- With regard to the testing techniques applied, a distinction can be made as follows (see Figure 5) [Be90, Be95, Me79]:
 - Manual inspections (e.g., evaluation of program source codes, documentation, requirements specification, design documents)
 - Tool-supported program function tests (systematic execution of programs including actual/setpoint value comparison, testing of user interfaces)
 - Structural tests (e.g., tests based on data/control flow, determination of test coverage)
 - Tool-supported static analyses of the source code (e.g., check whether programming rules are observed, data flow analyses)
- With regard to the technical objective of software tests, distinctions that can be made include:
 - Test of functionality
 - Test of conformity with standards or normative regulations

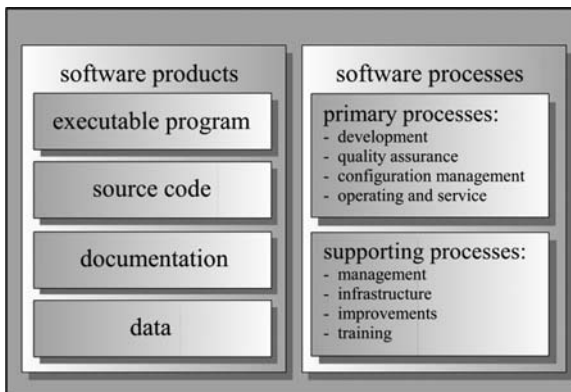


Fig. 4. Components of the software under test.

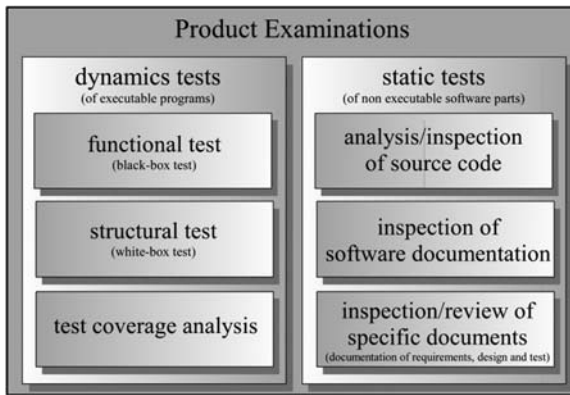


Fig. 5. Techniques of software product evaluation.

- Test of the programming technique
- Test of ergonomic characteristics
- Test for data and program security
- Test for functional safety

Although not complete, the bullets regarding the test objectives serve as a guideline for the explanations below.

3.3.1 Functionality test

All functions of a program are checked for compliance with an individual set of requirements. To test the functionality, a specification, a list, or catalogue of requirements must be available or at least precise ideas about the required or desired scope of functions. Experience has shown that the formulation of such requirements is often imprecise (see Section 3.2). As testing techniques, mainly functional and structural tests are applied [Be90, Be95, Me79]. Code inspections, too, are of some help.

3.3.2 Conformity test

The software is tested for conformity with software-specific or domain-specific standards or normative regulations. Tests are carried out, for example, with regard to the following regulations.

- ISO/IEC 25051/12119: Testing whether a software is suitable as COTS software product.
- ISO 9241: Testing whether a software is fit for its intended use (see further below).
- OECD Principles of Good Laboratory Practice [GLP95] and NPL/PTB Guidelines (see Section 2.4): Testing of compliance with regulations prescribed or recommended for the use of computers and software in (accredited) laboratories or in metrology-related applications.

If necessary, compliance with other software- and domain-specific standards or guidelines for laboratories is tested. Testing techniques frequently applied are functional and structural tests as well as the inspection of program source code and documentation.

3.3.3 Test of the programming technique

By testing the quality of the programming technique, evidence is to be given that the software works in a robust and reliable manner. For this code-specific test of non-functional requirements, code inspections and tool-supported static analyses, especially data flow analyses [GS99], are used. The following questions can be answered by this category of test.

- Has the software been implemented according to the state of the art, and has the software complied with any specified programming rules?
For example, the following error categories can be detected by the tests: initialisation problems, memory access problems, and side effects.
- Are the data processed correctly from the moment of their input until the moment of their output? Is it possible to reconstruct the relationship between output value (e.g., price of the product indicated on the weighing instrument) and input values (sensor value, measurement range) by means of the program source code?

To answer the questions of the last bullet, methods of data flow analysis [GS99] are invaluable. This special method can, for example, be used to test whether security against software manipulation in measuring systems subject to legal control is assured. Software which has to be protected against tampering is found, for example, in electronic weighing instruments and in fuel dispensers at petrol stations. To protect the software, it cannot, however, be ‘sealed’ completely. From time to time, it must be possible for the shop assistant or service-station attendant to adjust the price of his products per kilogram or litre. For this purpose, he must have limited access to certain data storage areas of the software. He is not, however, supposed to influence the formation of the weight values, for example. When the weighing instruments are type-approved, evidence must be provided that those components of the software which are subject to legal control and form the measurement value are not influenced by the price adjustments made, as is permitted, by the shop assistant and that no criminal manipulation of the software is possible or at least it is not possible by simple means. For this purpose, extensive software-specific tests have to be carried out based on the analysis of the program data flow.

This basic understanding of securing software can be transferred to types of instruments, not only to those ones that are subject to legal verification.

3.3.4 Test of ergonomic characteristics

Here, the usability and the fitness for the intended purpose of a software product according to the international standard ISO 9241-10/11 is tested for

a specific context of use. Essential criteria to evaluate how the interaction between user and software is designed and implemented are as follows.

- *Suitability for the task:* The user shall be supported in such a way that he can accomplish his task effectively and efficiently.
- *Self-descriptiveness:* All dialog steps shall be immediately understandable and explained on request.
- *Controllability:* The user shall be able to start the dialog and to influence its direction and speed until he reaches the objective.
- *Conformity with user expectations:* The dialog shall be consistent and tailored to the knowledge of the user.
- *Fault tolerance:* Even if the user makes inputs which obviously are erroneous, he shall be able to achieve the intended objective with only very little additional effort.
- *Suitability for individualisation:* The program shall be so designed that it can be adjusted to the requirements of the task and to the individual capabilities and preferences of the user.
- *Suitability for learning:* The user shall be supported and guided in learning the program dialogs.

The test of ergonomic characteristics includes a functionality test according to the above-mentioned ISO/IEC 25051/12119.

3.3.5 Security test

The testing and evaluation of software security within the scope of information technology security is carried out according to special security criteria laid down in international regulations, such as the following.

- ISO/IEC 15408, Common Criteria for Information Technology Security Evaluation.
- ITSEC: Information Technology Security Evaluation Criteria [ITS91],

Conformity tests are mainly carried out for the following quality characteristics.

- *Availability:* Data and services must at any time be available to authorised users.
- *Confidentiality:* Information shall be available to authorised users only.
- *Integrity:* Data and programs must be protected from unintended or unauthorised modifications (including complete loss).
- *Authenticity:* Programs must clearly identify the communication partner (user, process) of protected transactions.

In particular, the existence and efficiency of security measures is tested. Such measures can be as follows.

- Controls of access to programs, password systems
- Restrictions of access to certain data storage areas or documents

- Role concepts, graded granting of rights
- Measures for antivirus protection, firewalls
- Plausibility checks for all input data
- Database integrity rules
- Archiving and backup measures
- Disaster recovery

3.4 Execution of software product tests

To warrant the repeatability and comparability of tests, all requirements, test data, and methods used in the tests must be well defined and documented. Also, the organisational sequences which are repeated in each test have to be laid down as working instructions. For example, the following questions must be clarified in advance for each test.

- What requirements are to be met by the software to be tested (see Section 3.2)?
- Which criticality (risk/integrity level) is allocated to the software system, to its components, and, if appropriate, to each individual requirement to be checked?
- What normative documents (standards, guidelines, regulations) are to be observed?
- Which test concept or testing procedure is used (testing depth, selection of test data and methods, selection and adjustment of procedures, analyses, of testing results)?

Figure 6 shows the basic procedure and the essential elements of a software validation process.

3.5 Process assessment and improvement

It is the precondition for high quality of a software product that its overall development process and all its subprocesses are also of high quality. This is achieved by assessing the overall process and its subprocesses and improving them on the basis of the assessment results. The assessment of processes and their improvement are cyclical elements.

Formal or informal assessments of software development processes may be applied in the form of audits at the software producer's site. Formal assessment systems are offered by the International Standards series ISO/IEC 15504 (SPICE) or the assessment models CMMI (Capability Maturity Model Integration) [CMM02]. On the other hand, in informal assessments, made-to-measure lists of questions based on relevant process standards (such as ISO/IEC 12207) can be compiled and applied. If the answers to the audit questions are positive and if the overall assessment of the development process is also positive, this can be regarded as an important indication that the producer's software is in compliance with the latest state of the art. All in

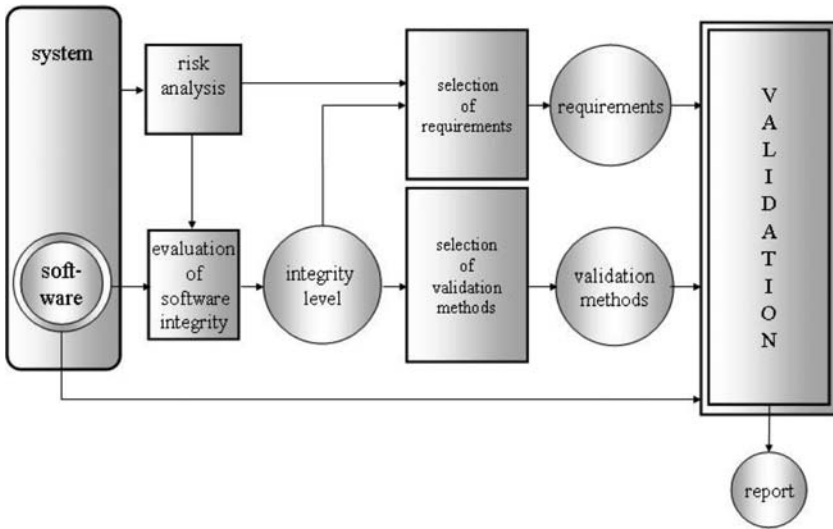


Fig. 6. Schematic representation of the software validation process.

all it can be assumed that in the case of an overall positive judgement, good testability and maintainability of the software product is ensured.

To illustrate the procedure and thoroughness of an audit, some questions are given in the following.

3.5.1 Audit – Software testing:

- Are adequate software testing procedures used in the project (test planning, test case determination, test data generation, performance of tests, test analysis, test documentation)?
- Are the methods used documented?
- Who carries out the tests? Is there a separate group for this?
- Are reviews of the results of the different stages carried out at regular intervals (e.g., review of the requirements specification, design documents, test schedules)?
- Who draws up the test schedules and by what means?
- Who carries out the functional test?
- Are there test records and summarising test reports?
- Is there a procedure to record and archive the test results?
- By what means is it assured that all user requirements and product functions have been checked?
- Who authorises test schedules, test records, test reports?
- Are confidence-building internal audits of the software development processes carried out?
- Can the producer guarantee access to the test environment (testing tools, test data) and to the test documentation?

3.5.2 Audit – Software documentation:

- Are there any company-specific standards or guidelines for the preparation and maintenance of the software documentation?
- Is the documentation produced parallel to the project?
- Is the documentation subject to regular reviews?
- What is the scope of the documentation supplied?
- For what time is the documentation valid?
- Is the documentation subject to version management?
- Can the producer guarantee access to the software development documentation?
- Can the producer guarantee access to the source code (if necessary)?

3.5.3 Audit – Error messages and change management:

- Is there a formal problem message procedure with feedback to the customer or to the test team?
- How does the producer react to error messages and suggestions for improvement and how does he deal with them?
- In what way are the customers or the test team informed about how the matter is handled?
- Are error statistics (error types, frequencies) compiled?
- Who initiates software modifications?
- Who approves software modifications?
- Are there documented methods for change and version management?
- Do the procedures of the change management also apply (apart from programs) to the software documentation and other documents such as test schedules, test reports, or design schedules?
- Are new software and document versions systematically identified?
- How are the customers informed about new versions?
- Is a configuration management tool used?
- Are there any provisions as to which tests have to be repeated in case a version is modified, and in which way they have to be carried out?

If necessary, these audit questions have to be extended or refined to allow producer- or domain-specific checklists to be derived.

Finally, the quality of the software development processes can be improved by the following means.

- The results of the audits carried out at the producer's site are directly put into practice (this has a direct impact on the processes and, consequently, on the products).
- Use of auxiliary means in software development: standards, working instructions, checklists, process models, guidelines (e.g., for design, programming, documentation, testing).
- Use of validated software tools (compiler, configuration management).
- Use of validated requirements catalogues.
- Reuse of high-quality software or software components.

3.6 Future work

Main elements of conformity assessment are testing and certification of products or processes, the producer's declaration, and accreditation. For these elements, two trends become apparent in software quality assurance which are of importance for metrology in particular.

The first trend is the increasing orientation towards preventive quality assurance for the development processes. Extensive time- and labour-consuming product tests are simplified or replaced by process audits, producer's declarations, or the testing of samples. In the field of legal metrology, this trend is accounted for by the introduction of specific conformity assessment modules H and H1 within the framework of the MID [MID04]. In future it will be a task to support this shift of emphasis by organisational and technical means. Besides, it is necessary to work out instructions for process audits at the producer's site. For areas requiring a high degree of measurement correctness or running an increased risk of manipulation, it is advisable to carry out systematic tests of intermediate products based on specific risk analyses.

The second trend regards product tests. There is an increasing demand that these tests should be repeatable and comparable. This would ensure the traceability of the test results and the objectivity of the evaluation. This can, however, be achieved only if the tests are carried out on the basis of detailed working instructions. During the test, relevant intermediate results, all final results, and the respective test environment (test engineer, hardware, software, etc.) are to be documented. Accredited software test centres must meet these requirements. In future, greater attention will have to be paid to the reliability of product testing. For this purpose, studies are necessary in the field of test coverage and software metrics in conjunction with the particular needs of metrology. The following questions will have to be answered more profoundly for metrology. When can the test be terminated? Is the entire scope of the program covered by test cases? Which test cases are still missing in which program components and have therefore to be carried out subsequently? Which software metrics help evaluate the software product and the test results?

4 Software validation in metrology

The validation of methods used for measurement and testing is crucial for the quality assurance of metrological processes. For software validation in metrology, however, there is not yet any overall guidance material or textbook available. Examples where selected aspects of software validation in the metrological context are dealt with are [GSS07], [WPB07], [Ri06], and [Ja06].

Validation is an essential aim of the standard ISO/IEC 17025. This standard also provides the conceptual basis for the validation of software. In this

section special interpretations of the standard for software are outlined in Section 4.1, testable requirements for software products and for the handling of software in laboratories are systematically derived in Section 4.2, and methods of software validation are presented in Section 4.3.

4.1 Software-oriented interpretation of the standard ISO/IEC 17025

An essential requirement of the standard is the necessity to validate laboratory-developed or changed test and calibration methods. If these methods are based on software or are supported by it, the software, too, must be subjected to the required validation. The basis of the software validation is the definition of testable software requirements.

Interpreting the general requirements of ISO/IEC 17025 with special regard to software, it is assumed that software programs used in the laboratory can implement or support both the items ‘equipment’ and ‘method’ according to the terminology of the standard. Therefore, the requirements for both – ‘equipment’ and ‘methods’ – are considered in the interpretation process. Altogether, the aspects this standard contains with regard to software products or software processes are extremely multifaceted. In the following, we have listed those aspects of the standard which are taken into account for the derivation of software requirements.

- Equipment used for measurements, testing, and/or calibrations
- Methods and procedures (measurements, tests, calibrations, data analyses)
- Handling of test and calibration data (control, archiving, etc.)
- Handling of electronic documents and records (technical and management scope)

For different reasons, extracting the requirements that apply to software products and the use of software from the text of ISO/IEC 17025 turns out to be a difficult task. First of all, the standard is, of course, not structured according to the needs of a software-oriented system. This means that the requirements are often ‘hidden’ or mentioned several times in different contexts. Secondly, many requirements are of a very general nature and were defined without taking the special software aspects into account. Their application to software products and processes leaves some room for interpretation.

4.2 Systematisation of software requirements

When interpreting the requirements of the standard with special regard to software, it is assumed that software programs used in the laboratory can implement or support both the ‘equipment’ and the ‘methods’ of the laboratory. In the systematisation of the software requirements, therefore, not only the requirements for the equipment of the laboratory will be taken into account

but also those for the relevant laboratory processes. For each individual application within the scope of a software validation, however, the total number of requirements derived for software must be reduced to the corresponding amount of requirements relevant for the special case. Not every requirement – and by no means the total number of requirements – is relevant for a specific validation case.

The total amount of requirements is subdivided into two fields. The first field encompasses the requirements for software products. The term software product comprises the executable computer program itself, but also the software documentation with all its variants, the program source code, and any associated data or databases.

The second field encompasses the requirements for the management of software. Thereby, the complete life cycle of the software, from its development to its validation and maintenance, is taken into account.

4.2.1 Requirements for software products

This subsection describes the requirements which a testing or calibration laboratory must fulfil when it uses software programs for testing or calibration within the scope of its facilities or procedures.

By the term ‘software’, we understand software as being a component of a measuring chain. It receives measurement values, measurement parameters, or intermediate results from the hardware, from users or from data storage media, processes them, and issues processed data or measurement results.

The input data of a program are measurement values (raw data), measurement parameters, and configuration data (facility parameters). The data it issues are the measurement results. If the program does not cover the whole measuring chain, intermediate results are entered or issued instead of measurement values and results.

Measurement values influence the measurement result directly and are newly acquired several times for each measurement. Measurement parameters influence the measurement result directly or indirectly and are usually entered or determined once for each measurement. Configuration data also influence the measurement result, but they are modified only rarely or never. Thus, configuration data can be logged and/or administrated independently of the measurements, whereas measurement values and measurement parameters must be logged for each measurement.

In the following we list all the requirements derived from ISO/IEC 17025. They are ordered in groups of similar subjects and numbered for easy reference.

Software identification

(P1) The software must be identifiable. The identification must consist of the name of the program, the version number and the author/contact person, and it must be shown at the beginning of the program or whenever required.

Functional scope, fitness for purpose, and conformity

(P2) The software must allow the input, logging and, if necessary, storage of measurement values and measurement parameters as well as the calculation, logging and, if necessary, storage of measurement results.

(P3) The software must comprise all functions that are necessary for the determination of the measurement result.

(P4) The software must not generate any numerical errors which reach the order of magnitude of the measurement uncertainty.

(P5) The implemented algorithms must as far as possible correspond to the specification. If approximations are used, estimations or test series must prove how large the difference is to the exact solution. The difference must be taken into account when determining the measurement uncertainty. The approximate solution methods must be described.

(P6) The values supplied by the measurement hardware must be processed as specified.

Plausibility checks

(P7) The software must survey the status of the measurement hardware. The values supplied by the hardware must be checked for plausibility. If the measurement values are implausible or show unexpectedly large variation, the user must be notified.

(P8) The values entered by the user must be checked for plausibility.

Confidentiality and integrity of data

(P9) It must not be possible to alter the configuration data by simple means. The user must agree to an alteration before it is carried out. Alterations of the configuration data must be logged together with the date.

(P10) Admissible deviations, variation ranges, plausibility limits etc. must either be implemented firmly or must only become effective when the user has agreed to any alterations.

(P11) After the measurement has started, it must no longer be possible to alter measurement parameters and other values entered by the user.

(P12) The deletion of data must either be reversible or may only become effective if the user has agreed to it.

(P13) It must not be possible for the user to alter the values supplied by the measurement hardware. If they can be altered, the modifications must be registered.

(P14) If there is a manual mode for measurement values which are normally captured automatically, or if there is a learning or demonstration mode, the values captured in this way must clearly be marked as such.

(P15) If the software rejects or re-formats entered data, a display (if it exists) has to be updated immediately.

(P16) For each number which has to be entered and for each number which has to be indicated, stored or printed out, it must become clear which unit of measurement is valid for it.

(P17) Special technical terms must be used in the user interface and in the documentation in a clear and consistent way and in compliance with the standard.

(P18) If the configuration data are stored as a file, the file must be protected against unintentional alterations.

(P19) If the measurement values, measurement parameters and/or measurement results are stored as a file, the file must be protected against unintentional alterations.

(P20) The program must not store or issue measurement values if the measurement was not complete or has been disrupted.

Documentation

(P21) There must be documentation available which has to be valid for the software version to be put into operation.

(P22) The documentation must contain the following information: name of the program, version(s) for which the documentation is valid, date of issue, author/contact person, purpose of the program or brief introduction, description of the main functions of the program, list of all error messages along with instructions for the further recommended action, system requirements, constraints or limitations and a description of the file formats, page numbering, total number of pages or a mark showing the end of the document.

(P23) The program and the documentation must related to each other.

Logging and archiving

(P24) For each complete measurement and each measurement not having been rejected, the software must issue a measurement protocol. The protocol must contain all measurement values and measurement parameters captured by the program, including the date.

(P25) The software must allow the recording and logging of all configuration data.

(P26) All measurement values and measurement results having been issued, stored, printed out and dispatched must at any time allow the corresponding measurement parameters and configuration to be identified.

(P27) The software must issue measurement values and measurement parameters in a form suitable for archiving. If the data are archived electronically, a format for the stored data must be determined and described in the documentation. If the measurement results are not archived, it must be possible to process retrospectively, using software or otherwise, the archived measurement values (i.e. raw data), measurement parameters and configuration data to recalculate the measurement results.

(P28) Subsequent alterations of the stored or issued measurement values, measurement parameters or measurement results must be clearly marked. The date of and the reason for the alteration must be noted, as well as the name of the person in charge. The original value must remain identifiable (readable).

(P29) If the software issues test reports or calibration certificates, these must comply with the prescriptions for test reports or calibration certificates.

4.2.2 Requirements for the management of software

The requirements defined in this section apply to the management of software programs which are used in laboratories and contribute essentially to obtaining test and calibration results. Mainly, requirements with regard to various procedures or processes during the life cycle of the software products are derived. Processes of the life cycle of software regarded under this aspect are: contract design, purchase, development, documentation, validation, archiving, installation, and maintenance. For the individual software processes, binding regulations and documented procedures must exist in the laboratory. These can be worked out by the laboratory itself or can be part of an existing quality management system. Such required prescriptions can also be focused on certain essential subjects such as, for example, the aspect of data and program security.

In the following, the requirements are subdivided into the fields of security, documentation, validation, purchase/development, and installation as well as archiving.

Program and data security

(M1) The laboratory must have regulations at hand which guarantee the protection of the clients' confidential information and of the ownership rights in the

- processing of confidential data by software;
- issuing of confidential data in electronic files;
- archiving of electronic files or printouts;
- sending of confidential data by means of software.

It is best to summarise the security measures in a security concept. Some possible measures are:

- regulation of the access to the software (passwords, file access control, access to the computer or laboratory, stand-alone operation, firewall, etc.);
- regulation of the access to archives;
- encoding of electronic files on hard disks and on other electronic storage media;
- virus protection measures;
- encoding of transmitted data;
- sealing of computer printouts.

(M2) The laboratory must have regulations at hand which ensure the integrity of the data and programs and guarantee a protection against intentional and unintentional alterations. Measures to ensure program and data integrity can also be included in the security concept. Further possible measures - besides the ones mentioned above - are:

- activation of the write-protection function for files;
- ensuring that only the latest versions of programs and documentation are used;
- measures offered by the program itself (requests, plausibility checks, etc.).

Documentation process

(M3) The laboratory must have regulations for the documentation of software at hand. These regulations and the documentation of the software programs are part of the management system of the laboratory and must be taken into account for the control of documents.

(M4) (to be applied only for laboratory-developed software or software that was specially developed by order of a particular customer:) The software documentation must be checked and approved before use. This also applies to each revision. The checking and approval must be carried out by authorised (competent) personnel.

(M5) (to be applied only for laboratory-developed software or software that was specially developed by order of a particular customer:) If the software is modified, the documentation must also be revised. Electronic elements of the documentation (e.g. on-line help functions) must also be updated.

(M6) The latest valid version of the documentation of the software programs must be available in all places where it is needed to comply with the requirements for work quality.

(M7) Outdated versions of the documentation must be destroyed or clearly be identified as being invalid. This also applies to electronic elements (e.g. on-line help functions).

(M8) There must be a list in which the latest valid version of the documentation, its date of release and all locations where it is kept must be entered. This shall ensure that always the latest version of the documentation is used and all older ones can be localized and exchanged.

Validation

(M9) Before being used, the software has to be validated. The validation must be carried out in accordance with the regulations of the laboratory and must be documented.

(M10) After software modifications or larger updates the validation must be repeated.

Purchase/development and installation

(M11) The laboratory must have regulations covering how software is specified, developed, selected, installed, purchased and/or used. Here, it can be laid down:

- that purchases and programming orders have to be carried out on the basis of a requirements specification,

- who is going to compile the requirements specification and what will be its minimum content
- which standards and guidelines will have to be generally complied with
- according to which criteria the selection amongst several products or tenders is made
- who is responsible for the installation and validation
- who will give the permission to use the software

(M12) Software purchases and programming orders must be carried out based on a specification which must be checked and approved.

(M13) The installation and operation of the software program have to be documented. This includes the following information:

- software identification (name, version/variant);
- manufacturer;
- type of documentation and, if applicable, locations where the documentation is kept;
- date of installation;
- installation procedures, installation result, installation problems;
- computers on which the installation was carried out;
- validation documents or reference to these;
- validation result and directions for use having arisen from the validation;
- problems and modifications (configuration alterations, patches, smaller updates, etc.) during the running operation;
- if applicable, modifications made to the previous version and their impacts.

(M14) If the contribution of the software program to uncertainty is not negligible, it must be taken into account in the measurement uncertainty evaluation of the test or calibration result.

Archiving of software versions

(M15) All program versions (electronic files and printouts) that are relevant for the reconstruction of tests or calibrations must be archived. The laboratory must have archiving regulations with the following minimum content:

- type of the data to be archived (what is archived?);
- type of identification, registration and storage of the records;
- periods of record-keeping;
- measures against the loss of the physical and logical readability of electronic files;
- measures for the protection of confidentiality;
- measures for the protection against intentional or unintentional alterations.

(M16) Old program versions which were used for testing must be archived.

(M17) Before software programs are modified, the impact of the alterations on the already archived data must be estimated and, if need be, measures to maintain the readability of the data (migration) must be taken.

4.3 Methods of software validation

Testing or calibration methods having been developed or modified by a laboratory itself must be adequately validated. From the software point of view, this means that the software products used to implement the methods have to be validated. Validation means, as interpreted by the standard, to confirm, by examination and by provision of objective evidence, that the particular requirements for a specific intended use are fulfilled. The applied validation methods shall be appropriate for the kind and scope of the respective examination. Commercial off-the-shelf software (e.g., word processing, spreadsheet, database, and statistical programs for general use) which is used within its intended range of application can be considered to be sufficiently validated. For software having been specially developed for a customer, the validation can be regarded as a part of the acceptance test. During validation, the testing documents of the manufacturer should be taken into account. Software which is already in use must be validated, too. For this purpose, first of all, all those documents which establish confidence in the correctness of the software must be collected. It is only after this that a decision can be taken with regard to appropriate validation procedures. Also the evidence of long years of satisfactory operation of the software can be a suitable method of validation. When software products or processes have been modified, it is necessary to carry out a new validation (partial validation). All records with regard to the validation procedures and validation results must be preserved.

The spectrum of applicable validation methods is very broad. It extends from the simple checking of documents (e.g., program descriptions, test documentations) to functionality tests of computer programs, through to systematic inspections of program source codes. For each application, an individual and suitable validation method has to be determined. In the following, examples of validation methods are listed.

Functional test of the software product

- Operation of the software with determined input values (measurement standards) whose corresponding output values are known
- Repeated tests of one or several retained specimens with known characteristic values
- Repeated tests using the same or different procedures, and comparison of the results
- Comparison with a reference software, with reference calculation processes, with (certified) reference materials (programs, data)

Examination/inspection of the documentation

- Detailed inspection of software documentation (e.g., user manual, development documentation, test documentation, requirements specification, design specification)

- Examinations regarding the compliance with all relevant standards and (legal) regulations
- Examinations aimed at the preservation of the traceability to SI units by the software
- Inspection of specific requirements (limits of use, sensitivity, uncertainty, linearity, etc.)

Further investigations

- Audits of quality management systems and/or realisations of software process models (e.g., for software development, configuration management, total life cycle)
- Evaluation/acceptance of corresponding audit reports or declarations (self-declarations) of manufacturers/developers
- Evidence of satisfactory operation of the software; documentation of long-term correct operation of the software
- Interlaboratory comparisons regarding software validation

In case of high criticality of individual computer programs or program components, software engineering methods with higher testing depth can be of advantage. Such methods are, for example, dynamic functional tests of individual program modules (black-box/white-box tests), tool-aided static analyses of the program source code, or manual inspections of the program source code (see Section 3.3). If these methods are used, it is recommended to involve software engineering experts.

5 Outlook: Towards a GUM-like validation guideline

Software validation is an interdisciplinary task; it is much more than just testing a program. Software validation begins with the definition and refinement of requirements and ends with the generation of objective validation reports. To a great deal, validation consists in the specification of requirements.

In this activity, metrologists and software engineers have to work closely together. At decisive points in the process, for example, when it comes to refining the requirements, their collaboration is an absolute must. This is particularly true with the example of the GUM-supporting software product, whose validation is discussed in a case study presented on the accompanying DVD, where also a list of International Standards on the issue is reported.

To be able to support metrology over the entire validation process, the existence of a comprehensive ‘Guide for Software Validation in Metrology’ is indispensable. A guide being of equal significance as the GUM should provide at least:

- A common understanding of software validation in metrology
- A uniform terminology for software validation

- A description of the metrology-specific methodology
- Recommendations for carrying out validations
- Recommendations for selecting appropriate methods
- A means to evaluate the validations having been carried out
- ‘Best-practice’ examples

Thus, such a ‘Guide for Software Validation in Metrology’ should become a supporting companion for metrologists.

Basic questions to be dealt with when starting to elaborate such a guide are as follows.

- Basic aim: Should the focus be preferably on the support of constructive, preventive quality assurance, or rather on guidance for the analytical software validation?
- Basis: Should we start completely a new or should we just take an existing approach and develop it further?
- Range of software: Which range of metrological software should be covered by the guide?
- Basic standards: Selection of the relevant standards to which the guide shall be traced back.
- Risk management: Role and assessment of different risk evaluations.
- Variety of requirements: Should a methodology be provided to derive and refine the requirements, or shall the requirements be developed for typical cases?
- Variety of methods: Should the full range of validation methods be offered or just a preselection?
- Methods: Should new methods be developed or old ones adapted, or should only the existing ones be used?
- Restrictions: Will there, in practice, be any restrictions as to time, costs, and so on which will have to be taken into consideration, and to which extent?

It is to be expected that a Software Validation Guide will become much more complex than the GUM, for example. It is high time to start a corresponding initiative for developing such a Software Validation Guide.

References

- [Be90] Beizer, B.: *Software Testing Techniques*, Van Nostrand, New York, 1990.
- [Be95] Beizer, B.: *Black-Box Testing*, John Wiley, New York, 1995.
- [Be00] Beizer, B.: Software is Different, in: *Annals of Software Engineering*, Volume 10, Numbers 1–4, pp. 293–310 (18), Springer, 2000.
- [CC04] *Common Evaluation Methodology for Information Technology Security (CEM)*, version 2.4, 2004.
- [CMM02] *Capability Maturity Model Integration (CMMI)*, <http://www.sei.cmu.edu/cmmi/cmmi.html>, 2002.

- [DS97] Development Standard for IT Systems of the Federal Republic of Germany (EStIT), Part 1: Software Lifecycle Process Model (V-Model), General Offset Reproduction No. 250, 1997.
- [FDA02] U.S. Department of Health and Human Services, General Principles of Software Validation; Final Guidance for Industry and FDA Staff, January 11, 2002.
- [GLP95] OECD Series on Principles of Good Laboratories Practice, No. 10: The Application of the Principles of GLP to Computerised Systems, OECD/GD(95)115, 1995.
- [GR00] Greif, N., Richter, D.: Software Engineering Related Standards and Guidelines for Metrology, in: P. Ciarlini, A. Forbes, F. Pavese, D. Richter, Advanced Mathematical & Computational Tools in Metrology IV, 109 - 121, World Scientific, London, ISBN 9810242166, 2000.
- [Gr04] Greif, N.: Richtig funktionierende Software: Anforderungen und ihre berprfung, in: Sensoren und Messsysteme, VDI-Berichte 1829, 85 - 95, ISBN 3-18-091829-2, VDI-Verlag, 2004.
- [Gr06a] Greif, N.: Software Testing and Preventive Quality Assurance for Metrology, in: Computer Standards & Interfaces 28 (2006), 286-296, ISSN 0920-5489, 2006.
- [Gr06b] Greif, N.: Validierung messtechnischer Software, in: Sensor und Test 2006, PTB-Bericht PTB-Q-3, 69-77, 2006.
- [GS99] Greif, N., Schrepf, H.: Towards Secure Measurements Through Software Analysis, in: D. Richter, V. Granovski (eds.), Methodological Aspects of Data Processing and Information Systems in Metrology, PTB Report PTB-IT-7, 52-59, June, 1999.
- [GS00] Greif, N., Schrepf, H.: Software requirements for measuring systems - Examples for requirements catalogues, PTB Laboratory Report, PTB-8.31-2000-2, Braunschweig and Berlin, July, 2000.
- [GS06] Greif, N., Saborrosch, D.: Guideline for Software Documentation, PTB-IT-14, Braunschweig und Berlin, July 2006.
- [GSS07] Greif, N., Schrepf, H., Saborrosch, D.: Several Guidelines for the Development and Validation of Software in Metrology, www.softwarepruefstelle.de
- [GSR99] Greif, N., Schrepf, H., Richter, D.: Software Evaluation in Calibration Services: Requirements and Testing Procedure, in: D. Richter, V. Granovski (eds.), Methodological Aspects of Data Processing and Information Systems in Metrology, PTB Report PTB-IT-7, 60-72, ISBN 3-89701-379-7, Braunschweig and Berlin, June, 1999.
- [GSR06] Greif, N., Schrepf, H., Richter, D.: Software Validation in Metrology: A Case Study for a GUM-Supporting Software, Measurement 39(2006) 849-855, 2006.
- [GUM95] Guide to the Expression of Uncertainty in Measurement (GUM), ISO/BIPM Guide, 1995.
- [ITS91] Information Technology Security Evaluation Criteria (ITSEC), Version 1.2, Commission of the European Communities, 1991.
- [ITS92] Information Technology Security Evaluation Manual (ITSEM), Commission of the European Communities, 1992.
- [Ja06] Jacobsen, J.: Validation of software in measuring instruments, in: Computer Standards & Interfaces 28 (2006), 277-285, 2006.

- [KHL01] Kitchenham, B.A., Hughes, R.T., Linkman, S.G.: Modelling Software Measurement Data; IEEE Transactions on Software Engineering, 27, 9 (2001).
- [Me79] Meyers, G.J.: The Art of Software Testing, John Wiley, New York, 1979.
- [MID04] Measuring Instruments Directive (MID), Directive 2004/22/EC of the European Parliament and of the Council on Measuring Instruments, Official Journal of the European Union, L135/1, 30 April 2004.
- [NOR03] NORDTEST Technical Report 535, Method of Software Validation, 2003.
- [Ri02] Richter, D.: Software in der Metrologie - Chancen und Risiken, in: PTB-Mitteilungen, Heft 2/2002, 99–113, 2002.
- [Ri06] Richter, D. (ed.): Special Issue “Validation of Software in Metrology”, Computer, Standards & Interfaces, Elsevier, 28, 3, January (2006).
- [SS97] Sommerville, I., Sawyer, P.: Requirements Engineering, Wiley, New York, 1997.
- [TG06] Tasic, T., Grottker, U.: An overview of guidance documents for software in metrological applications, in: Computer Standards & Interfaces 28, 256–269, (2006).
- [WEL05] WELMEC Guide 7.2: Software Guide (Measuring Instruments Directive 2004/22/EC), <http://www.welmec.org>, 2005.
- [WPB07] Wichmann, B., Parkin, G., Barker, R.: Validation of Software in Measurement Systems, Version 2.2, Software Support for Metrology, Best Practice Guide No. 1, NPL Report DEM-ES 014, January, 2007.
- [Zu97] Zuse, H.: A Framework for Software Measurement; de Gruyter, Berlin and New York, 1997.

Virtual Instrumentation

Octavian Postolache^{1,2}, Pedro S. Girão¹, José M. D. Pereira^{1,2}

¹ Instituto de Telecomunicações,

Instituto Superior Técnico, DEEC, UTL, Av. Rovisco Pais 1, 1049-001

Lisbon, Portugal octavian.postolache@ist.utl.pt, psgirao@ist.utl.pt

² Escola Superior de Tecnologia, IPS, Setúbal joseper@est.ips.pt

Summary. Traditional instruments and instrumentation have been replaced at an increasing pace by hardware/software mixed measurement-oriented systems. The software component gives the hardware extended measuring capabilities and the instruments are thus named virtual instruments. This chapter deals with virtual instruments and with the different forms of their implementation. Both the hardware and software components of a virtual instruments are detailed. A short reference to virtual laboratories is also included.

Key words: Virtual instrument, virtual instrument software architecture, Interchangeable Virtual Instrument (IVI), graphical programming language, VI, virtual laboratory

1 Introduction

According to the International Vocabulary of Basic and General Terms in Metrology (VIM), a measuring instrument is *a device or combination of devices designed for measurement of quantities*. This definition does not distinguish a real instrument from a virtual instrument. Therefore, we chose to present here an extensive definition of a *virtual instrument* (VI), that is to say, what we think a virtual instrument is and what it is not.

Everyone agrees that the *virtual* concept is related to the fact that a virtual instrument is able to provide more information than the one immediately available from its underlying hardware. This means that a virtual instrument is not only hardware but *must have* a software component. However, several traditional instruments, like some network analyzers or even bench-top multi-meters, have both hard- and software components and no one classifies them as virtual instruments. This is because the virtual concept before mentioned is not satisfied: those instruments just do what they were meant to do. Thus, a virtual instrument is more than just hardware and software. Other aspects are also characteristic of a virtual instrument or system:

- Hardware reconfigurability, that is, the same hardware must be able to be reprogrammed for different purposes. The user can customize, up to a certain point, the hardware.
- Software portability, that is, it should be possible for the software developed in one software environment for a particular application to be used for the same application in another software environment.
- Virtual instruments do not have displays or knobs. The ‘normal’ user–instrument interface is through the display unit (e.g., monitor) used to ‘create’ the instrument’s GUI.

From what has been said, it is clear that virtual instruments and virtual instrumentation systems can be developed using different types of hardware (e.g., sensors, data acquisition boards, general purpose instruments) and different programming environments (e.g., textual languages, graphical languages).

Virtual instruments (VIs) ([1]–[4]) entered the instrumentation and measurement world in the late 1980s when the technology of personal computers and object-and visual-oriented programming languages allowed the combination of a general-purpose computer and a generic data acquisition system to emulate several traditional measurement instruments. Virtual instruments are characterized by their versatility and low cost so they are very suitable for different areas, including education. The primary difference between real instrumentation and virtual instrumentation is the software component of a virtual instrument. The software enables complex and expensive equipment to be replaced by simpler and less expensive hardware. The whole system is characterized by its versatility because software and hardware modules change dynamically according to the application requirements.

Virtual instrumentation or instruments are truly instruments, such as a digital voltmeter, digital counter, or oscilloscope. However, the difference between a virtual instrument and conventional instrumentation is huge as can be concluded from Table 1.

Without the displays, knobs, and switches of conventional, external, box-based instrumentation products, virtual instruments use a device with processing and display capabilities such as a personal computer (PC) for all user interaction and control. In many common measurement applications, a data acquisition board or card, with a personal computer and software, can be used to create an instrument, which can reduce instrumentation costs by a factor of four.

With the increasing demand of VIs and the technological developments in the last decade in telecommunications and computer electronics, joined by a significant development in digital signal processing, several things happened.

- Development of user-friendly software for VI design using graphical programming languages to make programming easier. This software was designed in a manner that guides the user intuitively to choose system elements which are appropriate for the user’s objective, and which

Table 1. Comparison between traditional and virtual instruments.

Traditional Instruments	Virtual Instruments
Vendor-defined	User-defined
Function-specific, standalone with limited connectivity	Application-oriented system with connectivity to networks, peripherals, and applications
Hardware is the key	Software is the key
Expensive	Low-cost, reusable
Closed, fixed functionality	Open, flexible functionality leveraging off familiar computer technology
Slow turn on technology (5–10 year life cycle)	Fast turn on technology (1–2 year life cycle)
Minimal economics of scale	Maximum economics of scale
High development and maintenance costs	Software minimizes development and maintenance costs

automatically precludes the user's selection of system elements that are mutually incompatible

- Possibility of implementing VIs using processing units less powerful, smaller, and less expensive than a PC (e.g., ARM microcontrollers)
- Possibility of increasing the accuracy of measurement through digital signal processing
- Migration of the VI concept to other environments (e.g., to industry, that traditionally favours solutions tested over time but recognises the positive features associated to VIs such as high performance, flexibility, productivity and low cost)
- Integration of the VI technology and other technologies (e.g., sensing, communications, information technology)
- Viability of wired and wireless networked measuring systems for applications requiring measurements in different locations – distributed measurement systems
- Viability of dealing with applications requiring real-time measuring systems using low cost off-the-shelf equipment

It is in a context modeled by the above-mentioned events that the present chapter is written. Telecommunications and information technology are strongly involved in virtual instrumentation and consequently the text abundantly uses terminology typical for those domains.

It is the synergy between hardware and software that confers on VIs their appealing characteristics. By discussing hardware in Section 2 and software in Section 3 it looks as if we are trying to divide something that makes sense only when it is considered as a whole. Far from it! If we do that it is because we believe that the presentation becomes clearer dealing separately with each

component. We do hope that the reader agrees with us after reading the whole chapter.

A final comment in this introductory section, about VI drawbacks, three in particular, is in order.

- The software component of a VI is often its main asset and the VI performance critically depends on its software. As we all know, software is extremely prone to problems, both due to programming but also when it runs in a processing system. Particular attention must be paid to potential problems related to software.
- The flexibility of VIs often comes at a cost in the sense that an equivalent traditional instrument, whose design has been optimized for a single or a restricted number of tasks, may outperform a VI.
- As with any other measuring instrument or system, virtual instruments and systems must be calibrated. Software validation is essential to ensure the overall precision of the instrument or system.

2 Hardware

As already mentioned, virtual instruments combine rapid development software and modular flexible hardware to create user-defined test and measurement systems. The VI main concept is related to the utilization of a computing platform for test and measurement tasks including general-purpose hardware and specific software. By loading different software, different functionality will be loaded into the system.

The most widely used computing platform is the IBM-compatible personal computer (PC), but other platforms such as PXI (Peripheral Component Interconnect (PCI) bus eXtension for Instrumentation), VXI (VME Bus Extension for Instrumentation), Macintosh, VME (Versa Module Eurocard Bus)-based and UNIX (UNiplexed Information and Computing System) based systems are also popular. The VI makes use of the services and architecture provided by these computing platforms, namely of standard I/O hardware (RS232, USB, IEEE 1394). General-purpose hardware such as multifunction I/O boards (MIOB), GPIB controllers, and RS232-to-RS485 bridges permit the extension of the VI hardware functionality. At the same time, networking (wired or wireless) with other computers and devices is used to extend the functionalities to remote measuring and data publishing.

According to the number of hardware elements involved in the VI materialization,

single unit virtual instruments and *multiple unit virtual instruments* are considered. The first category includes multifunction input/output (I/O) boards with different computer bus compatibility (PCI, PCMCIA, USB, FireWire, RS232, CompactFlash). The second category of VIs is characterised by different hardware units including computational modules (e.g., real-time controllers), I/O modules (e.g., NI CompactRIO 9215) and networking modules

(e.g., RS485, Ethernet, Bluetooth). Real-time distributed systems with analog and/or digital I/O (e.g., VXI, PXI, FieldPoint, Compact RIO) and general-purpose ‘classical’ instruments such as multimeters and oscilloscopes with remote control interfaces, GPIB, USB, or Ethernet, can be also considered multiunit virtual instruments.

The next section discusses single unit virtual systems developed mainly around multifunction I/O boards, also known as data acquisition boards.

2.1 Single unit based virtual instruments – Multifunction I/O boards

Data acquisition systems, as the name implies, are systems used to collect information to document or analyse some phenomenon. As technology has progressed, data acquisition has been made more accurate, versatile, and reliable through electronic equipment. One of the most representative devices, the *multifunction I/O board (MIOB)* [5, 6], is designed to meet the needs of Virtual Instrumentation Systems (VISs) by providing multiple measurement modes. These different measurement modes include analog input, digital I/O, counter/timer functionality, external timing, triggering, and control all in a single package that is compatible with different types of computer buses including PCI. Examples of MIOB buses are PCI, ExpressPCI, and PCMCIA for plug-in computer multifunction I/O boards, and RS232, RS485, USB, Ethernet, and FireWire for external wired MIOBs.

The MIOB serves as a focal point in a VI system, tying together a wide variety of products, including sensors. The main blocks implemented in a MIOB are: sample-and-hold block, analog-to-digital converter block (ADC), digital-to-analog converter block (DAC), digital input/output (DIO) block, timing and synchronisation block (TS), bus controller block (BC), plug and play block (PnP), and interrupt interface block (II).

The ADC block has usually one of the designs: multiplexed–single ADC (M-SADC), or parallel–multiple ADCs (P-MADC). The M-SADC type is a low cost solution and permits the digitization of signals from multiple channels at a low sampling rate. The P-MADC type is more expensive and presents better dynamic characteristics and simultaneous sampling capabilities, because each channel has its own sample-and-hold, a dedicated programmable gain instrumentation amplifier and an ADC. Figure 1 represents the acquisition part of a MIOB for multiplexed and parallel architectures.

In Figure 1 one can identify the Programmable Gain Instrumentation Amplifiers (PGIA) that adapt the level of the input signals to the range of the ADC. To avoid losing samples, the ADC requires dedicated memory (e.g., 256-word first-in-first-out (FIFO) memory). The ADC timing control is provided by the TS block that ensures the correct scheduling of the complex interaction among the memory, the bus controller, and other external timing signals. The main electronic device of the TS block is a programmable universal timer (e.g., OKI MSM82C54-2RS/GS/JS).

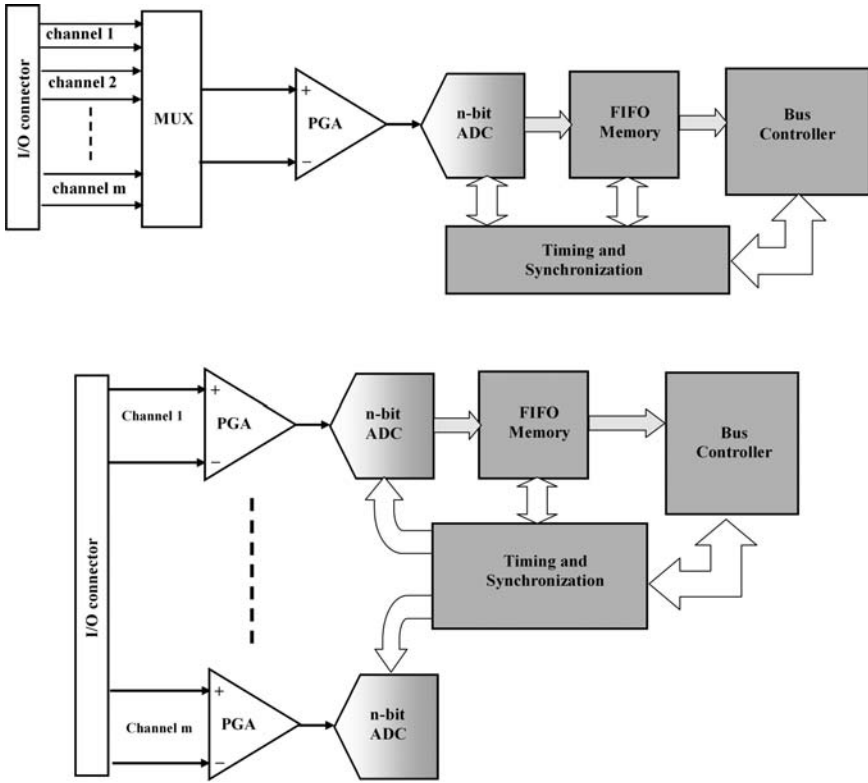


Fig. 1. ADC block for the M-SADC (top) and P-MADC (bottom) architectures.

The bus controller provides a high rate (usually improperly called high speed, but that we also use throughout this chapter) transfer communication path to other storage units such as the computer host memory.

For MIOB selection as part of a VI, the following characteristics are important [7]:

- Computer bus compatibility
- Number of analog inputs
- Input resolution (number of bits)
- Maximum sampling rate (samples/s)
- Multichannel sampling rate (samples/s)
- Number of analog outputs
- Output resolution (number of bits)
- Maximum output rate (samples/s)
- Output range (V)
- Number of digital I/O lines
- Number of counter/timers
- Triggering type – analog, digital

The plug-in computer MIOB presents the advantages of transfer speed (because they are connected directly to the bus) and cost (because the overheads of packaging and power are provided by the computer). Each board installed in a computer has a unique input/output map location. The I/O map in the computer provides the address locations that the processor uses to gain access to the specific device as required by a program.

When field measurements are required, external MIOBs with RS232, USB, PCMCIA, CompactFlash, or FireWire bus interfaces are used to develop appropriate virtual instruments. A portable MIOB offers a number of advantages over PCI, ExpressPCI, and PXI MIOBs. external MIOBs are easy to use, because one does not have to open a computer chassis to install it, are easier to move from one application to another, and have the capacity to enable a user to begin taking measurements very quickly (when the board is combined with ready-to-run DAQ software) because of the plug-and-play capability of PCMCIA, CompactFlash, USB, and FireWire. Considering the latest developments in the area of portable MIOBs, elements of CompactFlash, USB, and FireWire buses and particular MIOB features as part of portable virtual instruments are discussed in more detail next.

USB protocol and MIOBs with USB bus controller

The *Universal Serial Bus (USB)* [8] was introduced in 1995 to address a number of connectivity issues associated with existing serial communication standards. USB supports multiple devices and provides easier installation, faster transmission speeds, and simpler cabling requirements than conventional parallel or serial ports. USB was also designed to supply modest operating power directly to peripherals, eliminating the need for external power supplies in some cases.

Many PCI-Pentium-based PCs include up to four USB ports. Using USB hubs, up to 127 USB peripherals can be connected. USB connections work well at distances up to 50 feet (15m). The maximum speed for USB version 1.1 is 12 megabits per second (Mbits/second or Mbps). However, the USB 2.0 standard released in 2000 raised the maximum USB speed to 480 Mbits/second.

USB peripherals are hot plug devices, so they can be attached or removed from an energized PC without damage and without the need to reboot the system. USB also incorporates plug-and-play, so a compatible PnP operating system will automatically recognize and reconfigure PC resources to handle the addition or removal of a USB peripheral.

USB modules offer good noise immunity, with performance benefits for noise-sensitive measurements. USB cables are typically 1 to 5 metres long, so the I/O circuitry is located further away from the computer's noisy motherboard and power supplies, and closer to the signals they will be measuring. USB is fully supported by Windows 98, Windows Me, Windows 2000, Windows XP, and Windows Vista.

Concerning the USB MIOB as a component of virtual instruments, different solutions are commercially available from manufacturers such as National Instruments, Keithley Instruments, and Agilent. USB MIOBs that comply with USB 2.0 and are connected to PCs with a high-speed USB 2.0 port can attain the full data transfer rate of USB 2.0. The increased bandwidth enables multiple I/O operations simultaneously at throughput rates up to 1.2 MS/s (e.g., NI 6251). MIOBs are also compact and portable (Figure 2), allowing their use in field applications.

FireWire (IEEE 1394) protocol and MIOB with FireWire controller

IEEE 1394 [9, 10] is a serial bus standard that was developed collaboratively by Apple, Intel, Texas Instruments, Microsoft, Sun Microsystems, Compaq, and National Semiconductor. FireWire is the Apple Computer trademark for its implementation of IEEE 1394.

FireWire, built to achieve high-speed data transfer, uses a peer-to-peer architecture in which the peripherals are intelligent and can negotiate bus conflicts to determine which device can best control a data transfer. Data transfer rates of 100, 200, and 400 Mbps are currently supported by the bus architecture.

FireWire was also developed for the purpose of creating an inexpensive alternative to parallel buses. Some problems with parallel buses are: (a) they are confined to a small physical area; (b) they do not have plug-and-play




Fig. 2. USB-MIOB implementations: Keithley KUSB-3160 (left), National Instruments USB-6251 (right).

support; (c) and they do not support isochronous applications. The IEEE 1394 bus was engineered to solve these problems.

FireWire MIOBs are developed now by National Instruments (e.g., DAQPad-6070E, DAQPad-6052E) and are available for Windows 2000 and Windows XP (but not yet for Mac OS). Several advantages of IEEE 1394 are: a tiered-star topology allowing up to 63 devices (only 16 devices are allowed by the National Instruments NI-DAQ drivers), long distance (up to 500 m with extender), hot plug-and-play, easy configuration, no need to open the PC for installation, high throughput (up to six DAQPad-6070Es at full rate), higher than PCMCIA. As IEEE 1394 disadvantages can be mentioned: FireWire MIOBs are more expensive than MIOB supported by PCI and PCMCIA; only one manufacturer commercializes boards.

Concerning MIOBs' performance, the maximum aggregate sampling rate is about 1.25 MS/s for 12 bit resolution. Although the USB2.0 protocol was already upgraded, FireWire will be upgraded (1394b) to support communication rates up to 3200 Mbits/s.

CompactFlash protocol and MIOB with CompactFlash controller

CompactFlash (CF) [11, 12] was originally a type of data storage device used in portable electronic devices. As a storage device, it typically uses flash memory in a standardised enclosure, and was first specified and produced by SanDisk Corporation in 1994. The physical format is now used for a variety of devices including MIOB devices (e.g., Talisman DATAQ-CF2, National Instruments NI CF-6004). The Compact Flash Standard organization (www.compactflash.org) promotes two CF standards. According to the CompactFlash Specification Revision 2.0 (2003), the data transfer rate associated with CF devices is 16 Mbps, whereas for the devices compatible with CompactFlash Specification Revision 3.0 the data transfer rate is up to 66 Mbps.

MIOBs were designed and implemented according to the CF standard specifications and can be used as the hardware component of virtual instruments mobile systems (Figure 3).

Several characteristics of commercially available MIOBs are presented in Table 2.

Table 2. CompactFlash MIOBs characteristics.

MIOB	BUS	Analog Input	Res- olution (bits)	Analog Out- put	Max Sam- pling Rate (kS/s)	Max Input Range (V)	DIO
NI CF-6004	CF	4SE	14		200	± 5	4
TS-F CF241	CF	4SE	24	2	38.4	± 2.5	4
DATAQ CF2	CF	4SE	24	2	38.4	± 2.5	4



Fig. 3. CompactFlash MIOB solutions.

2.2 Multiunit based virtual instruments

Multiunit virtual instruments are those that use input/output modules with networking interfacing capabilities. One of the first types of virtual systems developed for test and measurement application were constructed using general-purpose traditional instruments connected in a GPIB or a RS485 network. In order to provide connectivity between the processing unit (e.g., PC) and the measurement modules, the computer must be compatible with the implemented test and measurement network (e.g., GPIB traditional instrument network, Ethernet traditional instrument network). Because a GPIB controller is not available on general-purpose computers (only USB, RS232, Centronix, and Ethernet communication ports are usually included), additional interfacing modules (e.g., NI PCIe-GPIB, NI GPIB-USB-HS) have to be included as part of the hardware. For Ethernet compatible instruments, the remote control is performed using the computer Ethernet interface according to the measuring node IP (Internet protocol) address and port number (e.g. IP=192.168.1.2; port=5555)

Distributed virtual instrument systems based on general-purpose instruments

Traditional general-purpose instruments such as multimeters, oscilloscopes, programmable power supplies, and signal generators can be used as hardware components of virtual instrumentation systems as long as they are provided

with some sort of interface that allows remote operation. GPIB (IEEE 488) is perhaps the most successful of those interfaces. This standard interface has the following characteristics: up to 15 instruments, called devices, can be connected to a computer, usually called the controller, because it is responsible for coordinating bus traffic; 1 Mbps transfer rate; 8 bit parallel data bus; 3 wired or handshake lines; additional control lines (SRQ – Service ReQuest, IFC – InterFace Clear, REN – Remote ENable, EOI – End Or Identify, ATN – ATtention).

An example of a GPIB-based distributed virtual system to test a device, DUT, is shown in Figure 4.

Nowadays, in general-purpose measuring instruments, the GPIB interface has been replaced by Ethernet (local area network (LAN) interface). The primary benefit in using Ethernet is cost. In nearly all cases, the Ethernet network precedes the measurement system, so it often adds little cost to the measurement system itself. Ethernet provides a low-cost, moderate-throughput method for exchanging data and control commands over long distances. However, due to its packet-based architecture, Ethernet is not deterministic and has relatively high latency. The lack of determinism and high latency are incompatible with the requirements of some instrumentation systems (e.g., real-time systems). These situations are better served with a dedicated bus such as PXI, VXI, GPIB, and more recently LXI (LAN eXtension for Instrumentation).

LAN extension for instrumentation-based systems

LAN eXtensions for Instrumentation (LXI) [13, 14] is a test and measurement system architecture based on proven, widely used standards such as Ethernet. LXI enables fast, efficient, and cost-effective creation and reconfiguration of test systems. The main reasons to consider LXI for present and

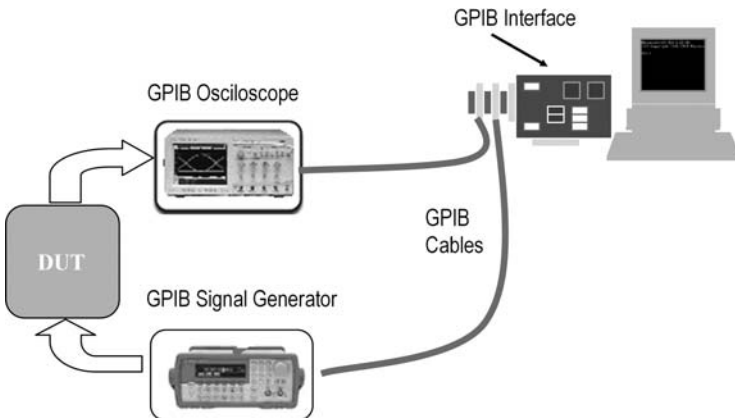


Fig. 4. GPIB Distributed virtual instrument (GPIB interface): DUT device under test.




Fig. 5. Distributed virtual instrument based on LXI network architecture with WWW connection.

future test systems are: ease of use, performance, cost, scalability, longevity, flexibility, and IEEE-1588 [15] synchronization. It is possible to use LXI to develop distributed virtual instruments. Figure 5 represents a distributed virtual instrument hardware architecture that includes a router to separate the instruments from the network.

Because LXI instruments use Ethernet and TCP/IP protocols to communicate, they can also connect to the Internet. A local server assigns each instrument an IP address consistent with its subnet and any user who knows the instrument's address can communicate with it from almost anywhere. This type of distributed virtual instrument poses performance and security challenges. Regarding performance, commands sent to the instrument, for instance, according to the standard commands for programming instruments, SCPI, must travel through a number of intervening hubs, routers, and switches that can decrease throughput because they also handle other Internet traffic. Concerning security, because the instruments operate usually under Windows operation systems, antivirus and firewall protection programs are required.

Distributed systems with real-time capabilities

A real-time system (RTS) responds in a (timely) predictable way to unpredictable external stimuli. In an RTS, timing is crucial and tasks have to start and finish within time boundaries. To meet this requirement, RTS must operate under real-time operating systems (RTOS). A distributed system can be defined as a system that uses two or more linked computing subsystems to achieve a task.

Real-time distributed applications such as simulation, control, or high-performance data logging must be built in a manner that guarantees that all elements of the system operate in a deterministic fashion.

Real-time distributed virtual instrument systems include different embedded architectures. Examples from National Instruments are PXI (PCI eXtension for Instrumentation), FieldPoint, and CompactRIO [16].

FieldPoint. *FieldPoint* is a proprietary solution for interfacing devices to computers developed by National Instruments that follows the fieldbus concept [17]. The idea of the fieldbus grew out of the problem of interfacing hundreds or thousands of sensors and actuators to programmable logic controllers (PLCs) and process control computers in large industrial plants. Rather than connect each sensor or actuator to a central plant computer, requiring hundreds or thousands of kilometers of wiring, the idea of *FieldPoint* was to connect related groups of sensors and actuators to a local microcomputer that communicates with the central plant computer via an Ethernet local area network. Earlier *FieldPoint* units used serial interface lines (RS232 and RS485 protocols) for communication.

PXI. *PCI Extension for Instrumentation systems (PXI)* [18, 19] consist of a rugged chassis, embedded controller, and plug-in I/O modules. To develop a real-time virtual instrument based on PXI, a real-time software component must be developed using a real-time programming language (e.g., LabVIEW Real-Time). For the particular case of LabVIEW Real-Time for ETS targets, the embedded controller is converted into a real-time controller by downloading the RTOS and application software to a dedicated microprocessor. The embedded software then has access to all I/O in the PXI system, taking advantage of the PXI advanced timing and synchronization features to achieve precise I/O triggering and multimodule synchronization.

CompactRIO. The *NI CompactRIO programmable automation controller (PAC)* [20, 21] is a low-cost reconfigurable control and acquisition system. The CompactRIO architecture is composed of an embedded real-time controller (e.g., cRIO-9002), a reconfigurable embedded chassis containing an FPGA (e.g., NI cRIO-9104), and hot-swappable I/O modules (e.g., NI 9233). The direct connection between the I/O modules and the FPGA enables the developer to tightly integrate timing and triggering between I/O modules through the FPGA. The embedded FPGA in CompactRIO is programmed using LabVIEW FPGA Module functions. The programmed areas are then accessed by specific functions of the LabVIEW Real-Time Module that runs in the CompactRIO Real-Time Controller.

3 Software

Various software technologies and development environments exist for developing virtual instruments. They can be classified in two categories: *textual programming languages* and *graphical programming languages*. Graphical programming languages are visual programming languages and textual programming languages can also be of the visual type (e.g., VisualC++) but that also generate program lines.

3.1 Textual programming languages

Two subcategories can be identified in textual programming languages: *procedural programming languages* and object-oriented languages.

Procedural programming for virtual instrumentation implementation

Procedural programming is by far the most common form of programming but not in the VI domain. A program is a series of instructions which operate on variables. It is also known as imperative programming. Examples of languages used to develop the software component of the virtual instruments, which despite their differences are considered procedural programming languages, include Basic, C, LabWindows/CVI, and MATLAB.

Procedural programming handles a problem by carrying out a sequence of operations. Structures such as IF–THEN–allow the control of the sequence of operations. Variables play a central role in programs, and their scope is an important notion in block-structured languages.

Advantages of procedural programming include its relative simplicity, and ease of implementation of compilers and interpreters. Disadvantages include the difficulties of reasoning about programs and, to some degree, difficulty of parallelization. Procedural programming tends to be relatively low level compared to some other paradigms, and as a result can be very much less productive. Because software development is so important to VIs, we next present a summary of the characteristics of several relevant procedural programming languages.

C and LabWindows/CVI. C is a powerful flexible language that provides fast program execution and imposes few constraints on the programmer. C's power and fast program execution come from its ability to access low-level commands, similar to assembly language, but with high-level syntax. Its flexibility comes from the many ways the programmer has to accomplish the same tasks. C includes bitwise operators along with powerful pointer manipulation capabilities.

Modularity is another good feature of C. Sections of code can be stored in libraries for reuse in future programs. This concept of modularity also helps with C's portability and execution speed. The core C language leaves out many features included in the core of other languages. These functions are instead stored in the C Standard Library where they can be called on when needed. An example of this concept is C's lack of in-built I/O capabilities. I/O functions tend to slow down program execution and also to be machine independent when programmed to run optimally. For these reasons, they are stored in a library separately from the C language and only included when necessary.

```

#include <stdio.h>
#include <bios.h>
#define com1 0
#define settings (0xE3)
main( )
{ /* declare an0 as a floating point number */
    float an0 ;
    /* configure com1 9600 baud, 8 bit words, no parity */
    bioscom (0,settings,com1);
    /* send RA0 command to ADR101 on com1 */
    fprintf (stderr,"RA0 \xD");
    /* initialize com1 buffer */
    fscanf (stderr,"%e",&an0);
    /* print data on screen */
    rewind (stderr);
    /* read data from com1 and store it at address of an0 */
    printf ("ANALOG PORT 0= %e PERCENT \n",an0); }

```

Fig. 6. Virtual voltmeter C routine.

For VI implementation, C languages offer the possibility to develop test and measurement instruments for most of the hardware platforms mentioned in Section 2 but are normally used for programming multifunction input/output boards with PCI, RS232, or USB interfaces. Figure 6 is a commented list of a routine for a virtual voltmeter based on a ADR101 RS232 compatible MIOB from Ontrak Control Systems.

C is, however, not a good solution to implement a graphical user interface (GUI) in a virtual instrument. A better choice is *LabWindows/CVI* [22] developed by National Instruments. It is a software development environment for programming in C. It is equipped with powerful function libraries and a comprehensive set of software tools for data acquisition, analysis, and presentation that are used to interactively develop data acquisition and instrument control applications. With *LabWindows/CVI* it is possible to edit, compile, link, and debug C programs. Additionally, compiled C object modules, dynamic link libraries (DLLs), C libraries, and instrument drivers can be used in conjunction with C source files to develop the software part of virtual instruments.

Typical *LabWindows/CVI* applications include the following elements: user interface, data acquisition, data analysis, and program control.

As mentioned above, an important advantage of *LabWindows/CVI* is its graphical user interface (GUI) capability. As shown in Figure 7, *CVI* offers predefined graphical elements that permit us to develop in a short time a user interface. The GUI is easy to implement but programming requires from the VI developers a good knowledge of C.

MATLAB. *MATLAB* from The MathWorks, Inc. [23,24] is a general-purpose procedural programming language whose use for VI programming increased with the inclusion in version 6.5 onwards of a Data Acquisition toolbox and an Instrument Control toolbox. The type and number of hardware supported by

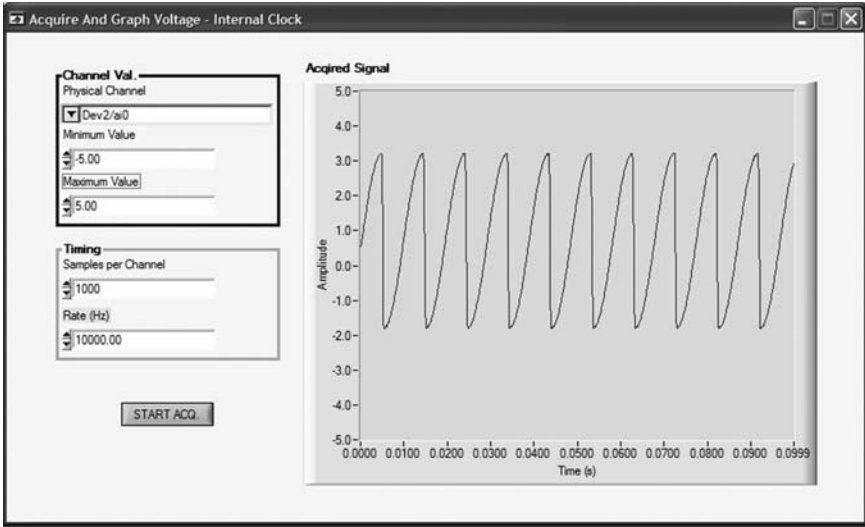


Fig. 7. GUI using the LabWindows/CVI.

this toolbox is continuously increasing and includes MIOBs, general-purpose remote-controlled instruments and digital signal processor (DSP) development cards (e.g., TI C64x). Developed in MATLAB, the software component is able to process signals acquired with the hardware in operation, the processing parameters being controlled from the graphical user interface. The user interface of a virtual oscilloscope and a virtual modulation analyzer developed in MATLAB are presented in Figure 8.

The Instrument Control toolbox includes functions that enable communication with instruments such as oscilloscopes, function generators, and analytical instruments directly from MATLAB. Thus, it is possible to generate data in MATLAB to send out to an instrument, or read data into MATLAB for analysis and visualization. This toolbox provides a consistent interface to all devices independent of hardware manufacturer, protocol, or driver. The toolbox supports VXIplug&play, GPIB, TCP/IP, and UDP communication protocols.

To combine different types of hardware in a VI, MATLAB (usually thought of as a procedural language) uses the device object concept to represent instruments, that is, the properties and methods specific of an instrument are encapsulated within device objects. The block diagram that expresses the device object relation with the GUI, the instrument driver, the interface object and the hardware component of the virtual instrument is presented in Figure 9. The device object represents a bridge between the MATLAB Instrument Driver and the MATLAB GUI whose elements were exemplified in Figure 8. Other important features of the MATLAB Instrument Control toolbox are that it permits: (i) synchronous and asynchronous (blocking and nonblocking)

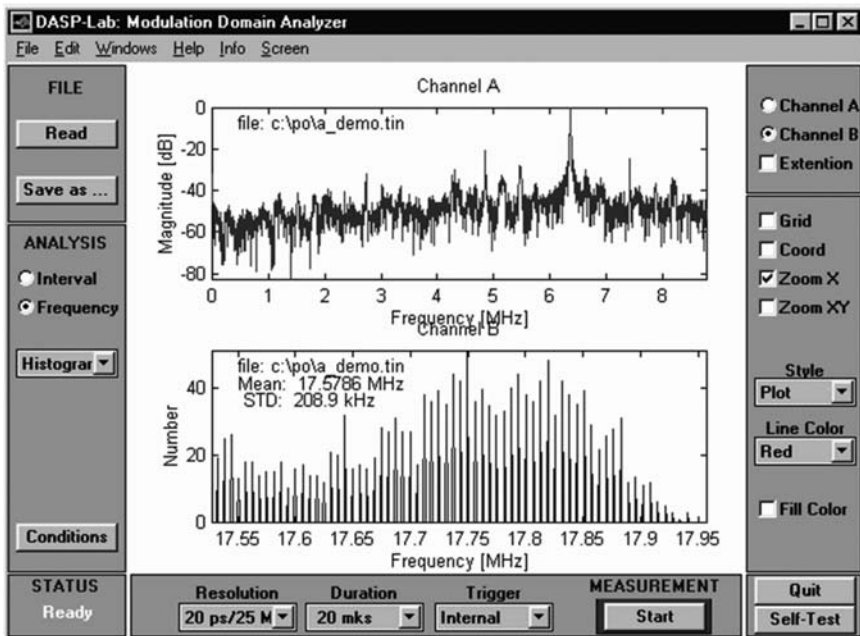



Fig. 8. Virtual instrument GUI based in MATLAB: virtual oscilloscope (top), virtual modulation analyzer (bottom).

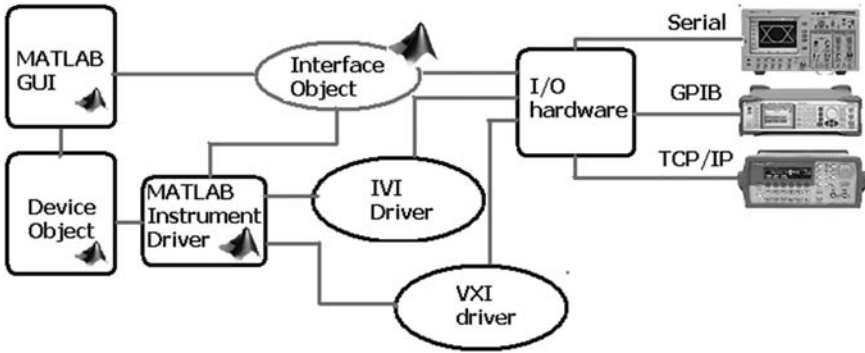


Fig. 9. MATLAB device object block diagram.

read-and-write operations; (ii) event handling for time-out, bytes read, data written, and other events.

MATLAB Data Acquisition toolbox provides a complete set of tools for MIOB based VIs. The toolbox allows the configuration of external hardware devices, reading data into MATLAB and Simulink for immediate analysis, and sending out data. GUI design tools can be used to design the virtual instrument front panel according to the implemented functionalities. As the Instrument Control toolbox, the Data Acquisition toolbox provides functions for creating device objects that are directly associated with MIOBs. These objects include base properties that apply to all supported hardware, such as sample rate, trigger settings, and channel properties. Device-specific properties permit us to access the specific features and capabilities of the used hardware. Data Acquisition Toolbox analog input functions can control the digitization of the analog input signals. An analog input object with the number of the input channels can be created and data transferred to the memory or read to MATLAB workspace. The analog output functions can control the hardware through an analog output object. Based on Data Acquisition toolbox functions and MATLAB GUI functions, a VI with an appropriate user interface can be implemented. In conclusion: MATLAB is a procedural programming language that provides an easy way for VI software implementation. However, the flexibility, portability, and MATLAB GUI elements are limited when compared to LabWindows/CVI or to graphical programming languages such as LabVIEW, VEE Pro, or TestPoint.

Object-oriented programming for virtual instruments implementation

Object-oriented programming (OOP) [25] is a programming language model organized around *objects* rather than *actions* and *data* rather than *logic*. Object-oriented programming takes the view that what we really care about

are the objects we want to manipulate rather than the logic required to manipulate them.

The first step in OOP is to identify all the objects and to generalize them as *classes of objects* by defining the kind of data they contain and the logic sequences that can manipulate them. Each distinct logic sequence is known as a *method*. A real instance of a class is an *object*. An object or a class instance runs in the computer. Its methods provide computer instructions and the class object characteristics provide relevant data. During operation a communication is made by objects – and the objects communicate with each other – with well-defined interfaces called messages.

Comparing object-oriented with a procedural language, one notes that the focus of procedural programming is to break down a programming task into a collection of data structures and subroutines, whereas in object-oriented programming it is to break down a programming task into objects.

Some differences between pure object-oriented languages and nonobject-oriented procedural languages are presented in Table 3.

The object-oriented languages used for VI software implementation include VisualC++, Visual C#, Java, and Visual Basic.Net. Visual Basic.Net (VB.NET) is part of a brand new platform, based on the .NET Framework with the result that VB.NET is fully object-oriented. This is a huge difference and is a powerful improvement over previous versions of Visual Basic.

Java [26, 29] as an object-oriented language is used only in particular applications involving RS232 compatible hardware or the new LXI devices. In fact, and to extend the capabilities of the websites to interact with different instruments, a Java extension for communication tasks, named Java Communications API, was recently developed. For example, the case of `javax.com` extension package includes new data classes that manage the access and ownership of communication ports, provide an interface to physical communications ports, and provide an interface between the low-level classes and the underlying operating system. Java communications API enables access to serial (RS-232), parallel (IEEE 1284), and USB ports and LXI devices. The Java communication API features provides, in the RS232 case, access to: port settings (baud rate, parity, stop bits), port naming, mapping, enumeration (configurable), data transfer, hardware/software flow-control, and receive-buffer threshold control [30]. In the LXI case, Agilent and Keithley instruments developed proprietary Java communication APIs that provide enhanced Web capabilities for the instruments, which allow users to operate the instruments

Table 3. Object-oriented languages and procedural languages element comparison.

Pure Object-Oriented Language	Pure Procedural Language
Methods	Functions
Objects	Modules
Message	Call
Member	Variable

via the Web similarly to how they use the front-panel interfaces. Thus, with a Web server built into every LXI instrument and a standard Java-enabled Web browser on the computer, the full control of an instrument over the instrument's built-in LAN interface can be carried out using the Web virtual interface (Figure 10).

Web technologies such as Java applets [31] can provide a rich, responsive, and easy-to-use Web interface with only relatively modest memory and processing power requirements. Potential features of those technologies include: (i) full instrument configuration, (ii) full instrument control, (iii) tabular or graphical display of data, and (iv) extensive user help. Java applets' performance can be a problem for some virtual instruments system. However, a variety of techniques available for optimizing the performance of embedded applications [32].

Nowadays, the hardware for VI implementation comes with software drivers for Visual C++, Visual C#, Visual Basic, and Visual Basic.NET. Keithley and National Instruments offer different software packages compatible with those languages. One of them is Measurement Studio by National Instruments. It is an integrated suite of classes and controls for test, measurement, and automation applications in Microsoft Visual Studio 2005, Visual Studio .NET 2003, and Visual Studio 6.0. NI Measurement Studio dramatically reduces application development time by providing Windows Forms,

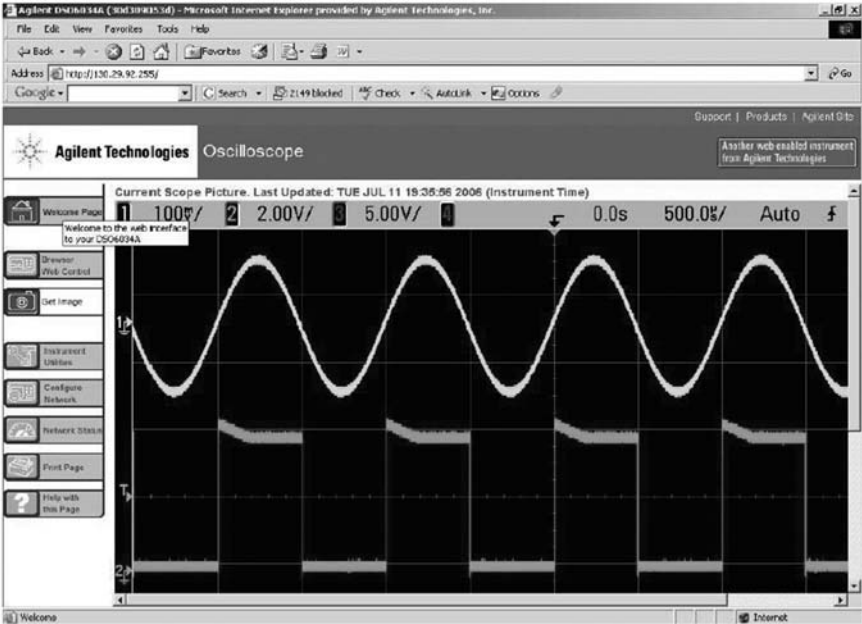


Fig. 10. Web page of an LXI virtual oscilloscope.

Web Forms, and ActiveX user interface components. Measurement Studio for Visual C# and Visual C.NET components provides the management .NET controls for creating rich Web and Windows GUIs, multithreaded application programming interface (API) for data acquisition, instrument control APIs and analysis libraries.

3.2 Graphical programming languages

A graphical programming language (GPL) is a visual programming language and thus, programs are built by manipulating program elements graphically but without producing textual program lines.

In the context of VIs, graphical programming languages tend to integrate the approach of *dataflow languages* [33] to have immediate access to the program state for in-line debugging purposes.

LabVIEW and VEE Pro are the most common dataflow and graphical programming languages for VI implementation.

LabVIEW

LabVIEW from National Instruments [34,35] is a graphical programming language that looks like flowcharting. It allows complex tasks to be broken down visually into block diagrams. LabVIEW is the platform currently most used for the development of software for PCs, PDAs, real-time devices (e.g., PXI, CompactRIO), and embedded systems-based VIs. It utilizes graphics that look like real instruments for the program inputs and outputs and subVI (nodes) that are wired together to perform the program functions. LabVIEW is almost totally portable across Windows, Linux, and Macintosh platforms.

A LabVIEW program is called a VI which is short for virtual instrument. It is also possible to create sub-VIs, which are like subroutines. This makes modular and hierarchical programming very easy.

LabVIEW can simulate controls such as sliders and dials, many types of switches, indicators, and waveform graphs on its virtual workbench. LabVIEW is designed to take input data directly from the user through its virtual instrument interface or from input/output channels and input it into a LabVIEW virtual instrument. Data are analyzed/processed and displayed on the VI GUI or saved in a file using LabVIEW functions. The processed data can be used to drive controls. An example of the front panel of a virtual spectrum analyzer developed in LabVIEW is presented in Figure 11.

LabVIEW Front Panel is where the user interacts with the program by entering data and viewing results. It is called Front Panel because it is supposed to look like the front panel of an instrument or a group of instruments. The Front Panel includes user inputs called *controls* and program outputs called *indicators*.

LabVIEW's most important component for program development is the *Block Diagram*, a window where the *graphical code* is created. The Block

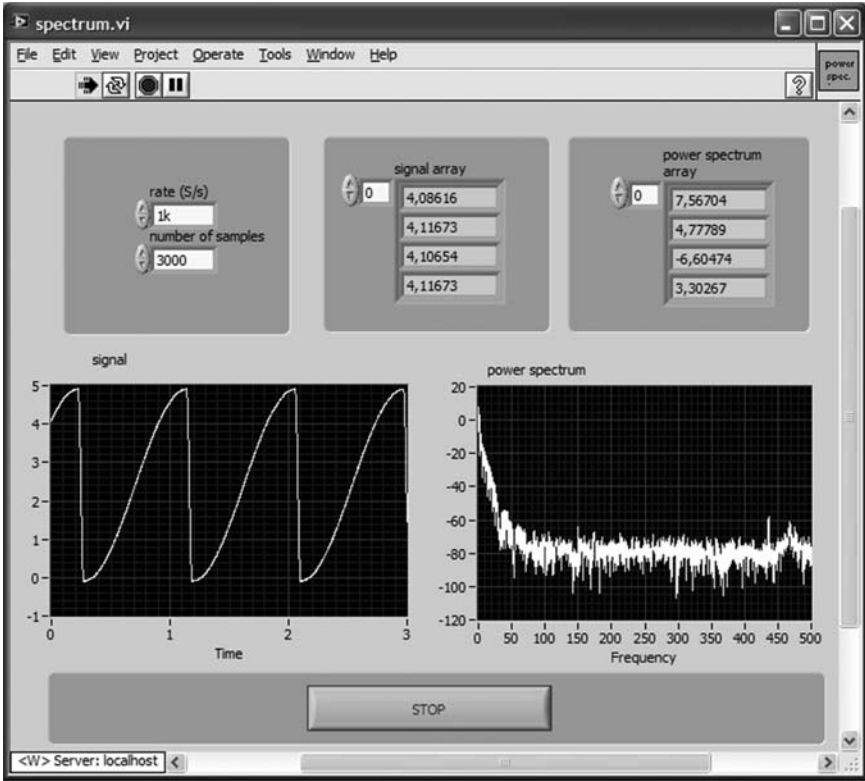


Fig. 11. Virtual spectrum analyzer front panel.

Diagram window (Figure 12) includes icons called *nodes*. There are many types of nodes: nodes that represent the controls and indicators of the Front Panel (e.g., stop, signal, and power spectrum in Figure 12) called *terminals*; nodes that represent *functions* (e.g., DAQ assistant, spectral measurements, Figure 12); and nodes that represent program execution structures (such as while loops, for loops, and case and sequence structures).

Most nodes are *wired* together for logical functioning of the program. All nodes have one or more input and/or output terminals called *connectors* to wire to and from. LabVIEW utilizes *dataflow* which means that a node executes its function only after it has received data from a previous node. Using one of LabVIEW features called *highlighted execution* the dataflow can be observed (Figure 13).

After the node is through with its function, the processed data are transferred via the wire to the next node. In the example of Figure 13, data go from the *Spectral Measurements Node* to the *Power Spectrum Node*. There are several different types of data and wire types to carry the data (e.g., integer, double, cluster, string) that are graphically expressed in different ways.

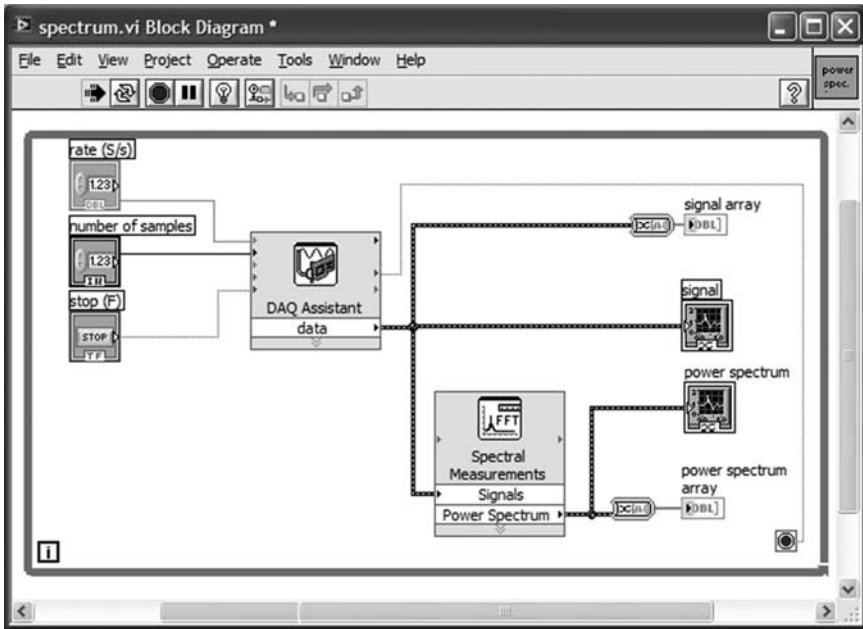


Fig. 12. Virtual spectrum analyzer Block Diagram.

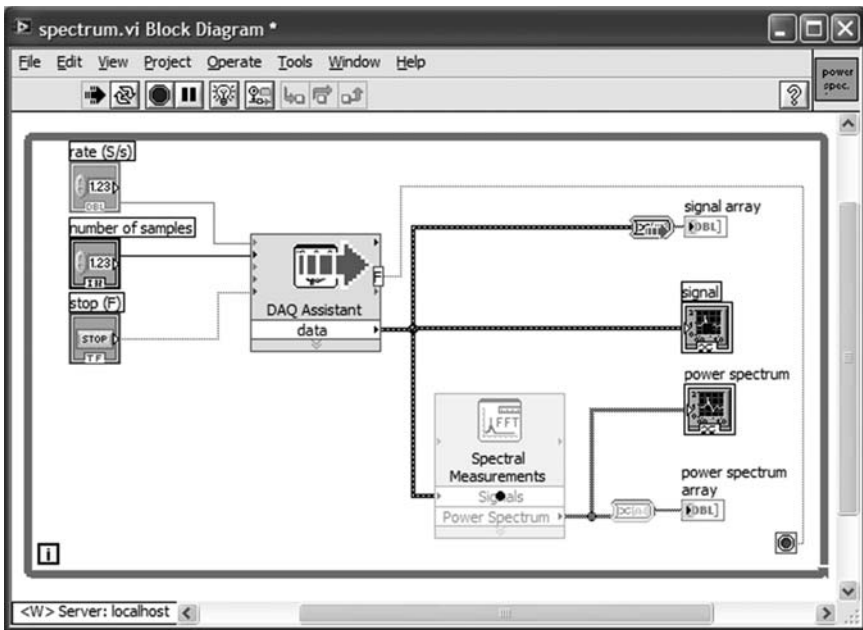


Fig. 13. The dataflow associated with LabVIEW program execution for the particular case of the virtual spectrum analyzer of Figures 11 and 12.

Table 4. LabVIEW programming terms compared to their conventional equivalents.

LabVIEW	Conventional Programming
VI	Program
LabVIEW	C, C++, Java, Pascal, BASIC, etc.
Function	Function or Method
Front Panel	User Interface
Block Diagram	Program Code
Sub VI	Subroutine, Subprogram

Table 4 includes a comparison among the main terms used in LabVIEW and in textual programming languages.

One of the main strengths of LabVIEW is its data analysis and data communication capabilities. Libraries for measurement, signal processing, mathematics, image processing, and control allow the extraction of information from acquired data within the LabVIEW environment without having to recur to external computational resources. Data communication is simplified by functions that allow the access of the program to communication channels (e.g., serial, TCP, Bluetooth).

VEE Pro

VEE (visual engineering environment) [36, 37] is a visual and dataflow programming language and development environment from Agilent Technologies. It is optimized to use with acquisition devices such as digital voltmeters and oscilloscopes, and source devices such as arbitrary waveform generators and power supplies.

VEE programs associated with virtual instruments are written by wiring together objects that represent *data sources*, *transforms*, and *sinks*. In the example shown in Figure 14, two function generators are the data sources and the waveform dataflows from their output terminals when the program runs. The ‘A+B’ object is a *data transform object*. It transforms the incoming data by adding the two waveforms, and then outputs the resulting waveform on its output terminal. From there, the dataflows to the waveform display which is a data sink; that is, it does not propagate the data any further and only displays it.

The VEE language is characterized by wide instrumentation connectivity and intuitive programming using a *block diagram*. As with the LabVIEW graphical programming environment, it permits us to select and edit objects from pull-down menus and connect them to each other by wires to specify the program’s flow, mimicking the order of tasks one wants to perform.

Agilent VEE Pro was designed to implement quickly virtual instruments associated with automatic measuring and testing. Hardware components can be controlled through GPIB, LAN, USB, RS-232, VXI, and other interfaces or buses, including PXI and SCXI data acquisition and modular instrumenta-

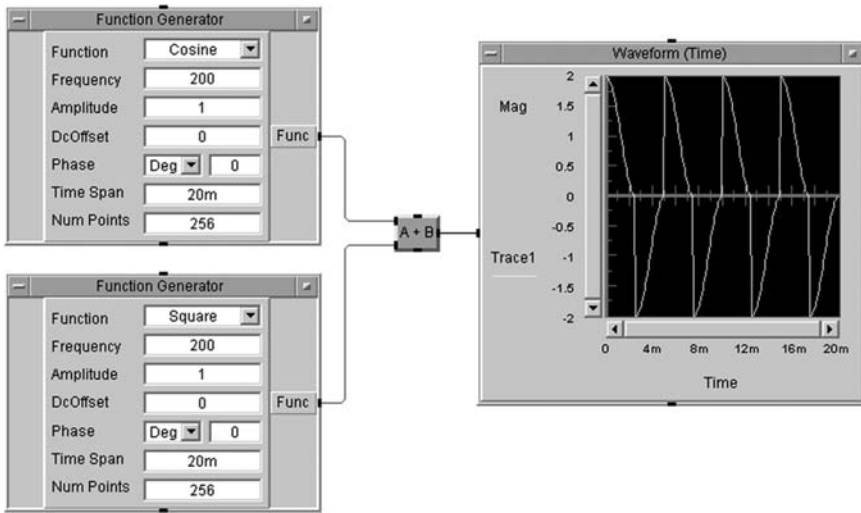


Fig. 14. VEE program for signal generation.

tion devices from National Instruments. Conventional instruments PC plug-in cards are controlled using an *instrument driver*. Instrument address and other parameters can be verified at runtime and changed on the fly, without reconfiguring programs. Agilent VEE Pro automatically handles different datatypes.

In order to use the VEE Pro functions in other programs, the ActiveX Automation Server (ActiveX AS) is used. It supports all popular programming languages, including Visual Basic, C/C++, Visual C#, all .NET-compliant languages, and National Instruments LabVIEW.

Concerning the GUI, VEE Pro permits the development of an interactive interface that includes different types of graphical elements. An example of a VEE Pro VI interface associated with a PID controller virtual instrument is presented in Figure 15.

TestPoint

TestPoint from Keithley [38] is a graphical programming language for test and measurement mainly used to develop virtual instruments based on MIOBs and on traditional instruments with RS232, GPIB, or Ethernet interfaces. Figure 16 shows the front panel of a VI developed with TestPoint. It is a virtual thermometer with three temperature measuring channels. The virtual system includes a Keithley 2700 multimeter with a plugged-in data acquisition module and its software carries out the hardware control, voltage-to-temperature conversion, and measured data display.

TestPoint programming is based on simply drag and drop objects representing graphs, displays, and other parts of measuring and testing VI in a display panel. On an *Action List*, all of the things that the virtual instruments

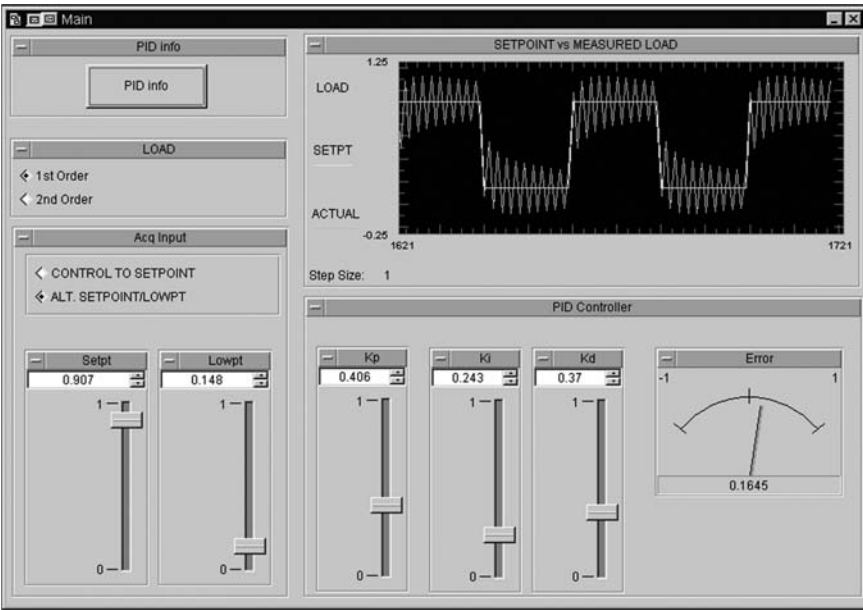


Fig. 15. VEE Pro front panel of a PID controller virtual instrument.

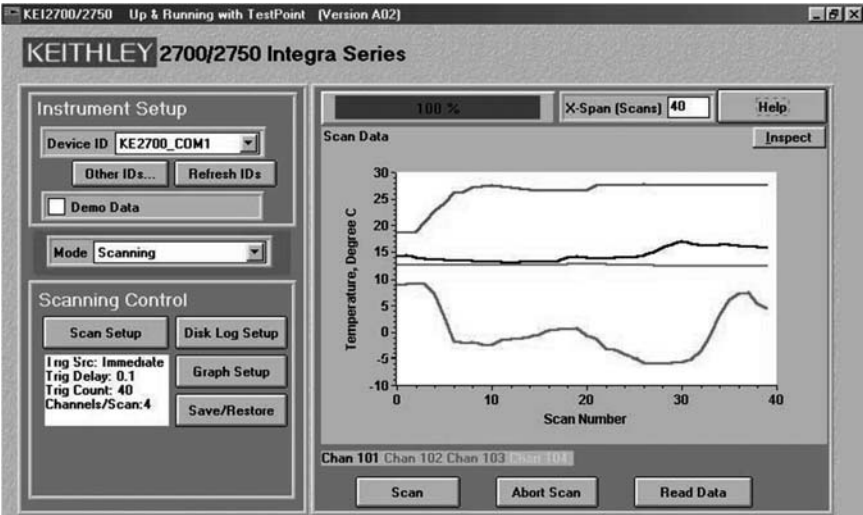


Fig. 16. Multichannel temperature measurement virtual instrument VEE Pro front panel.

will be able to do are included. After this stage, the TestPoint builds automatically the code to run in the VI, which makes this language different from other GPLs such as LabVIEW and VEE Pro. Figure 17 shows a TestPoint virtual PID controller.

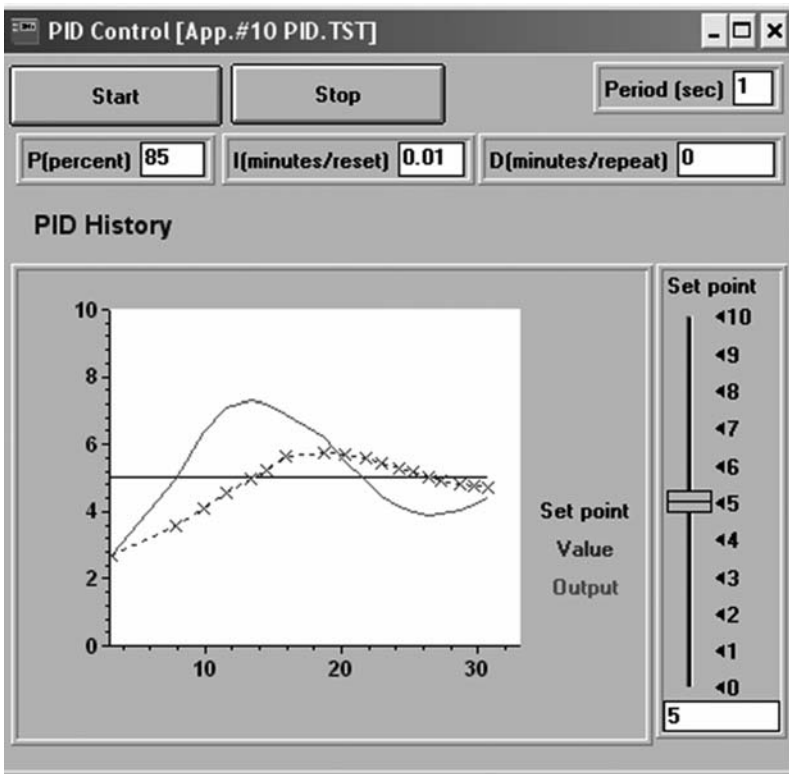


Fig. 17. Virtual PID controller developed in TestPoint.

TestPoint is compatible with industry standard software interfaces such as OLE, ActiveX controls, VBX controls, and MATLAB scripts.

TestPoint is an interesting alternative for virtual instrumentation implementation but is limited in terms of hardware targets (e.g., does not support real-time systems) and operating systems support (runs only on Windows OS).

3.3 VISA and IVI technologies

VISA and IVI are two concepts introduced to simplify virtual instruments programming through the inclusion of a layer between the instrument driver and the hardware. These technologies are oriented for applications based on traditional instruments and do not support MIOBs.

Virtual Instrument Software Architecture (VISA)

VISA [39, 40] is an industry standard implemented by several test and measurement companies widely used as an input/output application programming

interface (I/O API) for communicating with instruments from a PC. VISA's advantages are as follows.

- VISA provides interface independence. It uses many of the same operations to communicate with instruments regardless of the interface type. For example, the VISA command to write an ASCII string to a message-based instrument is the same whether the instrument is Serial, GPIB, USB or Ethernet. This can make interfaces switch more easily and also gives the users who must program instruments for different interfaces use of the same instrument programming technology (Figure 18).
- Programs written using VISA function calls are easily portable from one software platform to another. VISA does this by defining its own datatypes. This prevents problems such as those caused by moving from one platform to another where the size of an integer may be different. In other words, a LabVIEW application written with VISA calls can be easily ported to another platform that supports LabVIEW. Several operating systems are supported, including Windows XP/64, Windows 2000/ME/98, MAC OS X/9/8, some Linux distributions (i.e., Linux complete systems, not only the kernels), as well as some Solaris distributions;
- It provides a very simple-to-use API that has bus-independent functions for most of its I/O functionality. VISA provides the most commonly used functionality for instrumentation in a very compact command set, eliminating the need to learn low-level communication protocols for multiple bus types.

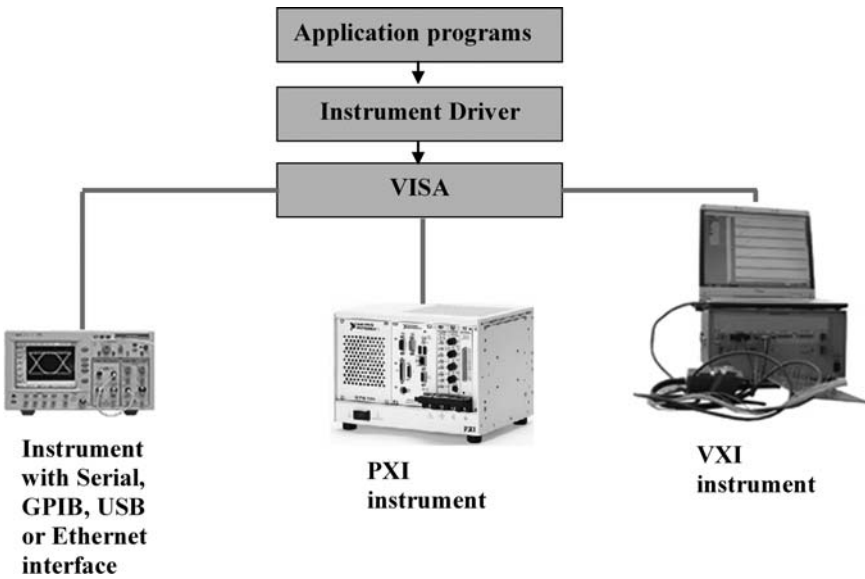


Fig. 18. Multinode virtual instruments based on VISA technology.

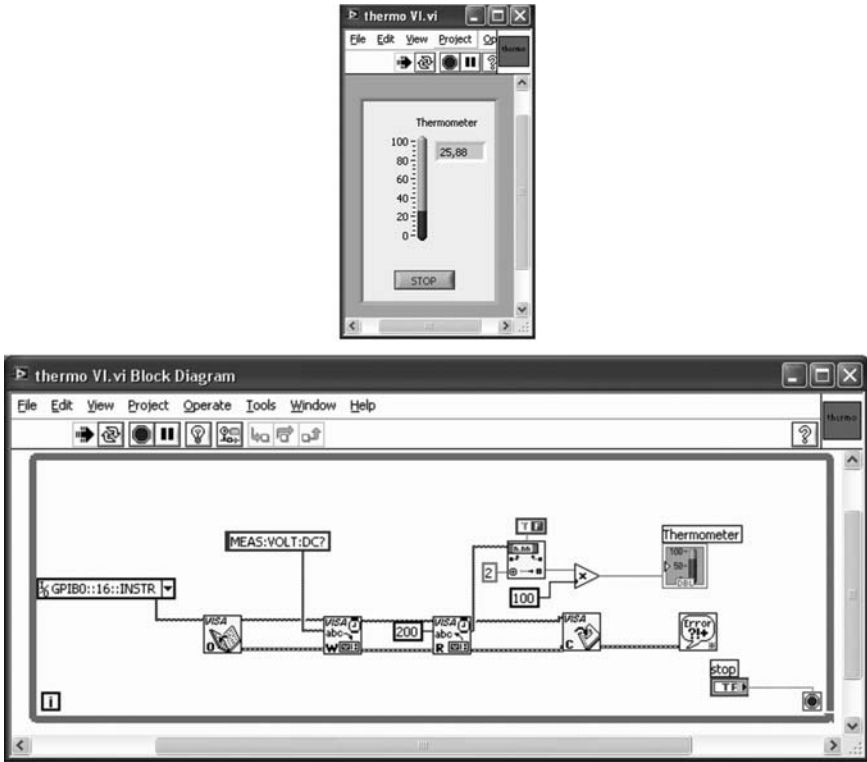


Fig. 19. Virtual thermometer GUI (top); VISA application sequence (bottom).

An example of VISA application where LabVIEW VISA functions are used is presented in Figure 19. It opens a session to a GPIB instrument, performs a write of `MEAS:VOLT:DC?-n` and then queries the device for its response. The string response is converted in double temperature data using the temperature sensor specifications.

The same format would be used in a text-based language such as C++ or Basic and if the used hardware had Serial, USB, Ethernet, IEEE-1394, or any of the other buses that VISA supports. All one would have to do would be to change the Instrument Descriptor connected to the VISA Open.

Interchangeable Virtual Instrument (IVI)

Interchangeable Virtual Instrument (IVI) ([40]– [42]) is also a standard whose goal is to create a driver standard that allows instrument interchangeability. *Interchangeability* means that software for different instruments is compatible with each other (i.e., the software developed for an arbitrary waveform

generator (ARB) from Fluke can be used with an Agilent ARB). According to IVI, the programs work because the IVI standard for ARB drivers requires ARBs from different manufacturers to support the same exact set of function calls.

In IVI, the user application program developed in textual (e.g., C) or graphical programming language (e.g., LabVIEW) communicates with an IVI class driver (e.g., IviDmm class driver) which is represented by a specific instrument driver (e.g., Keithley 2000 instrument driver). The interface between a specific instrument driver and the hardware is performed using the VISA I/O platform. A block diagram that expresses the interdependences on IVI technology is presented in Figure 20.

The IVI architecture breaks the traditional instrument driver into two parts: an instrument-specific driver and a class driver. The instrument-specific driver functions the way traditional instrument drivers had in the past, but with an underlying architecture that is optimized for performance and includes instrument simulation. The class instrument driver contains generic functions for controlling an instrument category and calls the corresponding instrument-specific driver functions at run-time. An VI for measuring or testing purpose can be created with either the class driver or the specific driver, but only programs written with the class driver are interchangeable. Figure 21 shows the application of Figure 19 implemented with IVI.

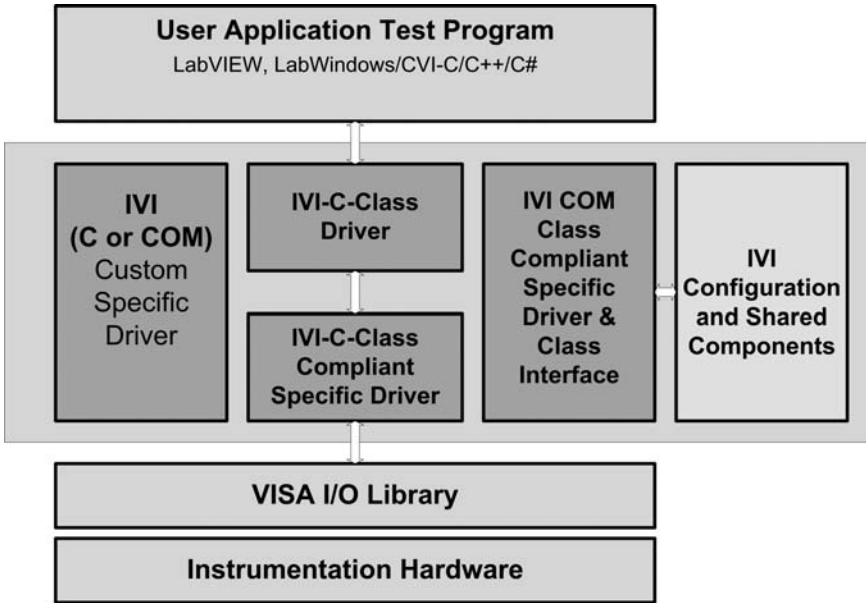


Fig. 20. IVI architecture block diagram.

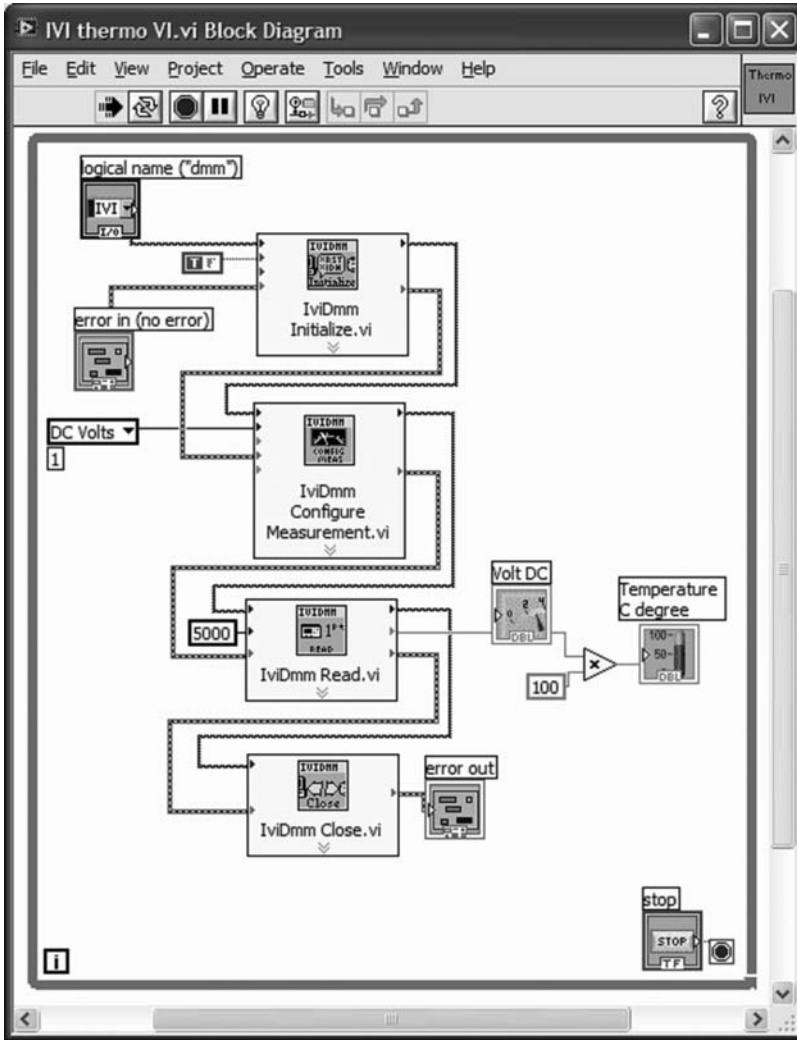


Fig. 21. Virtual thermometer: IVI implementation.

IVI provides the following benefits.

- IVI drivers behave intelligently and perform instrument I/O only when the value of a hardware setting needs modification.
- IVI drivers have a built-in instrument simulation capability that permits one to develop the test code while the hardware components of the VI are temporarily unavailable, or are in the process of being delivered.
- IVI drivers are interchangeable across instruments.

4 Virtual instrumentation operation

The software of virtual instruments is implemented in a computer or in a real-time controller (e.g., NI cRIO-9002) and assigns parameters associated with measurement and control processes. According to application requirements, virtual instruments components are placed at the same location or can be distributed in different locations.

In the first case, the virtual instrument is characterized by a concentrated architecture (software and hardware at the same location) and the user can control the measurement process through the GUI developed for the application. The presence of the VI hardware nearby the software assures a better understanding and verification of the measurement and control performed by the VI. Distributed VIs are multinode wired or wireless systems. Thus, additional software is required for communication among nodes. Usually, client server application architecture is implemented for that purpose.

Client server describes an application architecture in which the *client* requests an *action* or *service* from the provider of service, the *server*. Developing client server applications means developing modular programs. Modular programming separates large applications into smaller constituent pieces to ease development and provide better maintainability.

Client programs request service from a server by sending it a message and include a GUI for presentation of the service provided by an application. Server programs (e.g., LabVIEW Web Server) process client requests by performing the tasks requested by clients. Servers are generally passive as they wait for a client request. During these waiting periods, servers can perform other tasks or perform maintenance. Unlike the client, the server must continually run because clients can request service at any time. Clients, on the other hand, only need to run when they require service. Many server applications allow for multiple clients to request service. How the server manages service to the clients depends on the software designed for that purpose for the server, but generally it is on the basis of first-in, first-out, that is, the first client to be served is the first requesting service. Some applications may require different clients having different priorities. To do this the server must have a method to identify, differentiate, and authenticate the clients requesting service.

Figure 22 shows the general client server interaction and the GUI of a virtual vibration analyzer whose hardware is physically in the place measurements are made. Access to the instrument GUI for visualization is open, but access to instrument's controls can be restricted.

4.1 Virtual laboratories

Virtual laboratories (virtual labs) consist of virtual instrumentation that can be shared and accessed remotely. Usually, the virtual instruments are connected in a local network that includes a PC (local PC) and it is through the

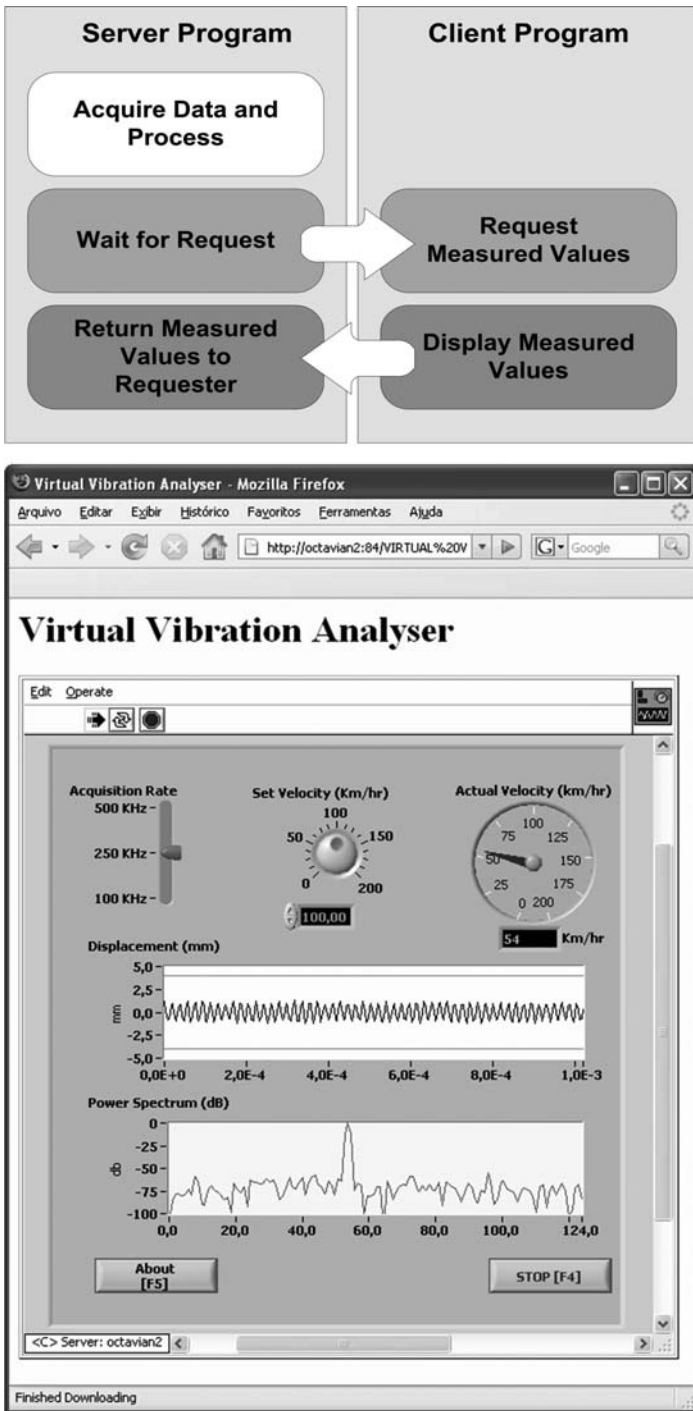


Fig. 22. Client server program interactions (top) and Web virtual vibration analyzer GUI (bottom).

local PC that users' PCs gain access to the instruments. This means that the local PC, operating as a server, and users' PCs (clients) must be connected, that is, must share a network. Because of its worldwide coverage, the more interesting network is, naturally, the Internet and virtual laboratories can be thus used as a hypermedia learning system resource.

Figure 23 represents the block diagram of a possible virtual lab implementation.

When users log in to the local PC over the Internet, they are able to control both the computer and the equipment. Video cameras and specific software are often included in the local system to assure the visual contact with the real lab.

Figure 24 shows the program block diagram of a typical virtual lab.

Virtual labs' design and implementation exploit new tools such as VRML [43] to construct virtual representations of the objects and Java, JavaScript, and Vrmlscript to implement the scripts.

The number of virtual labs available over the Internet increases as institutions operating in the education and formation fields change their educational paradigms to profit from the new communication and information technologies. Distance learning and particularly e-learning are the principal clients of virtual labs, a reality that opens the access to instrumentation and measurements to ordinary people.

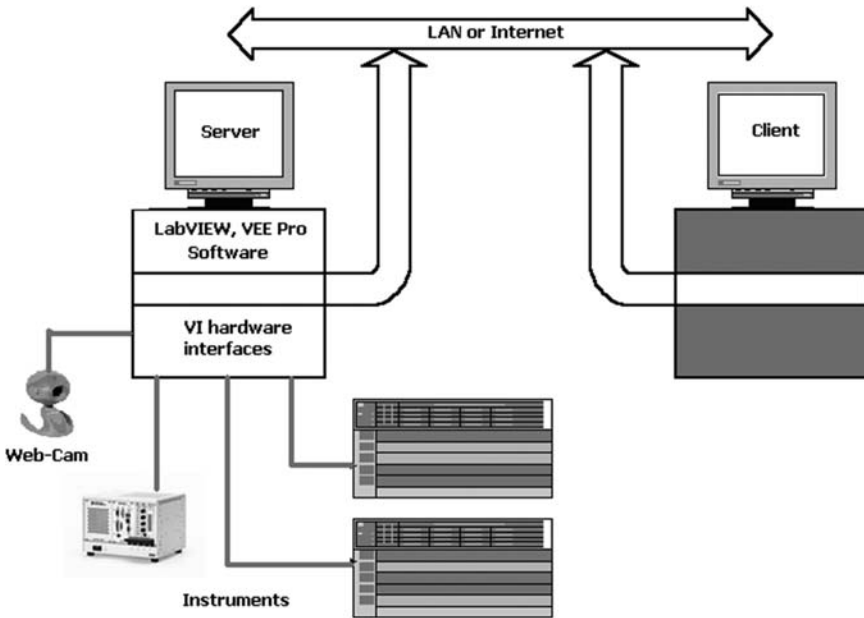


Fig. 23. Block diagram of virtual laboratory complex solution.

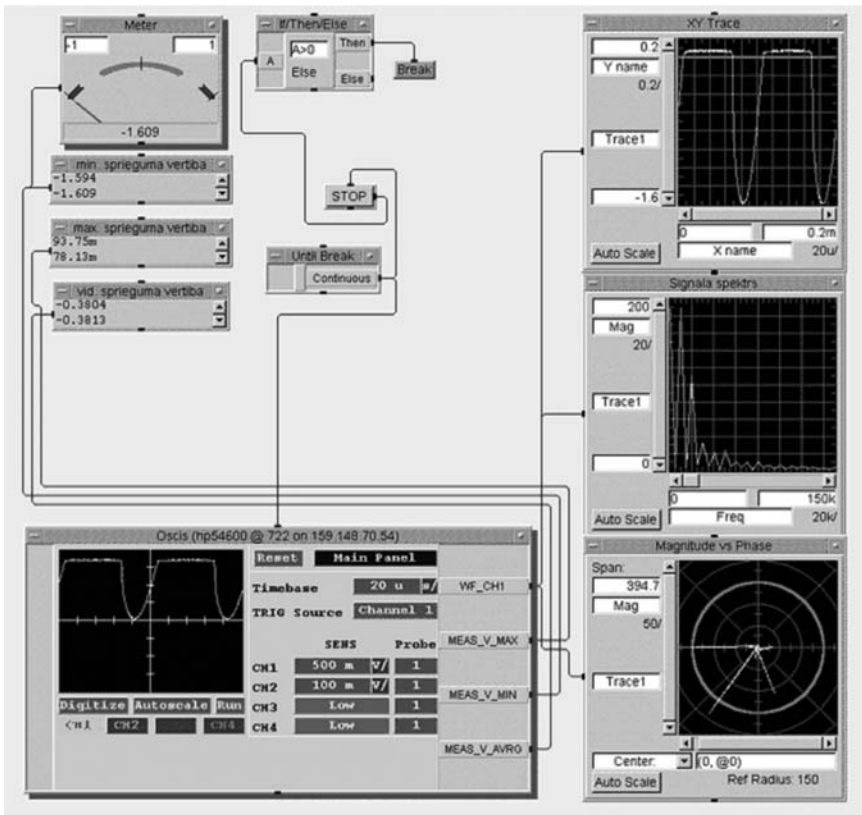


Fig. 24. Program block diagram of a virtual lab.

Table 5. List of acronyms in the chapter.

ADC	Analog-to-Digital Converter
API	Application Programming Interface
ARB	Arbitrary Waveform Generator
ATN	Attention
BC	Bus Controller Block
Bluetooth	Wireless Communication Protocol
CF	Compact Flash Protocol
CompactRIO	Compact Reconfigurable Input Output Technology
DAC	Digital-to-Analog Converter
DAQ	Data Acquisition
DIO	Digital Input Output
DLL	Dynamic Link Library
DUT	Device Under Test
EOI	End of Identify
Ethernet	IEEE803.2 Protocol

FIFO	First Input First Output
FireWire	IEEE1394 protocol
FPGA	Field Programmable Gate Array
GPIB	General-Purpose Interface Bus
GPL	Graphical Programming Language
GUI	Graphical User Interface
I/O	Input/Output
IFC	Interface Clear
II	Interrupt Interface Block
IP	Internet Protocol
IVI	Interchangeable Virtual Instrument
LAN	Local Area Network
LXI	LAN Extension for Instrumentation
MIOB	Multifunction Input Output Board
Mbps	Mega bits per second
M-SADC	Multiplexed single analog-to-digital converter
NI	National Instruments
NI-DAQ drivers	National Instruments Data Acquisition Drivers
OOP	Object-Oriented Programming
OS	Operation System
PC	Personal Computer
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PID	Proportional Integral Differential
PLC	Programmable Logic Controllers
P-MADC	Parallel multiple analog-to-digital converter
PnP	Plug-and-Play
PXI	PCI eXtension for Instrumentation
REN	Remote Enable
RS232	Point-to-Point Serial Protocol
RS485	Multipoint Serial Protocol
RTOS	Real-Time Operation System
RTS	Real-Time System
SCPI	Standard Commands for Programmable Instruments
SCXI	Signal Conditioning Extension for Instrumentation
SE	Single Ended
SRQ	Service Request
TCP/IP	Transfer Control Protocol/Internet Protocol
TS	Synchronization Block
UDP	User Datagram Protocol
UNIX	UNiplexed Information and Computing System (OS)
USB	Universal Serial Bus Protocol
VEE	Visual Engineering Environment
VI	Virtual Instrument

VIM	Vocabulary of general terms in Metrology
Virtual LABs	Virtual Laboratories
VIS	Virtual Instruments Systems
VISA	Virtual Instrument Software Architecture
VME	Versa Module Eurocard Bus
VRML	Virtual Reality Modeling Language
VXI	VME Bus Extension for Instrumentation
WWW	WorldWide Web

References

1. R. Baican, D. S. Neculescu. *Applied Virtual Instrumentation*. University of Ottawa, WIT Press, 2000
2. J.B. Olansen, E. Rosow. *Virtual Bio-Instrumentation: Biomedical, Clinical, and Healthcare Applications in LabVIEW*. National Instruments Series, 2001
3. J. Park, S. Mackay. *Practical Data Acquisition for Instrumentation and Control Systems*. Elsevier, 2003
4. Lingfeng Wang, K. C. Tan. *Modern Industrial Automation Software Design-Principles and Real World Examples*. IEEE Press, Wiley Interscience, New York, 2006
5. *Data Acquisition (DAQ) Hardware*. National Instruments. See at: <http://www.ni.com/dataacquisition/>
6. *Data Acquisition Modules/Cards*. Keithley Instruments, Inc.. See at: <http://www.keithley.com/products/dataacqmodules>
7. IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters. IEEE Std. 1241-2000, 2000
8. *Universal Serial Bus Revision 2.0 Specifications*. USB organization, 2006. See at: www.usb.org
9. *IEEE 1394 for Your Test and Measurement PC*. National Instruments. See at: <http://volt.ni.com/niwc/products/ieee1394.jsp>
10. T.J. Shea, J.A. Mead, P. Cerniglia, C.M. Degen. Evaluation of IEEE 1394 Serial bus for distributed data acquisition. *Proceedings of the 1997 Particle Accelerator Conference*, Volume 2, pp. 2502-2504, Vancouver, Canada. See at: <http://ieeexplore.ieee.org/iel4/6051/16227/00751254.pdf?arnumber=751254>
11. *Two-Plane Balancing Using LabVIEW PDA and NI CF-6004 CompactFlash Data Acquisition Card*. National Instruments. See at: <http://zone.ni.com/devzone/cda/tut/p/id/4381>
12. *CF+ & CompactFlash Specification Revision 4.1*. Compact Flash Association. See at: <http://www.compactflash.org/>
13. D. Pleasant. LAN-based measurement triggering using LXI instrumentation. *Proceedings of Autotestcon 2005*, pp 44- 48 , Orlando, USA, 2005
14. *LAN eXtension for Instrumentation standard. It's Time To Put the Power of Ethernet to Work in Test Systems*. See at: <http://www.lxistandard.org/home>
15. *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*. IEEE Std. 1588-2002, 2002
16. *Benchmarking Single-Point Performance on National Instruments Real-Time Hardware*. National Instruments. See at: <http://zone.ni.com/devzone/cda/tut/p/id/5423>

17. N.P. Mahalik. *Fieldbus Technology: Industrial Network Standards for Real-Time Distributed Control*. Springer-Verlag, Berlin, 2003.
18. *PXI: The Industry Standard Platform for Instrumentation*. National Instruments. See at: <http://zone.ni.com/devzone/cda/tut/p/id/3889>
19. *PXI overview*. PXI System Alliance. See at: <http://www.pxisa.org/overview.htm>
20. *cRIO - FPGA Design, Development and Programming Tutorial*. National Instruments. See at: <http://zone.ni.com/devzone/cda/tut/p/id/3358>
21. *NI CompactRIO Control and Acquisition System*. National Instruments. See at: <http://www.ni.com/compactrio/whatis.htm>
22. National Instruments. *Advanced Topics in Labwindows/CVIO*. Prentice Hall, upper Saddle River, NJ, 2004
23. D. Hanselman, B. Littlefieldand. *Mastering MATLAB 7*. Pearson/Prentice Hall, upper Saddle River, NJ, 2005
24. *MATLAB - The Language of Technical Computing*. MathWorks. See at: <http://www.mathworks.com/products/matlab/?BB=1>
25. B. Meyer. *An Object-Oriented Environment: Principles and Application*. Prentice Hall, upper Saddle River, NJ 1994
26. D. Flanagan. *Java in a Nutshell - 5th Edition*, O'Reilly, Sebastopol, CA, 2005
27. D. Flanagan. *Java Examples in a Nutshell - Third Edition*, O'Reilly, Sebastopol, CA, 2004
28. *Java Tutorials - Object-Oriented Programming Concepts*. Sun Microsystems. See at: <http://java.sun.com/docs/books/tutorial/java/concepts/>
29. *The Source for Java Developers*. Sun Microsystems. See at: <http://java.sun.com/>
30. *Java(tm) Communication API*. Sun Microsystems. See at: <http://www.sun.com/download/products.xml?id=43208d3d>
31. *Sun Developer Network (ISDN) Code Samples and Apps*. Sun Microsystems. See at: <http://java.sun.com/applets/>
32. P.F. Franklin. *LXI Clears the Way to Smarter Instruments*. Keithley Instruments, Inc.. See at: www.keithley.com/data?asset=50572
33. W.W. Wadge. *Lucid, the Dataflow Programming Language (Apic Studies in Data Processing)*. Academic Press, San Diego, CA 1985
34. *LabVIEW: 20 Years of Innovation*. National Instruments. See at: <http://www.ni.com/labview/>
35. G.W. Johnson, R. Jennings. *LabVIEW Graphical Programming*, 3rd edition. McGraw-Hill Professional, New york, 2001
36. T.E. Hulbert, R. B. Angus. Preparing a virtual engineering environment laboratory instructional package. *Proceedings of the 2002 American Society for Engineering Education Annual Conference and Exposition*, Boston, MA, 2002. See at: www.ni.com/pdf/academic/us/journals/lv02\51.pdf
37. *VEE Pro Briefs and Papers*. Agilent Technologies. See at: http://adn.tm.agilent.com/index.cgi?CONTENT_ID=771
38. *TestPoint Development Package*. Keithley Instruments, Inc. See at: <http://www.keithley.com/products/software/testpointdevpackage/?mn=testpoint>
39. W. Boyes. *Instrumentation Reference Book - Third Edition*. Butterworth Heinemann, Oxford, 2003

40. L. Gutterman. Integrating VISA, IVI and ATEasy to Migrate Legacy Test Systems. *Proc. of IEEE Autotestcon 2004*, San Antonio, TX, pp. 74-79, September 2004
41. *IVI -Driver Architecture. Interchangeable Virtual Instrument Foundation*. See at: <http://www.ivifoundation.org/About\%20IVI/Architecture.htm>
42. G.E. Geathers. *The IVI Foundation Signal Interface; A New Industry Standard*. See at: <http://ieeexplore.ieee.org/iel5/7542/20537/00949422.pdf>
43. A.L. Ames, D.R. Nadeau, J. L. Moreland. *The VRML 2.0 Sourcebook*. John Wiley & Sons, New york, 1997

Internet-Enabled Metrology

Tanasko Tasić

METREL d.d., Ljubljanska c. 77, 1354 Horjul, Slovenia
tanasko.tasic@metrel.si

Summary. State-of-the-art information technology enables a variety of specific metrological applications. The first benefit for the metrology community is in significantly increased remote functionality of the measuring systems (e.g., connected in distributed legal metrology measuring systems, or remote operation of the measuring instruments in severe environmental conditions). The next significant improvement comes from the introduction of new metrological services (e.g., time service, remote calibration, and remote software validation). Finally, it is important to mention the increased availability of the metrology-related information which is available on the World Wide Web. Presented approaches and solutions are not exclusively specific for the metrology community; however, it is important to mention them in order to make metrologists aware of their existence. This chapter is not an attempt at giving the complete overview of available solutions, but rather the first insight into the mentioned areas which will be the starting point for further investigations.

Key words: Distributed measuring systems, remote instruments operation, remote calibration, metrological database systems, Internet software validation services

1 Introduction

This chapter gives an overview of metrology-related applications whose development has been enabled by the spread of Internet services. It is not aimed at metrological IT experts but rather it is directed at metrologists (manufacturers of measuring instruments, scientists in metrological laboratories) who are not deeply involved in IT issues. The chapter is intended to give such readers an insight into this developing field and to help them solve problems with state-of-the art technologies.

The focus of the chapter is not in the explanation of the technical implementations but on giving an overview of the applicable benefits, thus constructional solutions are not explained in deep detail. Furthermore, it is not an attempt to give a complete presentation of all available services, but an informative overview of the common applications which may give the reader useful

information, indicate an existing solution for his or her problem, or suggest an idea for developing some new solution. There are many links to Internet sites with interesting contents; exploring these sites is left to the reader.

When mentioning Internet-enabled metrology, people usually have in mind issues related to Internet-enabled calibrations or Web sites with software validation tools. However, state-of-the-art IT enables a much wider variety of metrological applications. The benefits to metrology arising from the spread of use of Internet may be organized in three groups:

- Functionality of measuring instruments (distributed measuring systems)
- New metrological services (time service, remote calibration, software validation)
- Significantly increased availability of metrology-related information

In the context of this chapter, information technology covers more than the public communication networks, delivered using a range of hardware realisations and different operators: optical cable, cable TV networks, wireless, fixed wire and mobile telephony (GSM, GPRS, UMTS), or public networks using secure (e.g., VPN) channels. In distributed measuring systems, especially in cases when distribution companies possess their own distribution networks (e.g., electricity), common ways of data exchange are PLC (Power Line Carrier), DLC (Distribution Line Carrier), RADIANT (Radio Application Network), and ZigBee [2] (IEEE 802.15.4), for example. Additionally, the supporting communication protocols may be general-purpose internet protocols (SMTP, HTTP, FTP) or application specific, for example, DLMS (Device Language Message Specification; IEC 62056-46, IEC 62056-61).

From the point of view of IT security, metrology-related IT applications apply standard approaches (e.g., HTTPS, FTPS, public key infrastructure), mostly for the protection of data (to ensure correctness of measuring data) and for the authentication of involved parties.

In the background, very often as the central point of a distributed measuring system there is a database system. The applications again vary from data collection/billing or dynamic tariff calculation (electricity) in legal metrology, medical diagnostics (databases of particular health states' pattern signals intended for medical diagnostics), general information about metrological capabilities (e.g., national), monitoring and managing of distributed metrology systems (e.g., information about the bodies performing legal metrology tasks or information about the status of the measuring instruments, information about the instruments' calibration parameters during the life cycle), and so on.

Metrology institutions worldwide are fully aware of the importance of these issues. The International Organization of Legal Metrology (OIML) had organised the seminar on measuring instruments software as early as 1999. The International Bureau of Weights and Measures (BIPM) and leading national metrology institutes organize a series of conferences on the impact of the information technology in metrology:

- BIPM–NPL Workshop on the Impact of Information Technology in Metrology, Teddington, UK, 16–19 September 2002
- NMIJ–BIPM Workshop on the Impact of Information Technology in Metrology, Tsukuba, Japan, 18–20 May 2005 [13]
- PTB–BIPM Workshop on the Impact of Information Technology in Metrology, Berlin, Germany, 4–8 June 2007 [22]

It is possible to find much valuable metrology-related information in the proceedings of those conferences.

2 Functionality of measuring instruments (distributed measuring systems)

As said in a presentation at the FASIT [25] workshop, the development of legal measuring instruments began with the ‘iron age’ (mechanical measuring instruments), continued with the ‘electronic age’ followed by the ‘software’, and now we are already in the ‘communication age’. This statement is undoubtedly true for all areas of metrology that benefit from the state-of-the-art computer communication, Web, and database technologies.

As in all other modern devices, software is taking over more and more functions in measuring instruments and systems. Besides software running on the measuring instrument computer, communications and database systems are implemented to enable distributed measuring systems, measuring data storage and their subsequent processing. In addition to gathering, storing, and processing the measurement data, there are also the functionalities of maintenance of measuring systems, calibration, and surveillance of the measuring instruments metrological status. Most of these services are available via public communications networks; others will be available very soon.

This progress is undoubtedly useful for the users, because it enables faster measurements, higher accuracy, and opens up the possibility of various analyses and further processing. For the manufacturer the new technology simplifies realizing complex functions and gives flexibility for meeting the wishes of their customers and measuring instrument maintenance.

2.1 Legal metrology applications

Legal metrology instruments are subject to several instances of conformity assessments and inspections during their life cycle. Initial conformity assessment procedure (type approval) is the prerequisite for obtaining status of the legal metrology instrument. First verification is necessary before the first use, followed by regular and extraordinary verifications and inspection examinations during the measuring instrument’s life cycle.

Institutions that perform these conformity assessment procedures are both governmental and private bodies. In recent years, more and more conformance

assessment tasks are being entrusted to private bodies. Irrespective of the trends towards privatisation of conformity assessment services and towards removing barriers to trade, it is required to maintain an adequate level of consumer protection. The role of state institutions (who are responsible for operation of national metrology systems) is becoming more focused on the collection and analysis of reports on conformity assessment procedures (from private bodies performing, e.g., verification of legal measuring instruments) and occasional metrological spot checks of measuring instruments. As a consequence, huge amounts of data need to be exchanged between these institutions in order to maintain a suitable quality of the national legal metrology system. This leads to additional measuring instruments functionality, which is nowadays mostly implemented by software and information technologies.

An illustrative example is the operation of electricity meters in the deregulated energy market. In such a system, a variety of participants needs to have access to the electricity meter, each of them with its own access rights (Figure 1). The end customer receives electricity from one of the distributors. Energy consumption is processed by a metering data provider. The price of energy necessary contains shares of the electricity supplier, owner of the distribution network, operator of the IT network, and electricity distributor. Additional participants that also need to have access to the meter are members

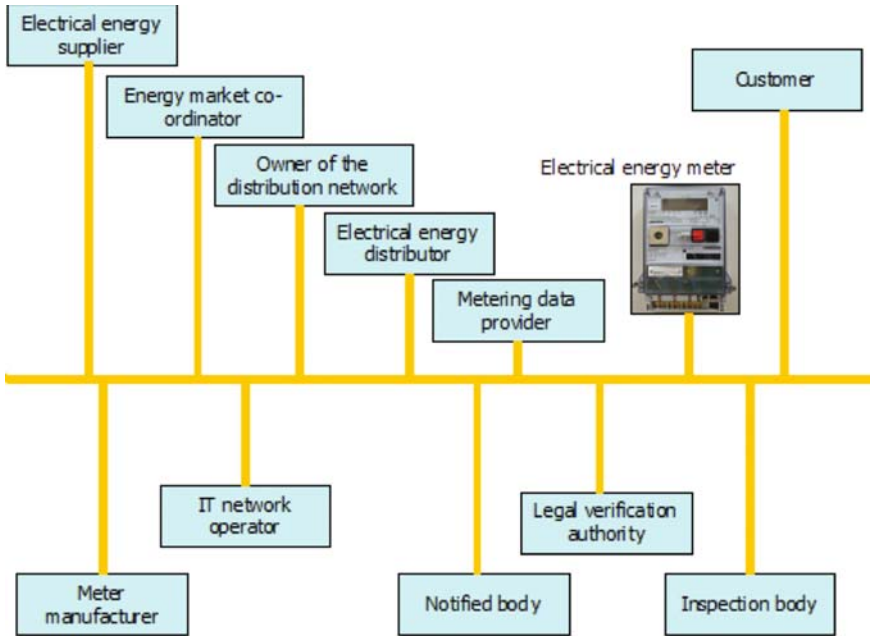


Fig. 1. Participants in modern distributed measuring system, an example from legal metrology.

of the legal metrology system: the notified body (performing initial conformity assessment), the legal verification authority (performing legal verifications), and the inspection body (usually independent body performing metrological surveillance of the measuring instruments in use). And of course, the meter manufacturer and/or her representative need to have access for maintenance purposes as well.

The usefulness of remote access to meters is well illustrated by the example of operators of measuring instruments. For utility companies (electricity distribution companies, for instance) remote access significantly decreases the costs of collection of measurement data and maintenance of measuring instruments. The highest benefit is definitely the possibility of remote update of software in measuring instruments already installed at the place of use. In the case of detection of a serious bug in the measuring instrument's software, the software in all instruments of that type has to be updated. If we take as an example electricity meters installed in remote mountain areas, the cost for a technician to visit an installation and physically load the new software might well be more than the original meter retail price.

Taking into consideration the variety of involved parties, such a system has to be adequately coordinated. Authentication of every participant is necessary to ensure that only entitled personnel have access to particular meter data or functionality and to assure that the correct customer is charged for the consumption. Generally speaking, the risks of fraud in metrological software applications are not so high, compared to other areas (e.g., e-banking). However, they are certainly not negligible, as they vary from possible attempts to amend energy consumption data by market competitors, monitoring of energy consumption in order to detect the absence of the residents from home (e.g., by burglars), or just-for-fun hacker attacks.

An implementation of such a system in praxis has been validated during the project SELMA [15, 23] '*Sicherer ELEktronischer Messdaten-Austausch*', secure electronic measurement data exchange). Participants in the project were the manufacturers of measuring instruments, universities, electricity distribution companies, and state institutes responsible for metrology and information security.

A similar approach has been implemented in road-traffic enforcement networks [6]. In such systems automated vehicle velocity measuring stations (police radar) transmit files with evidence of offences to the place of processing. For transmission of these files it is necessary to ensure their integrity and confidentiality, as well as proper authentication of participants in the process.

Being aware of the consequences of IT-related development in the field of measuring instruments [4], legal metrology institutions have prepared guidance documents that support harmonised software validation [27] like WELMEC [11] '*Software Guide (Measuring Instruments Directive 2004/22/EC)*', and OIML [18] '*General requirements for software controlled measuring instruments*'.

2.2 Remote operation of measuring instruments

Remote operation of measuring instruments is important in circumstances where environmental conditions are such that human presence at the time and place of performance of measurements is not possible (e.g., in a mine, in a weather station, near steel production furnaces, in space vessels) or in cases where permanent human presence is not practicable. There are many examples of the latter in industry (e.g., measurements after exposure of measuring instruments to stabilised environmental conditions or long-lasting measurements with periodic changes of reference values). It is important to stress that prerequisites for remote operation of measuring instruments are already present in the form of existing IT infrastructure. An example of remote-controlled factory testing of an electrical power monitoring system is illustrated in Figure 2. Several measuring instruments for monitoring the electrical power parameters are connected in a computer network by different communication interfaces. Testing the functionality of measuring instruments comprises several long-lasting measurements. One session of measurement of electrical energy as a rule lasts several days. During that time reference signals need to be changed several times.

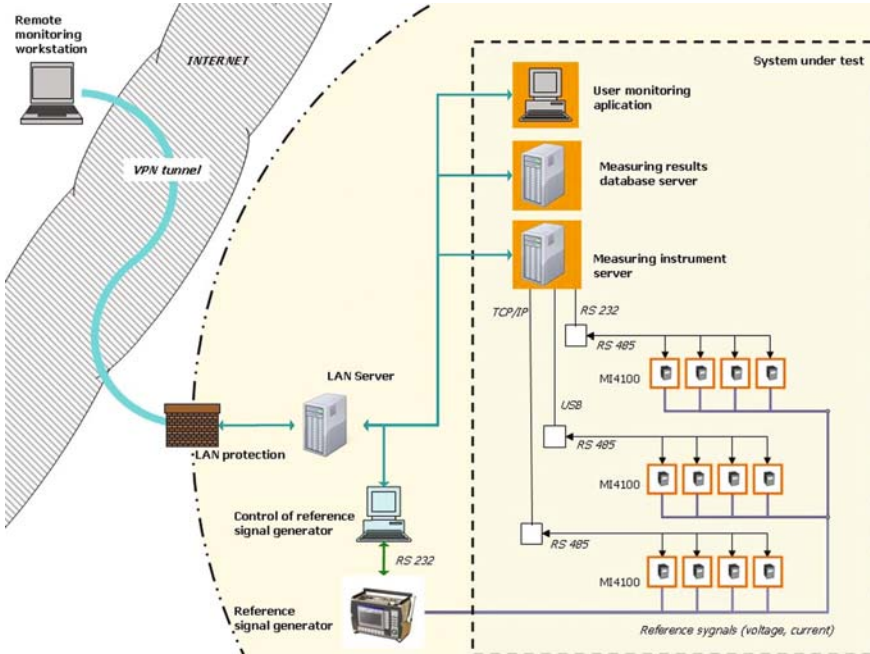


Fig. 2. Remote testing of measuring instruments: an example from the company METREL d.d.

In addition, it is necessary to check the functioning of the components of the system (particular measuring instruments, measuring instruments server, database server, and client application). These checks need to be performed every couple of hours. Once initiated all these checks may be performed remotely, for example, from home, generally much more practical than driving to the testing laboratory every time that some activity is necessary.

Some of the software functionalities necessary for support of the remote measurements are already embedded in operating systems (such as Microsoft Windows virtual private network and remote desktop connections in the explained example). Another solution is available in T&M automation tools, for example, publishing software control applications as Web pages, which can then be remotely accessed (e.g., this functionality is available in National Instruments LabVIEW).

3 Internet-supported metrological services

Besides distributed measuring systems that are intended for online uninterrupted service, there are very many Internet-supported services related to various metrological activities. These can be classified as follows.

3.1 General metrology related services

An example of such service is the Internet Time Service (ITS) [20], which is intended to synchronize computer clocks via the Internet. A reference clock runs on a WWW server. The service responds to time requests from any Internet client by sending the time and estimated delay information in several formats including the DAYTIME, TIME (older realisations), and NTP protocols. The Network Time Protocol (NTP) is one of the most accurate and flexible means of sending time over the Internet [16]. It can be used by almost any type of computer. The protocol is designed to compensate for some, but not all, network time delays between the server and the client. NTP is most successful across local area networks and can give an accuracy as good as a few milliseconds. On the World Wide Web, however, time transfer delays are at the mercy of server traffic and network bottlenecks, and accuracy figures cannot be quoted as easily. NTP conveniently supports security measures for users who want more reassurance concerning the origin of the timestamp (Authenticated NTP Services), rather than insecure NTP. On the client side, the user needs to have the software that can request time over the Internet. A version of the NTP client software used to synchronize computer clocks is called Simple Network Time Protocol (SNTP).

The accuracy of such services is rarely better than 0.1 s, for better accuracies it is necessary to use other methods and services, for example, Common-View GPS method.

3.2 Internet enabled/supported calibrations

The ‘old-fashioned’ procedure of performing calibrations, which is still used by the majority of customers and laboratories, involves transporting the instrument to be calibrated to the calibration laboratory [5]. The instrument has to be physically present in the calibration laboratory during the whole calibration process. Such an approach has several disadvantages, like:

- Long instrument downtime (for home laboratory);
- Costs of the transportation;
- Calibration of the instrument in conditions different from the conditions of its routine use (e.g., other personnel, instrumentation environment, climatic environment);
- Danger of damage during the transportation.

The spread of the Internet as a communication medium and the availability of the continuously improved measurement standards enable the dissemination of calibration values from higher-level standards laboratories in a different way. As the result, many institutions offer remote access over the Internet to more and more calibration and measurement services.

There are several ways in which the Internet is used for measurement process support. From the point of view of logistics, the simplest way is if there is no need to physically transport anything (no transfer standards). This way is possible for the dissemination of a limited number of physical quantities (e.g., time, frequency), or in situations when the measurement signal can be transmitted from the field laboratory to the reference laboratory. Other methods require either the presence of the reference standard (reference material) or presence of the transfer standard at the place where the calibration is performed.

From the point of view of laboratories that need calibrations of their instruments the benefits of Internet-enabled calibrations may be summarised as follows:

- Reduced costs;
- Direct access for more customers to the higher-level laboratories;
- The measurement standards belonging to the remote laboratory do not need to be transported;
- The standards are calibrated under normal conditions of use in their home laboratory;
- The ‘down time’ for the remote laboratory is kept to a minimum;
- Uncertainties are calculated online by reference software in the reference laboratory;
- The results can be reviewed online before the calibration is completed;
- Instructions and procedures can be conveyed over the Internet;
- The calibration conditions can be recorded using a video or digital camera;
- The calibration environment can be recorded;

- The calibration chain is cut down to one link;
- Calibration can be performed at any time, day or night;
- The expertise of the higher-level laboratory can be transferred.

Some generic IT-related issues have to be addressed:

- Security of calibration information and results as they are transferred over the Internet (data integrity, authentication, access control, and protection from viruses and worms);
- Smooth operation of Internet calibrations through site fire-walls, overcoming problems of two-way communication of data through fire-walls without compromising security;
- Division of software between Web pages and compiled measurement code;
- Appropriate choice of computer–instrument interfaces at the remote site together with future proofing to accommodate evolution in computer software and hardware;
- Storage of calibration results in a database for recall by users during the calibrations and the use of data warehousing to provide long-term access to calibration history;
- Guidance on the procedures necessary to meet accreditation requirements in general;
- Promotion of the technology to metrology areas where it may be applicable and advantageous but which have not yet taken it up.

For Internet-enabled metrology to be successful, it is essential that the software works reliably, both in terms of integrity of transmitted data and the smoothness of operation from the point of view of the user. It is also necessary to prove that the traceability of a calibration when it is carried out at a remote site can be maintained. This is a nontrivial issue as traceability covers factors such as the suitability of the calibration environment, the correct operation of the measurement equipment, and, perhaps of most importance in this case, the ability of the calibration staff at the remote laboratory to carry out the required tests under the guidance provided through Web pages.

Entirely remote calibrations

An excellent example of the application of the modern technologies in metrology is use of Global Positioning System (GPS) satellite signals and the internet for time synchronization and frequency calibration. The satellite constellation consists of 24 satellites (the first was launched in 1978 and the 24th in 1994). They transmit signals that can be detected by receivers on the ground. The satellites are positioned in six Earth-centred orbital planes with four satellites in each plane. This means that 100 percent of the time, any point on the Earth's surface can receive signals from at least six satellites. The main application of the GPS system is the determination of the position of objects on the Earth's surface (primarily for military purposes). However, this is not the only




Fig. 3. The common view method.

application of the system. Each satellite carries either rubidium or caesium oscillators, or a combination of both, which are synchronised with UTC [12] USNO [17] and UTC NIST [19]. Those UTCs are maintained within 100 ns of each other, and the frequency offset between the two timescales is less than $1 \cdot 10^{-13}$.

GPS receivers at locations A and B receive time information from the same satellite (which is in common view for both locations, A and B). The common-view method compares two clocks or oscillators located in different places. Unlike one-way measurements that compare a clock or oscillator to GPS, a common-view measurement compares two clocks or oscillators to each other. The first scientific publications describing this system, called ‘Common View’ appeared as early as 1980 [3].

There are several types of time and frequency measurements that utilize GPS signals; some of them require intensive calculations that used to take up to several weeks in the past. By combining the common-view technique with the Internet, it is possible to build a common-view network that processes data in near real-time [14].

The lowest uncertainties currently achieved using the GPS measurement techniques for carrier-phase common-view are less than 500 ps for time and less than $5 \cdot 10^{-15}$ for frequency.

This approach may be the starting point for realisation/dissemination of all physical quantities that are related to frequency (e.g., voltage).

Internet-supported calibrations with additional transmission of the measurement signal

This idea is explained using the example of the remote calibration of practical lengths by using low-coherence interferometry and optical fibre network [10]. The laboratory requiring calibration has its own low-coherence interferometer (LabI) and performs measurements on a unit under calibration (UUC). The measurement signal, in the form of the interferometer signal, is transmitted via optical fibre to the reference laboratory. The reference laboratory with its associated software then determines the parameters of the UUC.

A major limitation of this method is the signal loss in optical fibre between the customer premises and reference laboratory. With the state-of-the art technology affordable distance is about 20 km.

Internet-supported calibrations with the local reference standard




Fig. 4. Internet-supported calibrations with additional transmission of the measurement signal. ITI: IT Interface, UUC: unit under calibration, LabI: Instrument in customer's laboratory, here used as the light source, Std.: higher level measuring standard, Ref. SW: software associated with higher-level measuring standard.

This approach uses a local reference standard or material (Ref S) with stable, known characteristics for the calibration of an instrument under calibration (IUC). The IUC measures characteristics of the Ref S and sends the results via the Internet to the reference laboratory. After analysis of the measuring results, software in the reference laboratory calculates the calibration parameters for the IUC.

This approach is used for the calibration of instruments in many areas such as the measurement of impedance (iPIMMS [9], Primary Impedance Measurement Software for Impedance calibration of Vector Network Analysers), spectrophotometry (iColour Calibration Visible Diode Array Spectrophotometer [9]), pressure measurement [24], or ionising radiation [29].




Fig. 5. Internet-supported calibrations with the local reference standard. ITI: IT interface, Ref. S: Reference sample (material), IUC: Instrument under calibration (in customer’s laboratory), Ref. SW: Reference laboratory’s software.

Internet-supported calibrations with travelling reference standard

This approach is applied in areas where local reference materials or reference standards of appropriate quality are not available. The transportation of travelling standards from the reference laboratory to the customer’s laboratory is generally more convenient than the transportation of the instrument to the reference laboratory. Related costs are likely to be less as well. The procedure of calibration is the same as the one for the case of Internet-supported calibrations with local reference standard. Travelling reference standards are not yet available with very high precisions; the uncertainties for this type of calibration are usually higher than the uncertainties in reference laboratories.



Fig. 6. Internet-supported calibrations with travelling reference standard. ITI: IT Interface, Trans. Std: Travelling reference standard, IUC: Instrument under calibration (in customer’s laboratory), Ref. SW: Reference laboratory’s software.

3.3 Availability of specific metrological software validation services

Modern laboratory measurement systems are very often highly automated in data acquisition as well as in postprocessing of measured data [8]. Typically, software components of such systems may consist of both commercially available software packages and custom-made software, developed by the laboratory staff. In order to ensure confidence in the results provided by such systems, the software must be proven to be fit for purpose, so it is necessary to be able to properly verify it in whole and partially.

Web-based software validation tools make possible for different users around the world (developers, evaluators, or laboratories) to validate their

software modules using black-box testing methods. Two approaches to validation of software are identified as appropriate for Web-based software validation tools: validation with the reference datasets and validation with the reference software.

Examples of available tools are the following:

- Web pages with reference software for validation of the mathematical functions;
- Web pages with reference software for validation of the standardised metrological functions;
- Web pages with reference datasets for validation of the metrological software.

Reference software for validation of the mathematical functions

Reference software enables testing of the software components that realise standard mathematical/statistical functions (e.g., mean, standard deviation, square root, etc.). The user validates his software by comparing output parameters from his and reference software in response to the same input parameters. Phases in testing process are the following:

- Generate reference datasets and corresponding reference results consistent with the computational aim;
- Apply the software under test to the reference datasets to obtain test results;
- Compare the test results with the reference results.

Generation of the reference data may be organised in several ways:

- Applets that run on the client (or user's) machine. The user supplies values for a number of input parameters to the applet, which then generates its output in two windows. One window contains the reference dataset and the other the reference results and additional information;
- Servlets that run on the host machine. In the same way as for an applet, the user supplies values for a number of input parameters to the servlet, which then generates its output as a text page (viewed using an Internet browser) that may be saved to the user's machine;
- Web services that run on the host machine. Using this mechanism the data generators may be called directly (via the Internet) from software (such as a test harness) running on the client's (or user's) machine.

Reference software for validation of standardised metrological functions

This kind of reference software is suitable for validation of software components that realise algorithms based on well-known international technical

standards (e.g., IEC 60584-1,2 Thermocouples) [28]. Its purpose is to enable developers (metrologists from laboratories that develop their own software components) to validate their software components. The application is simple; the user just needs to enter input variables simultaneously to both the software under test and the reference software and afterwards compare the results from both.

Such an application also allows laboratories involved in an intercomparison to check their software modules before actual measurements begin and thus eliminate any potential source of difference due to software errors

Some additional issues arise in connection with this approach. The first one concerns the validation of the reference software [1]. The application that claims to be the reference needs to be suitably validated. The next question relates to IT security: the user needs to be 100% sure that the results come from the genuine reference software and not from some fake (malicious) copy. This question can be resolved with the use of appropriate Internet protocols (e.g. `https`).

The software characteristic which is usually validated with this approach is functionality (numerical correctness), one of the most important characteristics for metrologists.

Reference datasets for validation of metrological software

Instead of dynamically generated reference data, as explained in ‘Reference software for validation of the mathematical functions’, the reference datasets may be static, dedicated for validation of standardised metrological functions. Reference datasets must contain both input and output parameters. During the software testing the user applies input parameters, runs the software and then compares the results with the reference one(s).

4 Availability of the metrology-related data

The ease of distribution of information is the most important benefit of the Internet. In the metrological community, the most important information is provided by Web-hosted databases with relevant metrological information:

- Calibration/metrological resources and their capabilities. The most comprehensive database collection containing worldwide metrological information is available on the website of the Bureau international des poids et mesures (BIPM: <http://kcdb.bipm.org>). It contains information about scientific work, history and technical realisation of the primary measurement standards, calibration and measurement capabilities in particular countries, units of measurement, technical committees, measurement standards, publications, conferences, key comparisons, reference materials and other metrological information.

- Research-oriented databases, for example, PTB databases for vacuum-metrology medical research (heart bio signals).
- Information exchange centres, for example, ‘virtual institutes’. These sites are dedicated to experts in particular areas of metrology. Some examples are: the Virtual Institute for Reference Materials (VIRM, <http://www.virm.net/>), the Virtual Institute for Thermal Metrology (Evitherm, <http://www.evitherm.org/>), and the Virtual Institute of Energy Metrology (<https://bi.offis.de/viem/tiki-index.php>).
- Legal metrology databases are intended mostly for public information about nationally or internationally (EU-type approval or OIML-type approval certificates) approved measuring instruments, metrological standards and regulations, ongoing projects, and so on.
- Metrology-related public information: information about approved measuring instruments, about their intended use, about control bodies, and inspections.
- Databases of national metrology institutes contain a variety of data: organisational, calibration and verification capabilities, scientific background for metrological procedures, presentation of achievements, metrological advice for general public, and so on.
- Databases of reference materials or reference data, for example, at the Institute for Reference Materials and Measurements [7].

Various database technologies are implemented for the construction of the metrological database systems. The choice of which technologies to employ is influenced mostly by the amount of data in the databases and the available resources.

Metrological databases may be used as the tool for monitoring the processes in a distributed metrology system. An example of such a Web-based database system is realised in the Slovenian National Metrology Institute [21, 26]. The implementation is adjusted to the specifics of the organisation of Slovenian metrology. MIRS is responsible for the whole scope of national metrological activities, including maintaining the system of national and reference standards for physical quantities and chemical measurements, the system of legal metrology (type approvals, verifications, precious metals), metrological surveillance of legally controlled instruments and other issues, including the Slovenian business excellence prize. Monitoring of all processes in such a system requires acquiring, manipulating, and processing a large amount of data. The MIRS implementation is suitable for monitoring activities of the distributed metrology system in a small country, covering various aspects from the point of view of the responsible organisation, players, public, and international partners. Besides facilitating organisational issues it gives transparent information for all members of the metrology community, from high-end metrological laboratories to the users of measuring instruments.

The implementation is based on the Linux (<http://www.linux.org>) operating system with an Apache (<http://www.apache.org>) server and MySQL

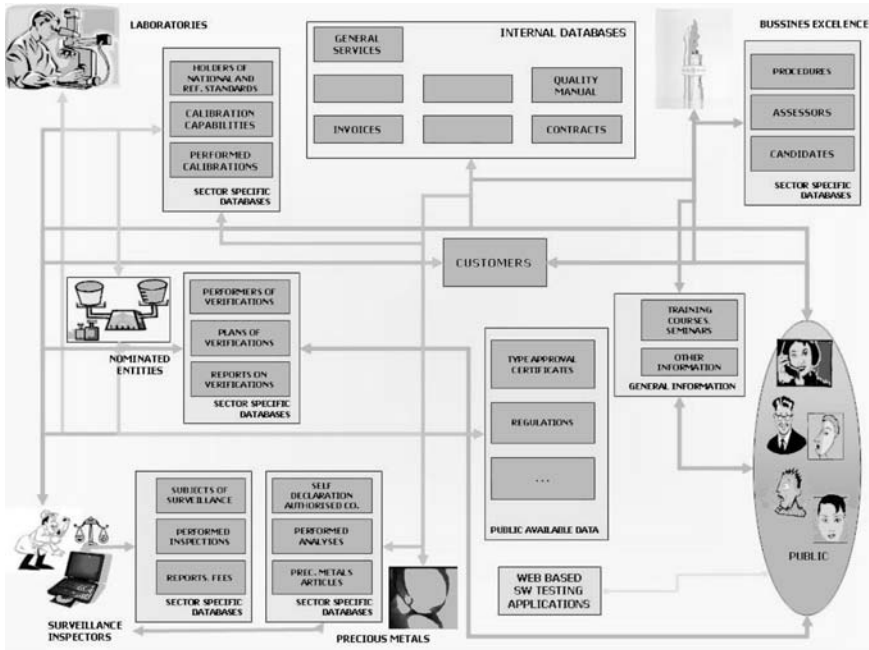


Fig. 7. System of databases for monitoring functioning of national metrology system, an example.

(<http://www.mysql.com>). Access and applications are realized with PHP (<http://www.php.net>) and sometimes JavaScript (<http://javascript.internet.com>) scripts. With selected open source technology it is possible to minimise initial expenses for building such a system. If necessary, the platform may be changed afterwards, when the basic concepts have been clarified.

5 Final Remarks

Benefits arising from the spread of the Internet for the metrology community are manifold. They may be summarised as follows.

- Implementation of new or improvement of the existing functionality of metrological systems (e.g., distributed measuring systems).
- Improvement of metrological services (e.g., remote calibrations, remote procedure validation tools). Benefits are both fundamental improvement in some areas (achieving lower uncertainties as in the case of time and frequency) as well as improved quality (functionality, reliability, and faster performance) of metrological services.
- Increased availability of metrology-related information.

In order to take full advantage of new possibilities, they have to be implemented carefully (taking into account IT security issues, in the first place) with a full understanding of the background of applied technologies.

The Web-enabled services presented here are only examples of metrological Internet services now available. In no case it is the complete list, but provides only a short overview and illustrations of the possibilities open to the users of the metrological services.

All websites referenced in this chapter were active at the time of writing the chapter. There is no guarantee that they are still ‘alive’ or correct, not transferred to another Web address. Certain commercial entities, equipment, or materials are mentioned in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

References

1. Premus A., Palmin U., Tasić T., and Bojkovski J. Validation of Web application for testing of temperature software. In *VII Conference on Advanced Mathematical and Computational Tools in Metrology, AMCTM 2005, Lisbon*, 2005.
2. ZigBee Alliance. Zigbee alliance - wireless control that simply works. <http://www.zigbee.org/en/index.asp>.
3. Allan D. and Weiss M. Accurate time and frequency transfer during common-view of a GPS satellite. In *Proceedings of 34th Annual Frequency Control Symposium, USAERADCOM, Ft. Monmouth, NJ*, 1980.
4. Richter D., Greif N., Zisky N., and Grottker U. IT in metrology: A view into the future. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin*, 2007.
5. Gareth F. Basic concepts of Internet metrology. In *Workshop on Internet Metrology and Self-Calibration, NMi, Delft, The Netherlands*, 2002.
6. Jaeger F. and Grottker U. Security of image and measurement data produced by traffic enforcement systems. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin*, 2007.
7. Institute for Reference Materials and Measurements. <http://www.irmm.jrc.be/html/homepage.htm>.
8. Parkin G. and Harris P. Ensuring numerical correctness using the internet. *Computer Standards & Interfaces: Special Issue on Validation of Software in Metrology*, 28(3):297–305, 2006.
9. Parkin G. and Barker R. M. *SSfM Good Practice Guide No 19: Internet-enabled Metrology Systems*. National Physical Laboratory, Hampton Road, Teddington, Middlesex, UK, 2006. <http://www.npl.co.uk/SSfM>
10. Matsumoto H., Sasaki K., and Hirai A. Remote calibration of practical lengths by using low-coherence interferometry and optical fibre network. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin*, 2007.

11. European Cooperation in Legal Metrology (WELMEC). <http://www.welmec.org/>.
12. Bureau international des poids et mesures (BIPM). Coordinated universal time (UTC). http://www.bipm.org/en/scientific/tai/time_server.html.
13. Bureau international des poids et mesures (BIPM). NMIJ-BIPM workshop on the impact of it in metrology (May 2005). http://www.bipm.org/en/bipm/nmij-bipm_workshop/.
14. Lombardi M.A., Nelson L.M., Novick A.N., and Zhang V.S. Time and frequency measurements using the global positioning system. In *Measurement Science Conference*, 2001.
15. Zisky N. Security concept for remote reading of measurement data. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin*, 2007.
16. National Research Council Canada (NRC). National Research Council of Canada, Conseil national des recherches du Canada. <http://www.nrc-cnrc.gc.ca/>.
17. US Naval Observatory. US Naval Observatory. <http://www.usno.navy.mil/>.
18. International Organization of Legal Metrology (OIML). <http://www.oiml.org/>.
19. National Institute of Standards and Technology (NIST). <http://www.nist.gov/>.
20. National Institute of Standards and Technology (NIST). NIST internet time service (ITS). <http://tf.nist.gov/service/its.htm>.
21. Metrology Institute of the Republic of Slovenia. <http://www.mirs.gov.si/en/>.
22. Physikalisch-Technische Bundesanstalt (PTB). PTB-BIPM workshop on the impact of information technology in metrology. http://e00051.berlin.ptb.de/portal/page?_pageid=42,1&_dad=portal&_
23. SELMA-Konsortium. Sicherer Elektronischer Messdaten-Austausch. <http://www.selma-project.de/>.
24. Kobata T., Kojima M., and Kajikawa H. Development of remote calibration technology for pressure standards. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin*, 2007.
25. Tasić T. International workshop: Future Aspects of Software and IT in legal metrology (FASIT). *OIML Bulletin*, XLV(1):49–51, 2004.
26. Tasić T., Urleb M., and Grgić G. System of databases for supporting co-ordination of processes under responsibility of MIRS. In *VII Conference on Advanced Mathematical and Computational Tools in Metrology, AMCTM 2005, Lisbon*, 2005.
27. Tasić T. and Grottler U. An overview of guidance documents for software in metrological applications. *Computer Standards & Interfaces: Special Issue on Validation of Software in Metrology*, 28(3):256–269, 2006.
28. Palmin U., Tasić T., and Bojkovski J. Web based application for polynomial thermoelectric calculation and verification of its correction curve. In *BIPM - NMIJ Workshop on the Impact of Information Technology in Metrology, Tsukuba, Japan*, 2005.
29. Sato Y., Yamada T., Suzuki T., Toramatsu C., Yunoki A., Hino Y., Harano H., Matsumoto T., Shimoyama T., Kurosawa T., Kato M., Kudo K.,

Matsumoto M., Kenji S., Wakitani Y., Itoh J., Hatakeyama S., Doi K., Saito Y., Suzuki K., Fukumura T., Miyamoto K., Iwamoto A., Endou K., and Matsubara S. The remote calibration of radioisotope calibrators and apparatuses for radiation protection. In *PTB-BIPM Workshop on the Impact of Information Technology in Metrology, Berlin, 2007*.

A

DVD Content

A copy of a DVD accompanies each copy of this book.

The DVD contains materials supplementing the subject matter illustrated in the textbook, provided by the authors or by additional authors contributing only to the DVD contents. It comprises examples, guidelines, reports, and software.

All material is included with the authors' agreement. The material is copyright-free and in some cases of software, a specific license is included in the relevant package with specified conditions of use.

DISCLAIMER: All authors, editors, and the publisher supply the material contained in the DVD as is. They will not answer questions in regard to it and will not make bug fixes or corrections to the programs or material.

DVD contents

Additional material to the chapters

- **Chapter 2** – Giovanni B Rossi and Francesco Crenna *Demonstrations of the probabilistic approach. Related educational software*
- **Chapter 3** – NIST e-Handbook (see below under heading “NIST Engineering Statistical Handbook”)
- **Chapter 4** – Scott Ferson, Vladik Kreinovich, Janos Hajagos, William Oberkampf and Lev Ginzburg *Experimental Uncertainty Estimation and Statistics for Data Having Interval Uncertainty* SANDIA REPORT SAND2007-0939 (May 2007)
- **Chapter 11** – Alan Steele and Rob Douglas *Simplifications from Simulations: Monte Carlo Methods* in Excel
- **Chapter 12** – Norbert Greif and Dieter Richter *Case Studies of Validation Procedures Performed at PTB.*
- **Chapter 13** – Octavian Postolache *Examples related to Chapter 13*

Reports and Guidelines

Measurement Software

- **LSMFE** – Franco Pavese and Daniela Ichim *Software code for implementing the Least Squares Method with Fixed Effect (LSMFE) – weighted and unweighted*, with Manual. Platforms: MS Excel for Windows and for Macintosh. Freeware (GNU license type)
- **SAODR** – Franco Pavese and Daniela Ichim *Software code for implementing the Sequence Analysis for Outlier Data Rejection (SAODR) routine*, with Manual. Platforms: MatLab (with runtime for LabView) and FORTRAN77. Freeware (GNU license type)
- **DATAPLOT** – NIST, **DATAPLOT** with Manual (NIST Handbook 148). Dataplot is a statistical graphical data analysis program that has been integrated with the NIST/SEMATECH e-Handbook of Statistical Methods. You can download the most current public version of the Dataplot software from the NIST Dataplot website:
<http://www.itl.nist.gov/div898/software/dataplot/>
 In the DVD: codes for Windows and Macintosh platforms.

NIST Engineering Statistical Handbook

NIST/SEMATECH *e-Handbook of Statistical Methods* (NIST Handbook 151)

Monte Carlo Simulation

NRC – Alan G. Steele and Robert Douglas *Monte Carlo simulation for expressing uncertainty of measurements: tutorial*

Statistical Software

- **Gum Benchmark** – Rüdiger Kessel: *GUM Benchmark* (v2.3) Educational Version. Freeware.
- **Calculation of expanded uncertainty** – Andrzej Hetman and Marian Jerzy Korczynski *Calculation of expanded uncertainty* Runtime version on LabView. Freeware
- **Bootstrap Calculator** – Patrizia Ciarlini and Giuseppe Regoliosi *Bootstrap Calculator* with Manual, Java. Freeware.
- **Semstat** – Semstat is a statistical software program developed by Jack Prins prior to his retirement from Sematech. Semstat is used for a number of examples in the e-Handbook. Semstat is only available for PCs. SEMSTAT may be downloaded from the International Sematech Web page: <http://www.ismi.sematech.org/stats/semstat.htm>.

Index

- E_n , 312–314, 326, 363, 365
 H_0 , 38, 41, 364
 L_1 -norm, 169
 L_2 -norm, 147
 χ^2 , 21, 81, 306
 δ -Dirac, 47
 t -Student, 40, 68, 78, 84, 95
 z -score, 21
<http://www.amctm.org>, VI
<http://www.imeko-tc21.org>, VI
<http://www.imeko.org>, VI
http://adn.tm.agilent.com/index.cgi?CONTENT_ID=771, 450
http://cholm.home.cern.ch/cholm/misc/gslmm/random_2generator-test_8cc-example.html, 368
http://e00051.berlin.ptb.de/portal/page?_pageid=42,1&_dad=portal&_, 470
<http://ieeexplore.ieee.org/iel4/6051/16227/00751254.pdf?arnumber=751254>, 449
<http://ieeexplore.ieee.org/iel5/7542/20537/00949422.pdf>, 451
<http://inms-ienm.nrc-cnrc.gc.ca/qde/montecarlo/>, 369
<http://iridia.ulb.ac.be/>, 253
<http://java.sun.com/applets/>, 450
<http://java.sun.com/docs/books/tutorial/java/concepts/>, 450
<http://java.sun.com/>, 450
<http://javascript.internet.com>, 468
<http://kcdb.bipm.org>, 466
<http://support.microsoft.com/kb/833855>, 368
<http://tf.nist.gov/service/its.htm>, 470
<http://volt.ni.com/niwc/products/iee1394.jsp>, 449
<http://www.R-project.org>, 115
<http://www.aonaware.com/>, 253
<http://www.apache.org>, 467
http://www.bipm.org/en/bipm/nmij-bipm_workshop/, 470
http://www.bipm.org/en/scientific/tai/time_server.html, 470
http://www.bipm.org/utills/en/pdf/news_jcgm-wg1.pdf, 114
<http://www.bipm.org>, 25
<http://www.cg.tuwien.ac.at/~theussl/DA/thesis.html>, 251
<http://www.compactflash.org/>, 449
<http://www.cs.hku.hk/~diehard>, 369
<http://www.cs.utep.edu/interval-comp>, 144
<http://www.csie.ntu.edu.tw/cjlin/papers/nusvmtutorial.pdf>, 253
<http://www.data-fusion.org/>, 250
<http://www.eng.tau.ac.il/~bengal/BN.pdf>, 254
<http://www.evitherm.org/>, 467
<http://www.geocities.com/adotsaha/NN/SOMinExcel.html>, 235
<http://www.imep.ws/>, 273

- <http://www.interscience.wiley.com>, 253
<http://www.irmm.jrc.be/html/homepage.htm>, 469
<http://www.iro.umontreal.ca/~simardr/testu01/tu01.html>, 369
<http://www.itep.ws/>, 327
<http://www.itl.nist.gov/div898/pubs/author/guthrie/guthrie-2002-01.pdf>, 114
<http://www.iupac.org/reports/provisional/abstract07/fajgelj\290208.html>, 273
<http://www.ivifoundation.org/About%20IVI/Architecture.htm>, 451
<http://www.keithley.com/products/dataacqmodules>, 449
<http://www.keithley.com/products/software/testpointdevpackage/?mn=testpoint>, 450
<http://www.linux.org>, 467
<http://www.lxistandard.org/home>, 449
<http://www.mathworks.com/products/matlab/?BB=1>, 450
<http://www.mathworks.com>, 176, 240
<http://www.mirs.gov.si/en/>, 470
<http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml>, 114
<http://www.mysql.com>, 468
<http://www.nag.co.uk/>, 176
<http://www.ni.com/compactrio/whatis.htm>, 450
<http://www.ni.com/dataacquisition/>, 449
<http://www.ni.com/labview/>, 450
<http://www.nist.gov/stat.handbook>, 28
<http://www.nist.gov/>, 470
<http://www.npl.co.uk/SSfM>, 469
<http://www.npl.co.uk/ssfm>, 24, 167
<http://www.npl.co.uk>, 24
<http://www.nrc-cnrc.gc.ca/>, 470
<http://www.oiml.org/>, 470
<http://www.patentstorm.us/patents/6108609.html>, 251
<http://www.php.net>, 468
<http://www.phy.duke.edu/~rgb/General/dieharder.php>, 368
<http://www.pxisa.org/overview.htm>, 450
<http://www.ramas.com/intstats.pdf>, 145
<http://www.sei.cmu.edu/cmml/cmml.html>, 409
<http://www.selma-project.de/>, 470
http://www.sm.luth.se/grip/Research/Publications/Wavelets/SoftwareComparison/Gri04S_Wavelet_Software_Comparison_Table.pdf, 220
http://www.stat.colostate.edu/research/2002_10.pdf, 114
http://www.stat.colostate.edu/research/2006_3.pdf, 114
<http://www.stat.ucla.edu/history/essay.pdf>, 114
<http://www.sun.com/download/products.xml?id=43208d3d>, 450
<http://www.usno.navy.mil/>, 470
<http://www.virm.net/>, 467
<http://www.wavelet.org>, 220
<http://www.welmec.org/>, 411, 470
<http://www.zigbee.org/en/index.asp>, 469
<http://zone.ni.com/devzone/cda/tut/p/id/3358>, 450
<http://zone.ni.com/devzone/cda/tut/p/id/3889>, 450
<http://zone.ni.com/devzone/cda/tut/p/id/4381>, 449
<http://zone.ni.com/devzone/cda/tut/p/id/5423>, 449
<https://bi.offis.de/viem/tiki-index.php>, 467
A2LA Guide, 7, 8, 11
absolute value minimisation, 169
accuracy, 2–4, 6, 11, 16, 77, 126, 139, 158, 159, 207, 210, 327, 343, 352, 415
bound, 139
computation, 133, 135
fractional, 343, 352, 359
lack of, 117
of approximation, 155

- of method, 359
 - of protocol, 459
 - relative, 126, 352
 - state estimate, 221
 - ADC, 180, 417
 - bandwidth, 185
 - input range, 180
 - multiplexed-single, 417
 - parallel-multiple, 417
 - resolution, 180
 - sampling rate, 180
 - Nyquist, 182
 - aliasing, 182, 213
 - Allan variance, 72
 - AMCTM, VI
 - analog-to-digital converter, 180
 - ANN, 225
 - feed-forward, 227
 - Kohonen, 227
 - multilayer perceptron, 227
 - radial-basis, 227
 - recurrent, 227
 - Elman, 227
 - scalability, 228
 - ANOVA, 42, 304, 305
 - one-factorial, 305
 - random effects
 - between*groups, 305
 - within*group, 305
 - approach
 - χ^2 test, 306
 - A, 17, 21
 - B, 18, 21
 - Bayesian, 13, 33, 39, 44, 46, 73, 75, 87, 104, 111, 157, 275
 - Birge, 307
 - cause-effect modelling, 277
 - engineering, 121, 122
 - fiducial, 13, 73, 75, 93, 107, 331
 - frequentist, 73, 74, 77, 101, 110
 - Gauss, 31
 - Laplace, 35
 - maximum entropy, 121, 275
 - other, 18
 - sequential quadratic program, 164
 - statistical, 74
 - systematic
 - to measurement modelling, 288
 - artificial neural networks, 201, 208, 225
 - training, 201, 226
 - assumptions, 2, 14
 - distributional, 75
 - backpropagation, 201
 - Bayes, 44, 236
 - risk, 249
 - theorem, 87, 105, 111, 325
 - Bayesian inference, 236
 - bias, 5, 7, 11, 14, 15, 17, 226, 300, 303, 304
 - between*groups, 304
 - between*laboratories, 320
 - additive, 319
 - consistent, 7
 - correction, 308
 - lack of, 150
 - significant, 304
 - term not included, 18
 - unbiased, 151
 - uncorrected, 9
 - unknown, 5, 302
- Birge factor, 308
- bootstrap, 74, 79, 81, 85–87, 102, 110
 - nested, 83
 - nonparametric, 82
 - parametric, 82
 - second level, 83
- box, 120, 126, 132
 - p-, 142
- Box–Muller method, 345
- BS 7925, 380
- calibration, 15, 48, 62, 107, 166, 268, 460
 - error, 37
- calibrator, 264, 268
- Campbell, 56
- CEM, 383
- central limit theorem, 35, 78, 122
- cepstrum, 214
- Chirplet Transform, 197, 199, 200
- chirplet: mother, 199
- Cholesky factorisation, 151, 156, 158, 162, 315
- CITAC Guide, 310
- classification, 227, 232
- clustering, 229, 233
 - algorithms, 229

- analysis, 314
- density-based, 229
- grid-based, 229
- hierarchical, 231
- model-based, 229
- partitioning, 230
- combination, 5
 - of measurement results, 265, 300
- combination arithmetic, 5
- comparability (metrological), 260, 262
- comparison, 3, 15, 256, 259
 - interlaboratory*, 3, 15–17, 262
 - intra*laboratory, 3
 - key, 4, 10, 15, 19, 361
 - calibration measurement capabilities (CMC), 21
 - key comparison reference value, 21, 362
 - significant unresolved deviation, 21
 - of measurement results
 - many, 265, 304
 - two, 256, 300
 - outcomes, 19, 22
 - protocol, 20
- compatibility, 21, 22, 258, 270, 302
 - lack of, 302
- complex coefficient, 179
- computer hardware
 - CompactFlash, 421
 - Ethernet
 - LAN, 423
 - FireWire (IEEE 1394), 420
 - GPIB, 423
 - LXI, 423
 - MIOB characteristics, 418
 - USB, 419
- computer interfacing
 - CompactRIO, 425
 - FieldPoint, 425
 - PXI, 425
- computer platforms, 416
 - bus compatibility, 416
 - IBM compatible, 416
 - Macintosh, 416
 - PXI, 416
 - UNIX, 416
 - VME, 416
 - VXI, 416
- confidence
 - interval, 19, 22, 75, 77, 78, 81, 84, 318
 - level, 22, 126, 364
- conformity, 2, 455
- consensus, 14, 270
- consistency, 1, 21, 22, 61, 307
 - degree of, 1, 301, 305, 306
 - lack of, 302, 306, 307
 - test of, 22
- correction, 11, 43, 46, 308
- correlation
 - type algorithm, 211
 - effect, 159, 301, 303, 314
- cost-benefit analysis, 243
 - BRAND, 244
 - decision-tree, 246
 - grid, 246
 - pair-compared, 245
 - Pareto, 245
 - SWOT, 244
- Cov, 81, 142
- covariance, 43, 80, 89, 142, 221, 222, 307, 313, 314, 336–338
- coverage
 - lack of, 306
 - level of, 52, 302, 304, 362
- CRM, 264
- data
 - ambiguity, 206
 - coordinate transformation, 208
 - experimental, 22, 148, 336
 - fusion, 205
 - architecture, 207
 - decision stage, 208
 - implementation, 239
 - raw, 207
 - space and time coverage, 206
 - terminology, 206
 - vector stage, 208
 - heteroscedastic, 318
 - homoscedastic, 3, 33, 35
 - mining, 227
 - modelling, 3, 276
 - multivariate, 78, 152, 157, 223, 276, 312, 336
 - nonnumerical, 2
 - normalised, 323
 - numerical, 2
 - of different accuracy, 207

- pooled, 304
- pooling, 265, 301
- processing, 119
- reference set, 465, 466
- robustness, 206
- security, 395
- univariate, 276, 300
- decision, 2, 6, 137, 235, 268
 - making, 207, 242
 - tree, 246
- degree of equivalence, 4, 21, 22
- degrees of freedom, 40, 78, 79, 81, 95, 101, 102, 108, 169, 332, 364
- deviation
 - additive, 289
 - multiplicative, 288
- DFT, 180, 185
 - algorithm, 181
 - inverse, 212, 214
- difference, 3, 256
 - extended coverage, 302
 - minimum coverage, 302
 - pair, 306
- dimensionality reduction, 223
- DIN 1319, 6, 8
- DIN 66270, 382
- Dirac-delta operator, 47
- Dirichlet condition, 179
- discriminant analysis, 314, 323
- dispersion
 - of measured values, 2, 258, 320
- distance
 - Chebyshev, 210
 - Minkowski, 210
- distribution
 - χ^2 , 96
 - t*-Student, 40, 81, 102, 108, 169, 276, 332, 333, 346
 - F*-Fisher, 38
 - arcsine, 108, 334
 - binomial, 324
 - bootstrap, 82
 - bounded, 124
 - Cauchy, 124, 125
 - conditional, 50, 66, 84, 87
 - cumulative function, 78, 94, 138, 349
 - double exponential, 169
 - empirical, 81, 120
 - fiducial, 75, 93
 - function
 - cumulative, 138
 - general, 334
 - skewed, 336
 - tabular form, 335
 - generating a specific one, 345
 - heavy tail, 171
 - joint, 47, 88
 - marginal, 47
 - multinormal, 157
 - multivariate, 336
 - normal, 15, 19, 33, 35, 52, 78, 81, 94, 122, 259, 266, 270, 309, 321, 332, 345, 354
 - multivariate, 153
 - objective prior, 39
 - parameter-free, 78
 - Pearson, 73
 - Poisson, 324
 - posterior, 87, 105, 152, 156
 - prior, 90, 105
 - rectangular, 275, 302, 334
 - simulated, 97
 - truncated, 99
 - U-shaped, 334, 348
 - uniform, 39, 45, 84, 85, 93, 98, 102, 121, 275, 334, 347
 - limitations, 122
- DoE
 - bilateral, 21
 - unilateral, 21
- domain
 - frequency, 293, 295
 - Laplace, 293
- DVD contents, 52, 65, 145, 367, 371, 408
- dynamic system, 283
- EA Guide, 7, 8, 11
- edge, 238
- Eisenberg principle, 214
- ELMEC-Guide 7.2, 391
- enclosure, 130, 133
 - for the values range, 127
- equation
 - algebraic, 280
 - linear system of, 149
 - nonlinear, 306
 - normal, 149
 - observation, 284

- ordinary differential, 284
- structural, 95, 112
- equivalence
 - degree of, 270
 - metrological, 259
- error
 - approach, 7, 11
 - bound, 119
 - guaranteed, 127
 - upper, 119
 - calibration, 37
 - component
 - lab specific, 12, 14, 18
 - shared, 12, 18
 - dynamic, 293
 - quick-changing response, 294
 - mean squared, 81
 - dynamic, 294
 - normalised E_n , 21, 312, 363
 - of linearity, 292
 - propagation, 43, 79, 100, 281, 301
 - random, 5, 11, 12, 14, 33
 - systematic, 5, 8, 10, 12, 13, 33, 139
 - as random variable, 6, 15
 - constant, 5, 19, 46, 139
 - correction of, 11
 - known, 6
 - neglected, 33, 37
 - nonprobabilistic, 19
 - randomising, 10, 15
 - unknown, 6, 19
 - upper bound, 139
 - theories, 11
 - classical, 33, 74
 - total, 12, 19
 - unforeseen, 12
- estimate, 5
 - 'best', 278
 - robust, 23, 79, 93, 143, 148, 169, 172, 173, 304, 308, 309
 - subjective, 79
 - unbiased, 150, 151
 - worst case, 5
- Euclidean
 - metric, 210
 - vector, 149
- EUROLAB, 310
- evidence theory, 32
 - examples, 37, 40, 49, 57, 63, 65, 76, 84–86, 89, 92, 93, 96, 97, 99–101, 104, 107, 109, 119, 130, 132, 134, 136, 139, 159, 166, 167, 184, 197, 201, 219, 222, 225, 232, 233, 235, 240, 241, 259, 277, 279, 285, 291, 294, 295, 310, 314, 338, 353, 362, 366, 456, 457, 465
 - expectation, 5, 6, 11, 14, 17, 24, 40, 43, 66, 104, 137, 139, 150, 276, 278, 306
 - client, 12
 - constant, 289
 - differences, 16
 - nonzero, 5, 13, 17, 22
 - of utility, 137
 - zero, 5, 8, 9, 11, 13–15, 18, 19, 40, 46, 305
 - expected utility, 137
 - expected value, 89
 - experiment design, 37
 - experimental unit, 9, 13
 - expert judgment, 10, 14, 19, 75, 87, 143
 - feature extraction, 224
 - FFT, 181–184
 - inverse, 191, 213
 - figure of merit, 210
 - filter
 - bandpass, 192
 - highpass, 197, 218
 - Kalman, 220
 - algorithm, 221
 - extended, 222
 - gain, 222
 - Kalman–Bucy, 222
 - Fisher, 37, 39, 75
 - footpoint
 - parameters, 164
 - problem, 164
 - Fourier
 - analysis, 179
 - Discrete Transform (DFT), 180, 212
 - Fast Transform, 181–183
 - Fast Transform (FFT), 181, 211
 - series, 179
 - Short Term Transform, 214
 - Short Time Transform, 199
 - Transform, 287

- frequency
 - angular, 179, 212
 - domain, 177, 182, 187, 200, 211
 - fundamental, 179, 197
 - harmonic, 179
 - Nyquist, 182, 213
 - sampling, 181, 213
- function
 - activation, 226
 - hard-limit, 226
 - linear, 226
 - sigmoidal, 226
 - autocorrelation, 286
 - basis (LS), 148
 - cost, 227
 - likelihood, 40
 - linear, 81
 - membership, 238
 - neighborhood, 234
 - range, 120
 - similarity, 210
 - transfer, 279, 280, 285
- functional
 - relationships, 316
- functional relationship
 - 2D averaging, 318
 - grid projection, 317
 - kernel estimation, 317
- fuzzy, 5, 32, 320, 324
 - inference, 238
 - Sugeno–Takagi type, 238
 - Zadeh–Mamdani type, 238
 - logic, 238
 - logic vs. Boolean logic, 238
- gate
 - elliptical, 211
 - rectangular, 211
- gating, 211
- Gauss problem, 31
- Gauss–Markov problem, 151, 153
 - nonlinear (NGM), 156, 161
 - structured generalised, 168
- Gauss–Newton algorithm, 154
- Gaussian probability distribution, 332
- geometrical interpretation, 149
- GPS, 461
- gradient, 154
- GUM, 6, 8, 10, 11, 13, 17, 19, 32, 42, 46, 71, 79, 81, 100, 104, 258, 264, 271, 272, 278, 281, 288, 301, 304, 310, 314, 321, 332, 347, 362
 - Supplement 1, 73, 329, 333
- harmonic total distortion (THD), 184
- Helmholtz, 55
- hierarchy, 22, 264
 - higher, 14
 - lack of, 8
- histogram, 109, 139, 336, 343, 350, 354–356, 358–360, 362, 367
 - ‘reflection’, 336
 - simulated samples, 330
- I/O hardware, 416
- identification
 - in chemistry, 322, 323
 - qualitative, 322
 - quantitative, 322
- IEC 61508, 380
- IEC 62056-46, 454
- IEC 62056-61, 454
- IEEE 1008, 380
- IEEE 1284, 431
- IEEE 1394, 420
- IEEE 1588, 449
- IEEE 488, 423
- IEEE 802.15.4, 454
- IEEE 829, 380
- IMEKO TC21, VI
- INC-1, 8, 42
- inference, 41, 45, 236
 - Bayesian, 54
 - deductive, 46, 53
 - inductive, 45, 53
 - probabilistic, 53
 - statistical, 73
- influence factors, 4, 8–10, 53, 148, 289
 - variability, 3, 6, 10, 11, 14, 15
- Internet
 - supported
 - metrological service, 459
 - supported
 - calibrations, 460
 - entirely remote calibrations, 461
 - with local standards, 463
 - with signal transmission, 463

- with travelling standards, 464
- Time Service, 459
- interoperability, 12
- interval, 5, 84
 - website, 128
 - addition, 135
 - affined Taylor techniques, 135
 - alternative definition, 130
 - arithmetic, 128
 - addition, 129
 - algorithm, 128
 - division, 128
 - multiplication, 129
 - operation function range, 129
 - reciprocal, 130
 - subtraction, 129, 135
 - bisection, 132
 - bound
 - guarantee, 127
 - on moments, 142
 - bounded, 19
 - centered form, 131
 - closed, 128
 - computation, 120, 127
 - excess width, 131
 - history, 127
 - refined, 130
 - straightforward, 130
 - confidence, 75, 77, 78, 321
 - bootstrap, 82
 - credible, 93
 - Eisenhart, 84, 86
 - error confining, 5, 19
 - fiducial, 97
 - general Taylor techniques, 134
 - half-width, 77, 80, 131
 - midpoint, 126, 131, 304
 - monotonicity, 133
 - multiplication, 135
 - of expectation, 140
 - general case, 141
 - knowing moments, 140
 - of integration, 214
 - open problems, 142
 - probabilistic partial information, 136
 - range, 120
 - uncertainty, 120
- inverse problem, 44, 117, 148, 176, 208, 241, 276
- inverse transform, 211
- ISO 12119, 381
- ISO 12207, 379
- ISO 14471, 379
- ISO 14598, 381
- ISO 14764, 379
- ISO 15408, 382, 395
- ISO 15504, 380
- ISO 15910, 379
- ISO 17025, 72, 340, 383, 400
- ISO 17050, 383
- ISO 18019, 379
- ISO 21748, 27
- ISO 21749, 5
- ISO 25000 Series, 380
- ISO 25051, 381
- ISO 2859, 324
- ISO 3534, 3, 5, 7, 26
- ISO 5725, 3–5, 7, 17, 310
- ISO 6592, 382
- ISO 9000, 379
- ISO 9001, 379
- ISO 9126, 390
- ISO 9127, 382
- ISO 9241, 393
- ISO TR15846, 379
- ISO TR9294, 379
- ITSEC, 382
- ITSEM, 383
- Jacobian matrix, 154
- KCRV, 16, 18, 21–23, 362
- Kinderman–Monahan–Ramage method, 346
- knowledge, 2
 - between* laboratories, 14, 17
 - imperfect, 2, 282, 288
 - prior, 16, 33, 87, 153
 - qualitative, 319
 - quantitative, 256, 300
 - subjective, 111
 - within* laboratory, 14, 15
- Lagrange multiplier, 121, 164
- learning
 - paradigm, 226
 - supervised, 232

- unsupervised, 229
- least squares method, 36, 74, 147
 - generalised, 318
 - generalised distance regression (GDR), 161
 - generalised Gauss–Markov regression (GGM), 167
 - linear, 148, 158
 - nonlinear, 154, 155, 160
 - orthogonal distance (ODR), 167
 - parameters estimation, 147
 - robust (RLS), 169
 - algorithms, 173
 - empirical implementation, 171
 - Huber M-estimator, 172
 - one-sided, 171
 - variable separation, 163, 166
- likelihood function, 40, 74, 87, 169
 - maximum, 125, 152, 156
- linear transmission element, 289
- linearisation, 123, 155, 281
 - no need for, 331
- Lipschitz exponent, 193
- loss, 248
- lti system, 280, 283, 284
- MANOVA, 314
- Markov formula, 276
- matrix, 149
 - block-angular, 162
 - covariance, 224
 - diagonal, 158
 - eigenvalues, 224
 - full rank, 149, 158, 164
 - Hessian, 154
 - identity, 162
 - invertible, 149
 - Jacobian, 154, 155, 162
 - lower triangular, 151
 - nondiagonal, 315
 - of uncertainties, 151
 - approximate, 155
 - nondiagonal, 167
 - structure, 158
 - positive definite, 151, 161
 - positive semi-definite, 315
 - structured, 162
 - transposed, 225
 - uncertainty, 159
 - approximate, 156
 - upper triangular, 149, 152, 163
- MCMC, 91, 111
- mean
 - common, 301
 - hypothesis test, 364
 - grand, 304, 305
 - weighted, 301, 303, 307, 353, 360, 365
 - departure from, 361
 - from Monte Carlo simulation, 360
 - zero, 46
- mean value, 5, 16, 122, 268, 301, 304
 - arithmetic, 5, 34
 - from Monte Carlo simulation, 359
 - posterior, 89, 91–93
- measurand, 5, 7, 12, 13, 17, 20, 43, 48, 74, 100, 110–112, 256, 276, 278, 301
 - multiple, 314
- measure of belief, 75
- measurement, 12
 - capability, 15
 - distributed system, 455
 - error, 2, 32, 119
 - underestimated, 122
 - formal theory, 62
 - non-repeated, 18
 - pair, 262, 316
 - probabilistic theory, 32, 55, 63
- procedure, 4
 - different, 4, 262
 - same, 3
- process, 2, 11, 32, 54, 65
- repeated, 35
- replicated, 9
- representational theory, 59
- results, 257
 - combining, 259, 265, 300
 - discrepant, 300–302
 - pooling, 265, 301, 303, 305
 - small number, 259
- scale, 56, 57, 59, 64
 - invariance, 57
- uncertainty, 32, 46, 119, 258, 264, 300
 - enflated, 270
 - evaluation, 264, 275, 300
 - model, 278, 302
- Measurement Instrument Directive (EU), 325

- measurements
 - pseudo, 329
 - interlaboratory*, 14
 - intralaboratory*, 13
 - autocorrelated, 72
 - cutting-edge, 120
 - direct, 43
 - in chemistry, 256
 - independent, 301
 - indirect, 43
 - modelling of, 275
 - nonrepeated, 2, 13, 14
 - procedure, 14
 - repeated, 1, 3, 5, 9, 13, 14, 37, 46, 140, 276
 - replicated, 1, 4, 13, 303, 336
 - shop-floor, 120
- measuring system, 62, 279
 - disregarding nonlinearity, 292
 - functional elements, 279
 - linearity, 280, 281
 - linearity time-invariance, 280
 - nonlinearity, 282
 - time variant, 283
- median, 66, 268, 309, 361, 362
 - from Monte Carlo simulation, 360
- metric
 - Euclidean, 210
 - Manhattan, 210
 - of association, 209
- metrology, 2, 13
 - Internet, 453
 - remote access, 457
 - legal, 376, 377, 391, 399, 455
 - databases, 467
 - Web-hosted databases, 466
- MID, 391
- minimisation, 36, 321
- MINPACK program, 161
- model, 2, 13, 37, 276, 317
 - basic categories of, 276
 - graphical type, 277
 - mathematical type, 277
 - verbal type, 277
 - cause–effect approach, 277, 282, 288
 - graph, 289
 - common random effect, 158, 160
 - data, 13, 15, 16, 155
 - data fusion, 207
 - Omnibus, 207
 - Dempster–Shafer, 237
 - deterministic, 49
 - equation, 278
 - fitting, 88, 148
 - piecewise, 317
 - imperfection, 292
 - input–output, 48, 79
 - inverse, 208, 278, 284
 - mixture, 23
 - nonexistent laboratory effect, 13
 - of bias, 318
 - of Kalman filter, 220
 - of randomness, 329
 - parameters, 148, 278, 288, 317
 - polynomial, 282, 317
 - probabilistic, 48
 - measurement error, 33
 - random effect, 161
 - random laboratory-effect, 5, 15
 - statistical, 74
 - structure, 288
 - systematic laboratory-effect, 5, 17
- modelling
 - concept, 288
 - of measurement process
 - for uncertainty evaluation, 277
 - systematic approach, 288
 - procedure steps, 290
 - standard components, 288
 - tasks, 275
- Monte Carlo method, 88, 97, 102, 124, 139, 329, 362
 - arbitrary pdf, 356
 - features, 331
 - finite number of resamples, 352
 - Gaussian pdf, 353
 - hypothesis test
 - common mean, 364
 - common mean (χ^2), 365
 - Markov chain, 91
 - non-Gaussian pdf, 356
 - number of simulated samples, 351
 - scalar quantities, 331
- MRA, 4, 10, 19, 22
- neuron
 - layers, 226
 - hidden, 228

scalar with bias, 225
 NIST e-Handbook, 5, 7
 NMI, 72
 nomenclature, 3, 9
 nonquantitative, 1
 non-repeated, 14, 17, 18, 21, 22
 NORDTEST, 310, 383
 normal equations, 149
 not-repeated, 10, 13
 NP-hard problem, 131, 142
 NPL Best Practice Guide N.1, 385
 null hypothesis (H_0), 38, 364

observations, 1, 48

- combination of, 1
- vector, 32, 149, 158

 OECD, 384
 offset, 159
 OIML Guide, 457
 operating-point vector, 278, 282
 optimisation problem, 121, 159
 ordering, 78
 orthogonal factorisation, 149
 outlier, 22, 148, 169, 171, 259, 265, 304

p-box, 138, 139
 parameter, 74

- of a model, 148, 317
- random variables, 87
- unknown constant, 78

 Parceval theorem, 190
 partial derivative, 80, 89, 124, 149, 154
 pdf, 16, 23, 169, 275, 276, 278, 332

- Gaussian, 276
- graphed, 335
- non-normal, 23

 Pearson

- coefficient, 313
- correlation matrix, 314

 percentile, 34, 78, 82, 102

- bootstrap, 82

 perceptron, 201, 226
 Popper, 54
 population, 22, 23, 268

- same, 268
- samples, 15

 precision, 4

- intermediate, 4, 9

principal components analysis, 223, 314, 323

- algorithm, 224
- factor analysis, 223
- principal co-ordinates, 223

 prior knowledge, 87
 probabilistic relation, 64
 probabilistic theory of measurement, 32
 probability, 119

- density function, 16, 23
- distribution, 13, 34, 65, 75, 137, 148
 - t -Student, 40, 78
 - moments, 16, 24, 88, 137, 139, 334
 - non-symmetric, 330
 - normal, 33, 35, 266, 332
 - objective prior, 39
 - of the measurand, 331
 - of the measurements, 331
 - partial information, 138
 - prior, 33, 44, 45, 87
 - symmetric, 34, 35
 - uniform, 39, 121
 - unknown, 120
- in metrology, 32
- measure of belief, 75
- of a relation, 63
- of causes, 45
- of detection, 325

 process, 2

- discrete-time, 222
- nonstationary, 72
- prediction-updating, 221

 pseudo-random number generator, 329, 339

- compact algorithm, 340
- distribution generation
 - for U-shaped distribution, 349
- dynamic range, 343
- generic, 340
- RANLUX, 341
- testing, 341
- user testing, 342

 PTB-IT-14, 410

QR factorisation, 149, 152, 155
 QUAM, 4, 6, 7
 quantile, 86, 97, 102, 108, 321
 quantitative, 1
 quantity, 1, 7, 55, 74

- input, 275, 276, 278, 289
 - vector, 289
- output, 276, 278
- pivotal, 78, 80
 - fiducial, 108
 - generalised, 94
- unit, 271

- R program, 86, 109
- random effect, 5, 148
- random number generator, 126
- random variable, 5, 9, 11, 13, 17, 22, 37, 43, 51, 65, 74–76, 122, 150, 305, 332, 337, 341, 366
 - Bernoulli, 320
- range enclosure, 127
- Rayleigh quotient, 316
- redundancy of information, 223, 249
- regression, 174, 317
 - generalised, 318
 - generalised distance (GDR), 161, 162, 167
 - generalised Gauss–Markov (GGM), 167
 - nonlinear, 313, 317
 - orthogonal distance (ODR), 162
 - weighted, 318
 - both variables, 318, 322
- remote operation, 423, 428, 444
 - measuring instrument, 458
- repeatability, 2–6, 9, 10, 15, 264, 310, 311, 320
 - conditions, 2, 13, 14, 18
- repeated, 1–3, 5, 7–10, 13–16, 18, 21, 22, 34, 35, 37, 39, 46, 52, 139, 140, 299, 316, 336, 337, 363
- replication, 1, 13
- reproducibility, 2–4, 6, 9, 10, 14, 15, 17, 264, 310, 318, 320, 324
- reproducibility, 320
- requirements, 376, 377
- response, 148
 - characteristic, 49
 - approximation, 208
 - normalisation, 208
 - forward characteristic, 284
 - function, 46
 - impulse, 194
 - inverse characteristic, 284
 - of a measuring system, 283
- restitution, 48
- risk
 - analysis, 325
 - level, 22
 - management, 249
 - contingency plans, 249
 - insurance, 249
 - metrological software, 371
- robust rank-based procedure, 79

- sample
 - bootstrap, 82
 - discrete, 78
 - finite number of, 180, 213
 - size, 79
 - standard deviation, 83
- second-order
 - approximation, 281
- Self-Organising Maps Networks, 233
- SELMA, 457
- sensitivity, 2
 - analysis, 89, 105, 111
 - coefficients, 43, 80, 81, 101, 282
- sensor network, 207
- signal, 279
 - amplitude, 180
 - analogue, 180
 - coherence, 286
 - digital, 180
 - discrete-time, 180
 - multidimensional, 214
 - periodic
 - continuous, 179, 183
 - stationary, 181
 - stochastic, 286
 - time variant, 283, 293
 - transients, 187
 - oscillatory, 197
- signal-to-noise ratio (SNR), 184
- significance
 - level, 302
- significance level, 41, 54
- similarity function, 209
- simulation
 - bootstrap, 81
 - Monte Carlo, 329
- SNR, 186

- software, 219, 222, 225, 228, 229, 232, 233, 235, 239, 240, 371
 - NORDTEST validation method, 386
 - ageing, 373
 - audits, 388, 397
 - cause–effect applicability, 374
 - corruption, 371
 - definition of, 372
 - development, 379
 - for metrology, 376
 - guidelines, 376
 - standards, 383
 - for virtual instruments, 425
 - installation, 405
 - functional tests, 407
 - measurability, 375
 - negative synergy effects, 374
 - peculiarities, 373
 - process
 - assessment, 388
 - product testing, 388
 - quality
 - assurance, 371, 387
 - standards, 377
 - quality management, 379
 - relations to space and time, 374
 - requirements, 380, 383, 389, 391, 397, 400
 - process, 384
 - products, 385
 - safety-related systems, 380
 - security, 371, 395
 - test, 391
 - test execution, 396
 - tested components, 392
 - ergonomic characteristics, 394
 - functionality, 393
 - programming techniques, 394
 - security, 395
 - tolerances, 375
 - usability, 394
 - validation, 371
 - archiving, 403, 406
 - confidentiality, 402
 - data security, 404
 - documentation, 403
 - for metrology, 399, 464
 - functional scopes, 402
 - integrity, 402
 - management, 404
 - method for, 407
 - of mathematical functions, 465
 - of metrological functions, 465
 - overall, 405
 - plausibility checks, 402
 - requirements systematisation, 400
 - towards a GUM-like, 408
- SOM
 - incremental-learning, 234
 - iterative algorithm, 234
- spectral
 - composition, 179
 - leakage, 182, 184
 - noise power, 197
 - power density, 287, 295
 - representation, 179
 - resolution, 181, 182
- standard
 - classes or types, 16, 20, 23
 - local, 15
 - measurement, 3, 56
 - same, 3
 - several, 3
 - single, 13, 14
 - written, 3, 4, 14
- standard deviation, 13, 16, 34, 72, 76, 78, 79, 81, 82, 88, 148, 210, 276, 302, 332, 345, 354, 359
 - of the mean, 305, 306, 308, 309, 311, 352
 - posterior, 89, 91–93
- standard method, 14, 17
- statistic
 - ortodox, 36
 - robust, 308
 - sufficient, 84
 - summary, 15
- steady state behaviour, 282
- Stevens, 57
- subband, 187, 194, 195, 197
 - multiresolution
 - decomposition, 196
 - reconstruction, 196
- subjective evaluation, 79
- subsampling, 182
- SUD, 21
- Support Vector Machines (SVM), 232

- systematic
 - approach
 - to measurement modelling, 288
 - difference, 3
 - effect, 5, 17, 46, 54, 111, 159
 - constant, 46
 - correction of, 9, 17, 43
 - of method, 310
 - recognised, 8, 9
- Taylor series expansion, 82, 89, 103, 106, 123, 131, 281, 292
 - higher terms, 132
- test, 3
 - false positive/negative, 322
 - of hypothesis, 41
 - Monte Carlo method, 363
 - of significance, 38, 54
 - proficiency, 3, 20, 259
 - statistical, 40, 78
 - true positive/negative, 322
- testing, 2, 7, 14
 - nondestructive, 325
 - nonquantitative, 319
 - qualitative, 319
 - inspection, 324
 - involving measurements, 320
 - tools, 320
 - uncertainty estimate, 325
- THD, 185
- time, 4
 - frequency analysis, 177
 - frequency representation, 188
 - as influence factor, 10
 - domain, 182, 187
 - invariance, 280
 - series, 227
 - short period, 4, 6
 - to frequency conversion, 177, 213
- traceability, 2, 12, 22, 260–262, 264, 269
- tracking system, 211
- transformation
 - functional, 125
 - inverse, 287
 - Laplace, 285
- true value, 5, 7, 14, 259, 270, 278
 - of the measurand, 5
 - trueness, 7
 - unknowable, 7
- Type A uncertainty component, 8–10, 43, 72, 73, 76, 79, 81, 258, 299, 313, 329
- Type B uncertainty component, 8–10, 43, 72, 73, 76, 79, 86, 87, 92, 93, 258, 326, 329
- uncertainty, 2
 - approach, 7, 11
 - assessment, 72
 - budget, 16, 258, 264, 301, 333
 - due to dynamic components, 293
 - excess, 302, 306
 - expanded, 52, 77, 80, 304, 362
 - expressions, 332
 - interval, 75, 77, 119
 - inverse evaluation, 158
 - lower bounded, 214
 - matrix, 151, 152
 - propagation, 43, 79, 100, 112, 152, 281, 301
 - standard, 74, 160, 300
 - combined, 301
 - Taylor series expansion, 79
 - Type A component, 8, 43, 72, 79, 81, 104, 337
 - Type B component, 8, 43, 72, 79, 81, 87, 92, 101, 104, 111, 271, 337
- UTC, 462
- utility function, 137
 - nonsmooth, 137
 - smooth, 137
- UTP, 185
- value
 - “best”, 265
 - consensus, 270, 271
 - reference, 7, 18, 259, 265, 266
 - representative, 15
- Var, 36, 51, 151
- variability, 4, 43
 - of influence factors, 11
 - withingroup*, 306
- variable
 - correlated, 223
 - internal state, 284
 - pseudo-random, 337, 338

- random, 5, 9, 13, 43, 74
- variance, 15, 19, 22, 35–38, 46, 47, 88, 93, 302, 314
 - inverse of, 303
 - minimum, 36, 150
- verifiability, 54
 - of ANN, 228
- VIM, 3, 4, 6, 10, 18, 21, 22, 73, 256–259, 264
- Virtual Institutes, 467
- virtual instrument, 413
 - comparison with traditional, 414
 - drawbacks, 416
 - operation, 444
 - virtual laboratory, 444
 - single
 - multifunction I/O (MIOB), 417
 - software, 425
 - C, 426
 - graphical programming, 433
 - IVI, 441
 - Java, 431
 - LabView, 433
 - LabWindows/CVI, 426
 - MATLAB, 427
 - object-oriented programming, 430
 - TestPoint, 437
 - VEE Pro, 436
 - VISA, 439
- virtual instruments
 - multiunit
 - distributed, 422
 - distributed with RTS, 424
 - networking capabilities, 422
- visualisation, 233
- voting, 235
- waveform, 182
 - distorted, 197
 - synthesised, 182
- wavelet, 187, 215
 - discrete, 216
 - father, 217
 - mother, 188, 194, 215, 219
 - complex conjugate, 189, 216
 - Daubechies filter, 196
 - discrete, 195
 - features, 191
 - networks, 201
 - nodes, 201
 - transform, 187, 188, 215
 - banks of filters, 193
 - continuous, 187, 189
 - discrete time, 188, 193
 - Fast, 218
 - implementation, 189
 - inverse, 216
 - modulus maxima properties, 192
 - properties of discrete time, 196
 - variance, 72
- weighting function, 170, 171, 173, 226, 303, 322, 353, 360, 363, 366
- Welch–Satterthwaite formula, 80, 81, 333
- WELMEC Guide, 457
- Wiener–Chinchine theorem, 287
- WinBUGS, 92, 106
- windowing, 183
 - rectangular, 183
 - triangular and other types, 184



© Birkhäuser Boston, a part of Springer Science+Business Media, LLC 2009

This electronic component package is protected by federal copyright law and international treaty. If you wish to return this book and the electronic component package to Springer Science+Business Media, LLC, do not open the disc envelope or remove it from the book. Springer Science+Business Media, LLC, will not accept any returns if the package has been opened and/or separated from the book. The copyright holder retains title to and ownership of the package. U.S. copyright law prohibits you from making any copy of the entire electronic component package for any reason without the written permission of Springer Science+Business Media, LLC, except that you may download and copy the files from the electronic component package for your own research, teaching, and personal communications use. Commercial use without the written consent of Springer Science+Business Media, LLC, is strictly prohibited. Springer Science+Business Media, LLC, or its designee has the right to audit your computer and electronic components usage to determine whether any unauthorized copies of this package have been made.

Springer Science+Business Media, LLC, or the author(s) makes no warranty or representation, either express or implied, with respect to this electronic component package or book, including their quality, merchantability, or fitness for a particular purpose. In no event will Springer Science+Business Media, LLC, or the author(s) be liable for direct, indirect, special, incidental, or consequential damages arising out of the use or inability to use the electronic component package or book, even if Springer Science+Business Media, LLC, or the author(s) has been advised of the possibility of such damages.