

Wiley Series on Parallel and Distributed Computing

Albert Y. Zomaya, Series Editor

The DATA Bonanza

Improving Knowledge Discovery
in Science, Engineering, and Business

EDITED BY

Malcolm Atkinson

Rob Baxter

Peter Brezany

Oscar Corcho

Michelle Galea

Mark Parsons

David Snelling

Jano van Hemert

 **IEEE**

 IEEE
computer
society

WILEY

The DATA Bonanza

**WILEY SERIES ON PARALLEL
AND DISTRIBUTED COMPUTING**

Editor: Albert Y. Zomaya

A complete list of titles in this series appears at the end of this volume.

The DATA Bonanza

*Improving Knowledge Discovery in
Science, Engineering, and Business*

Edited by

Malcolm Atkinson
Rob Baxter
Michelle Galea
Mark Parsons

University of Edinburgh, Edinburgh, UK

Peter Brezany

University of Vienna, Vienna, Austria

Oscar Corcho

Universidad Politécnica de Madrid, Madrid, Spain

Jano van Hemert

Optos PLC, Dunfermline, UK

David Snelling

Fujitsu Laboratories Europe Limited, Hayes, UK

WILEY

 **IEEE**
computer
society

Copyright © 2013 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services please contact our Customer Care Department with the U.S. at 877-762-2974, outside the U.S. at 317-572-3993 or fax 317-572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print, however, may not be available in electronic format.

Library of Congress Cataloging-in-Publication Data:

Atkinson, Malcolm.

The data bonanza : improving knowledge discovery for science, engineering and business / Malcolm Atkinson, Rob Baxter, Michelle Galea, Mark Parsons, Peter Brezany, Oscar Corcho, Jano van Hemert, David Snelling.

pages cm

ISBN 978-1-118-39864-7 (pbk.)

1. Information technology. 2. Information retrieval. 3. Databases. I. Title.

T58.5.A93 2013

001-dc23

2012035310

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

To data-to-knowledge highway engineers, everywhere.

Contents

CONTRIBUTORS	xv
FOREWORD	xvii
PREFACE	xix
THE EDITORS	xxix
PART I STRATEGIES FOR SUCCESS IN THE DIGITAL-DATA REVOLUTION	1
1. The Digital-Data Challenge	5
<i>Malcolm Atkinson and Mark Parsons</i>	
1.1 The Digital Revolution / 5	
1.2 Changing How We Think and Behave / 6	
1.3 Moving Adroitly in this Fast-Changing Field / 8	
1.4 Digital-Data Challenges Exist Everywhere / 8	
1.5 Changing How We Work / 9	
1.6 Divide and Conquer Offers the Solution / 10	
1.7 Engineering Data-to-Knowledge Highways / 12	
References / 13	
2. The Digital-Data Revolution	15
<i>Malcolm Atkinson</i>	
	vii

- 2.1 Data, Information, and Knowledge / 16
- 2.2 Increasing Volumes and Diversity of Data / 18
- 2.3 Changing the Ways We Work with Data / 28
- References / 33

3. The Data-Intensive Survival Guide 37

Malcolm Atkinson

- 3.1 Introduction: Challenges and Strategy / 38
- 3.2 Three Categories of Expert / 39
- 3.3 The Data-Intensive Architecture / 41
- 3.4 An Operational Data-Intensive System / 42
- 3.5 Introducing DISPEL / 44
- 3.6 A Simple DISPEL Example / 45
- 3.7 Supporting Data-Intensive Experts / 47
- 3.8 DISPEL in the Context of Contemporary Systems / 48
- 3.9 Datascope / 51
- 3.10 Ramps for Incremental Engagement / 54
- 3.11 Readers' Guide to the Rest of This Book / 56
- References / 58

4. Data-Intensive Thinking with DISPEL 61

Malcolm Atkinson

- 4.1 Processing Elements / 62
- 4.2 Connections / 64
- 4.3 Data Streams and Structure / 65
- 4.4 Functions / 66
- 4.5 The Three-Level Type System / 72
- 4.6 Registry, Libraries, and Descriptions / 81
- 4.7 Achieving Data-Intensive Performance / 86
- 4.8 Reliability and Control / 108
- 4.9 The Data-to-Knowledge Highway / 116
- References / 121

PART II DATA-INTENSIVE KNOWLEDGE DISCOVERY 123

5. Data-Intensive Analysis 127

Oscar Corcho and Jano van Hemert

- 5.1 Knowledge Discovery in Telco Inc. / 128

5.2	Understanding Customers to Prevent Churn / 130	
5.3	Preventing Churn Across Multiple Companies / 134	
5.4	Understanding Customers by Combining Heterogeneous Public and Private Data / 137	
5.5	Conclusions / 144	
	References / 145	
6.	Problem Solving in Data-Intensive Knowledge Discovery	147
	<i>Oscar Corcho and Jano van Hemert</i>	
6.1	The Conventional Life Cycle of Knowledge Discovery / 148	
6.2	Knowledge Discovery Over Heterogeneous Data Sources / 155	
6.3	Knowledge Discovery from Private and Public, Structured and Nonstructured Data / 158	
6.4	Conclusions / 162	
	References / 162	
7.	Data-Intensive Components and Usage Patterns	165
	<i>Oscar Corcho</i>	
7.1	Data Source Access and Transformation Components / 166	
7.2	Data Integration Components / 172	
7.3	Data Preparation and Processing Components / 173	
7.4	Data-Mining Components / 174	
7.5	Visualization and Knowledge Delivery Components / 176	
	References / 178	
8.	Sharing and Reuse in Knowledge Discovery	181
	<i>Oscar Corcho</i>	
8.1	Strategies for Sharing and Reuse / 182	
8.2	Data Analysis Ontologies for Data Analysis Experts / 185	
8.3	Generic Ontologies for Metadata Generation / 188	
8.4	Domain Ontologies for Domain Experts / 189	
8.5	Conclusions / 190	
	References / 191	
PART III	DATA-INTENSIVE ENGINEERING	193
9.	Platforms for Data-Intensive Analysis	197
	<i>David Snelling</i>	
9.1	The Hourglass Reprise / 198	

9.2	The Motivation for a Platform / 200	
9.3	Realization / 201	
	References / 201	
10.	Definition of the DISPEL Language	203
	<i>Paul Martin and Gagarine Yaikhom</i>	
10.1	A Simple Example / 204	
10.2	Processing Elements / 205	
10.3	Data Streams / 213	
10.4	Type System / 217	
10.5	Registration / 222	
10.6	Packaging / 224	
10.7	Workflow Submission / 225	
10.8	Examples of DISPEL / 227	
10.9	Summary / 235	
	References / 236	
11.	DISPEL Development	237
	<i>Adrian Mouat and David Snelling</i>	
11.1	The Development Landscape / 237	
11.2	Data-Intensive Workbenches / 239	
11.3	Data-Intensive Component Libraries / 247	
11.4	Summary / 248	
	References / 248	
12.	DISPEL Enactment	251
	<i>Chee Sun Liew, Amrey Krause, and David Snelling</i>	
12.1	Overview of DISPEL Enactment / 251	
12.2	DISPEL Language Processing / 253	
12.3	DISPEL Optimization / 255	
12.4	DISPEL Deployment / 266	
12.5	DISPEL Execution and Control / 268	
	References / 273	
PART IV	DATA-INTENSIVE APPLICATION EXPERIENCE	275
13.	The Application Foundations of DISPEL	277
	<i>Rob Baxter</i>	

13.1	Characteristics of Data-Intensive Applications / 277	
13.2	Evaluating Application Performance / 280	
13.3	Reviewing the Data-Intensive Strategy / 283	
14.	Analytical Platform for Customer Relationship Management	287
	<i>Maciej Jarka and Mark Parsons</i>	
14.1	Data Analysis in the Telecoms Business / 288	
14.2	Analytical Customer Relationship Management / 289	
14.3	Scenario 1: Churn Prediction / 291	
14.4	Scenario 2: Cross Selling / 293	
14.5	Exploiting the Models and Rules / 296	
14.6	Summary: Lessons Learned / 299	
	References / 299	
15.	Environmental Risk Management	301
	<i>Ladislav Hluchý, Ondrej Habala, Viet Tran, and Branislav Šimo</i>	
15.1	Environmental Modeling / 302	
15.2	Cascading Simulation Models / 303	
15.3	Environmental Data Sources and Their Management / 305	
15.4	Scenario 1: ORAVA / 309	
15.5	Scenario 2: RADAR / 313	
15.6	Scenario 3: SVP / 318	
15.7	New Technologies for Environmental Data Mining / 321	
15.8	Summary: Lessons Learned / 323	
	References / 325	
16.	Analyzing Gene Expression Imaging Data in Developmental Biology	327
	<i>Liangxiu Han, Jano van Hemert, Ian Overton, Paolo Besana, and Richard Baldock</i>	
16.1	Understanding Biological Function / 328	
16.2	Gene Image Annotation / 330	
16.3	Automated Annotation of Gene Expression Images / 331	
16.4	Exploitation and Future Work / 341	
16.5	Summary / 345	
	References / 346	
17.	Data-Intensive Seismology: Research Horizons	353
	<i>Michelle Galea, Andreas Rietbrock, Alessandro Spinuso, and Luca Trani</i>	

- 17.1 Introduction / 354
- 17.2 Seismic Ambient Noise Processing / 356
- 17.3 Solution Implementation / 358
- 17.4 Evaluation / 369
- 17.5 Further Work / 372
- 17.6 Conclusions / 373
- References / 375

PART V DATA-INTENSIVE BEACONS OF SUCCESS 377

18. Data-Intensive Methods in Astronomy 381

Thomas D. Kitching, Robert G. Mann, Laura E. Valkonen, Mark S. Holliman, Alastair Hume, and Keith T. Nodde

- 18.1 Introduction / 381
- 18.2 The Virtual Observatory / 382
- 18.3 Data-Intensive Photometric Classification of Quasars / 383
- 18.4 Probing the Dark Universe with Weak Gravitational Lensing / 387
- 18.5 Future Research Issues / 392
- 18.6 Conclusions / 392
- References / 393

19. The World at One's Fingertips: Interactive Interpretation of Environmental Data 395

Jon Blower, Keith Haines, and Alastair Gemmell

- 19.1 Introduction / 395
- 19.2 The Current State of the Art / 397
- 19.3 The Technical Landscape / 401
- 19.4 Interactive Visualization / 403
- 19.5 From Visualization to Intercomparison / 406
- 19.6 Future Development: The Environmental Cloud / 409
- 19.7 Conclusions / 411
- References / 412

20. Data-Driven Research in the Humanities—the DARIAH Research Infrastructure 417

Andreas Aschenbrenner, Tobias Blanke, Christiane Fritze, and Wolfgang Pemppe

- 20.1 Introduction / 417
- 20.2 The Tradition of Digital Humanities / 420
- 20.3 Humanities Research Data / 422

20.4	Use Case / 426	
20.5	Conclusion and Future Development / 429	
	References / 430	
21.	Analysis of Large and Complex Engineering and Transport Data	431
	<i>Jim Austin</i>	
21.1	Introduction / 431	
21.2	Applications and Challenges / 432	
21.3	The Methods Used / 434	
21.4	Future Developments / 438	
21.5	Conclusions / 439	
	References / 440	
22.	Estimating Species Distributions—Across Space, Through Time, and with Features of the Environment	441
	<i>Steve Kelling, Daniel Fink, Wesley Hochachka, Ken Rosenberg, Robert Cook, Theodoros Damoulas, Claudio Silva, and William Michener</i>	
22.1	Introduction / 442	
22.2	Data Discovery, Access, and Synthesis / 443	
22.3	Model Development / 448	
22.4	Managing Computational Requirements / 449	
22.5	Exploring and Visualizing Model Results / 450	
22.6	Analysis Results / 452	
22.7	Conclusion / 454	
	References / 456	
PART VI	THE DATA-INTENSIVE FUTURE	459
23.	Data-Intensive Trends	461
	<i>Malcolm Atkinson and Paolo Besana</i>	
23.1	Reprise / 461	
23.2	Data-Intensive Applications / 469	
	References / 476	
24.	Data-Rich Futures	477
	<i>Malcolm Atkinson</i>	
24.1	Future Data Infrastructure / 478	
24.2	Future Data Economy / 485	

24.3 Future Data Society and Professionalism / 489
References / 494

Appendix A: Glossary **499**

Michelle Galea and Malcolm Atkinson

Appendix B: DISPEL Reference Manual **507**

Paul Martin

Appendix C: Component Definitions **531**

Malcolm Atkinson and Chee Sun Liew

INDEX **537**

Contributors

- M. ATKINSON, School of Informatics, University of Edinburgh, Edinburgh, UK
A. ASCHENBRENNER, State and University Library Göttingen, Göttingen, Germany
J. AUSTIN, Department of Computer Science, University of York, York, UK
R. BALDOCK, Medical Research Council, Human Genetics Unit, Edinburgh, UK
R. BAXTER, EPCC, University of Edinburgh, Edinburgh, UK
P. BESANA, School of Informatics, University of Edinburgh, Edinburgh, UK
T. BLANKE, Digital Research Infrastructure in the Arts and Humanities, King's College London, London, UK
J. BLOWER, Reading e-Science Centre, University of Reading, Reading, UK
R. COOK, Environmental Sciences Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
O. CORCHO, Departamento de Inteligencia Artificial, Universidad Politécnica de Madrid, Madrid, Spain
T. DAMOULAS, Department of Computer Science, Cornell University, Ithaca, New York, USA
D. FINK, Cornell Lab of Ornithology, Cornell University, Ithaca, New York, USA
C. FRITZE, State and University Library Göttingen, Göttingen, Germany
M. GALEA, School of Informatics, University of Edinburgh, Edinburgh, UK
A. GEMMELL, Reading e-Science Centre, University of Reading, Reading, UK
O. HABALA, Oddelenie paralelného a distribuovaného spracovania informácií, Ústav informatiky SAV, Bratislava, Slovakia
K. HAINES, Reading e-Science Centre, University of Reading, Reading, UK
L. HAN, School of Computing, Mathematics & Digital Technology, Manchester Metropolitan University, Manchester, UK
J. VAN HEMERT, Optos plc, Dunfermline, UK

- L. HLUCHÝ, Oddelenie paralelného a distribuovaného spracovania informácií, Ústav informatiky SAV, Bratislava, Slovakia
- W. HOCHACHKA, Cornell Lab of Ornithology, Cornell University, Ithaca, New York, USA
- M. HOLLIMAN, Institute of Astronomy, University of Edinburgh, Edinburgh, UK
- A. HUME, EPCC, University of Edinburgh, Edinburgh, UK
- M. JARKA, Comarch SA, Warsaw, Poland
- S. KELLING, Cornell Lab of Ornithology, Cornell University, Ithaca, New York, USA
- T. KITCHING, Institute of Astronomy, University of Edinburgh, Edinburgh, UK and Mullard Space Science Laboratory, University College London, Dorking, UK
- A. KRAUSE, EPCC, University of Edinburgh, Edinburgh, UK
- R. MANN, Institute of Astronomy, University of Edinburgh, Edinburgh, UK
- P. MARTIN, School of Informatics, University of Edinburgh, Edinburgh, UK
- W. MICHENER, DataONE, University of New Mexico, Albuquerque, New Mexico, USA
- A. MOUAT, EPCC, University of Edinburgh, Edinburgh, UK
- K. NODDLE, Institute of Astronomy, University of Edinburgh, Edinburgh, UK
- I. OVERTON, Medical Research Council, Human Genetics Unit, Edinburgh, UK
- M. PARSONS, EPCC, University of Edinburgh, Edinburgh, UK
- W. PEMPE, State and University Library Göttingen, Göttingen, Germany
- A. RIETBROCK, School of Environmental Sciences, University of Liverpool, Liverpool, UK
- K. ROSENBERG, Cornell Lab of Ornithology, Cornell University, Ithaca, New York, USA
- C. ŠILVA, Department of Computer Science, Polytechnic Institute of New York, Brooklyn, New York, USA
- A. SPINUSO, School of Informatics, University of Edinburgh, Edinburgh, UK; Royal Netherlands Meteorological Institute, Information and Observation Services and Technology–R&D, Utrecht, The Netherlands
- D. SNELLING, Research Transformation and Innovation, Fujitsu Laboratories of Europe Limited, Hayes, UK
- C. SUN LIEW, School of Informatics, University of Edinburgh, Edinburgh, UK; Faculty of Computer Science and Information Technology, University of Malaya, Kuala Lumpur, Malaysia
- B. SIMO, Oddelenie paralelného a distribuovaného spracovania informácií, Ústav informatiky SAV, Bratislava, Slovakia
- V. TRAN, Oddelenie paralelného a distribuovaného spracovania informácií, Ústav informatiky SAV, Bratislava, Slovakia
- L. TRANI, School of Informatics, University of Edinburgh, Edinburgh, UK; Royal Netherlands Meteorological Institute, Information and Observation Services and Technology–R&D, Utrecht, The Netherlands
- L. VALKONEN, School of Informatics, University of Edinburgh, Edinburgh, UK
- G. YAIKHOM, School of Informatics, University of Edinburgh, Edinburgh, UK; School of Computer Science and Informatics, Cardiff University, Cardiff, UK

Foreword

This book is a systematic approach to help face the challenge of data-intensive computing and a response to the articulation of that challenge set out in *The Fourth Paradigm*, a collection of essays by scientists and computer scientists that I co-edited. It also recognizes the fact that we face that challenge in the context of a digital revolution that is transforming communities worldwide at Internet speed.

This book proposes a strategy for partitioning the challenge, both in the ways in which we organize and in which we build systems. This partitioning builds on natural foci of interest and the examples show that this approach works well in the context of groups and organizations. The technological strategy reflects the evolving pattern of provision driven by business models and the flourishing diversity of tools and applications that enable human innovation.

We all face the need to separate concerns every time we face a data-intensive problem. This is key to making data-intensive methods routinely available and to their easy application. This leads to the recognition of effective working practices that need supporting with better ‘datascope’ that are easily steered and focused to extract the relevant information from immense and diverse collections of data. The book calls for the introduction of “intellectual on-ramps” that match the new tools to well-understood interfaces, so that practitioners can incrementally master the new data-intensive methods.

This book calls for recognition that this notion of intellectual on-ramps is worthy of study. Data-intensive computing warrants an appropriate engineering discipline that identifies effective ways of building appropriate highways from data to required knowledge. It is a call to arms for a serious attempt at initiating the professionalization of this discipline.

We are very much at the start of the digital revolution. The growth in digital data will not abate, and in certain areas it will probably accelerate. There is much to be gained by exploiting the opportunities this bonanza of data brings but the extraction of insights and knowledge from ‘Big Data’ will also certainly transform our organizations and society. Responding effectively to these changes requires the availability of ready-made tools and reusable processes together with practitioners with the skills to deploy them precisely and safely. This book frames an approach as to how these tools, processes, and skills may be developed. Such a systematic approach is now urgently needed as the opportunities are rapidly outgrowing our capabilities to assemble and run data-intensive campaigns.

This book provides a vocabulary to facilitate data-intensive engineering by introducing key concepts and notations. It presents nine in-depth case studies that show how practitioners have tackled data-intensive challenges in a wide range of disciplines. It also provides an up-to-date analysis of this rapidly changing field and a survey of many of the current research hot-spots that are driving it. For all these reasons, I believe that this book is a welcome addition to the literature on data-intensive computing.

TONY HEY

*Redmond, Washington
July 2012*

Preface

The world is undergoing a digital-data revolution. More and more data are born digital. Almost every business, governmental, and organizational activity is driven by data and produces data. Science, engineering, medicine, design and innovation are powered by data. This prevalence of data in all that we do is changing society, organizations, and individual behavior. To thrive in this new environment requires new strategies, new skills, and new technology. This book is the first to expound the strategies that will make you adept at exploiting the expanding opportunities of this new world.

This book identifies the driving forces that are provoking change and proposes a strategy for building the skills, methods, technologies, and businesses that will be well adapted in the emerging data-wealthy world. This strategy will change the way in which you spend your organization's (country's, company's, institution's, and profession's) resources. You will invest more in exploiting data, even if that means spending less on creating, capturing or archiving it, as today most data are underused or never used, even though they frequently *contain the latent evidence that should be leading to innovation, knowledge and wisdom*. After reading this book, you will expect an understandable path from data, via analysis, evidence and visualization, to influential outputs that change behavior. When this is not happening, your organization is underperforming and is at risk. You will come to expect that all of those with whom you deal should be competent at getting good value from data.

This book will change your skills by developing your ingenuity in discovering, understanding, exploiting, and presenting data. You will acquire a compendium of tools for addressing every stage in the data life cycle. It will change education—everyone needs survival skills for the data-wealthy world. All professionals and

experts need experience and judgement for their part of the path from data to discovery, innovation, and outcomes.

This book will initiate the development of professionals who will engineer tomorrow's data highways. These highways will be designed to meet carefully analyzed and anticipated needs, to interwork with existing data infrastructure, to accelerate the journeys of millions from data to knowledge. These knowledge discovery highways will make it easy for people to get data from wherever they are stored and promptly deliver understandable information to wherever it is needed.

Who Should Read this Book

This book will be valuable to a wide range of information strategists, decision makers, researchers, students, and practitioners, from domains such as computer science (data mining, machine learning, statistics, databases, knowledge-based systems, large-scale computing), e-Science, and to workers in any discipline or industry where large-scale data handling and analysis is important.

The ideas presented are relevant to and draw on: data mining, knowledge discovery in databases, machine learning, artificial intelligence, databases and data management, data warehousing, information systems, distributed computing, grid computing, cloud computing, ubiquitous computing, e-Science (including a wide range of scientific and engineering fields dealing with large data), modeling and simulation.

This Book's Structure

This book consists of 24 chapters grouped into six parts; they are introduced here.

Part I: Strategies for success in the digital-data revolution Part I provides an executive summary of the whole book to convince strategists, politicians, managers, and educators that our future data-intensive society requires new thinking, new behavior, new culture, and new distribution of investment and effort. This part will introduce the major concepts so that readers are equipped to discuss and steer their organization's response to the opportunities and obligations brought by the growing wealth of data. It will help readers understand the changing context brought about by advances in digital devices, digital communication, and ubiquitous computing.

Chapter 1: The digital-data challenge This chapter will help readers to understand the challenges ahead in making good use of the data and introduce ideas that will lead to helpful strategies. A global digital-data revolution is catalyzing change in the ways in which we live, work, relax, govern, and organize. This is a significant change in society, as important as the invention of printing or the industrial revolution, but more challenging because it is happening globally at Internet speed. Becoming agile in adapting to this new world is essential.

Chapter 2: The digital-data revolution This chapter reviews the relationships between data, information, knowledge, and wisdom. It analyses and quantifies the changes in technology and society that are delivering the data bonanza, and then reviews the consequential changes via representative examples in biology, Earth sciences, social sciences, leisure activity, and business. It exposes quantitative details and shows the complexity and diversity of the growing wealth of data, introducing some of its potential benefits and examples of the impediments to successfully realizing those benefits.

Chapter 3: The data-intensive survival guide This chapter presents an overview of all of the elements of the proposed data-intensive strategy. Sufficient detail is presented for readers to understand the principles and practice that we recommend. It should also provide a good preparation for readers who choose to sample later chapters. It introduces three professional viewpoints: domain experts, data-analysis experts, and data-intensive engineers. Success depends on a balanced approach that develops the capacity of all three groups. A data-intensive architecture provides a flexible framework for that balanced approach. This enables the three groups to build and exploit data-intensive processes that incrementally step from data to results. A language is introduced to describe these incremental data processes from all three points of view. The chapter introduces ‘datascoptes’ as the productized data handling environments and ‘intellectual ramps’ as the ‘on ramps’ for the highways from data to knowledge.

Chapter 4: Data-intensive thinking with DISPEL This chapter engages the reader with technical issues and solutions, by working through a sequence of examples, building up from a sketch of a solution to a large-scale data challenge. It uses the DISPEL language extensively, introducing its concepts and constructs. It shows how DISPEL may help designers, data-analysts, and engineers develop solutions to the requirements emerging in any data-intensive application domain. The reader is taken through simple steps initially, this then builds to conceptually complex steps that are necessary to cope with the realities of real data providers, real data, real distributed systems, and long-running processes.

Part II: Data-intensive knowledge discovery Part II focuses on the needs of data-analysis experts. It illustrates the problem-solving strategies appropriate for a data-rich world, without delving into the details of underlying technologies. It should engage and inform data-analysis specialists, such as statisticians, data miners, image analysts, bio-informaticians or chemo-informaticians, and generate ideas pertinent to their application areas.

Chapter 5: Data-intensive analysis This chapter introduces a set of common problems that data-analysis experts often encounter, by means of a set of scenarios of increasing levels of complexity. The scenarios typify knowledge discovery challenges and the presented solutions provide practical methods; a starting point for readers addressing their own data challenges.

Chapter 6: Problem solving in data-intensive knowledge discovery On the basis of the previous scenarios, this chapter provides an overview of effective strategies in knowledge discovery, highlighting common problem-solving methods that apply in conventional contexts, and focusing on the similarities and differences of these methods.

Chapter 7: Data-intensive components and usage patterns This chapter provides a systematic review of the components that are commonly used in knowledge discovery tasks as well as common patterns of component composition. That is, it introduces the processing elements from which knowledge discovery solutions are built and common composition patterns for delivering trustworthy information. It reflects on how these components and patterns are evolving in a data-intensive context.

Chapter 8: Sharing and re-use in knowledge discovery This chapter introduces more advanced knowledge discovery problems, and shows how improved component and pattern descriptions facilitate re-use. This supports the assembly of libraries of high level components well-adapted to classes of knowledge discovery methods or application domains. The descriptions are made more powerful by introducing notations from the semantic Web.

Part III: Data-intensive engineering Part III is targeted at technical experts who will develop complex applications, new components, or data-intensive platforms. The techniques introduced may be applied very widely; for example, to any data-intensive distributed application, such as index generation, image processing, sequence comparison, text analysis, and sensor-stream monitoring. The challenges, methods, and implementation requirements are illustrated by making extensive use of DISPEL.

Chapter 9: Platforms for data-intensive analysis This chapter gives a reprise of data-intensive architectures, examines the business case for investing in them, and introduces the stages of data-intensive workflow enactment.

Chapter 10: Definition of the DISPEL language This chapter describes the novel aspects of the DISPEL language: its constructs, capabilities, and anticipated programming style.

Chapter 11: DISPEL development This chapter describes the tools and libraries that a DISPEL developer might expect to use. The tools include those needed during process definition, those required to organize enactment, and diagnostic aids for developers of applications and platforms.

Chapter 12: DISPEL enactment This chapter describes the four stages of DISPEL enactment. It is targeted at the data-intensive engineers who implement enactment services.

Part IV: Data-intensive application experience This part of the book is about applications that shaped the ideas behind the data-intensive architecture and methods. It provides a wealth of examples drawn from experience, describing in each case the aspects of data-intensive systems tested by the application, the DISPEL-based methods developed to meet the challenge, and the conclusions drawn from the prototype experiments.

Chapter 13: The application foundations of DISPEL The early development of DISPEL was influenced and assisted by research challenges from four very different data-intensive application domains. This chapter reviews these four domains in terms of their particular needs and requirements and how, as a suite, they provide an effective test of all key dimensions of a data-intensive system. It reviews the data-intensive strategy in terms of these applications and finds support for the approach.

Chapter 14: Analytical platform for customer relationship management This chapter demonstrates that the data-intensive methods and technology work well for traditional commercial knowledge discovery applications. Readers are introduced to the application domain through a scene-setting discussion, which assumes no prior knowledge, and are then taken through the process of analyzing customer data to predict behavior or preferences.

Chapter 15: Environmental risk management This chapter presents applications in the context of environmental risk management. The scenarios involve significant data-integration challenges as they take an increasing number of factors into account when managing the outflow from a reservoir to limit the effects downstream.

Chapter 16: Analyzing gene expression imaging data in developmental biology This chapter describes the application of data-intensive methods to the automatic identification and annotation of gene expression patterns in the mouse embryo. It shows how image processing and machine learning can be combined to annotate images and identify networks of gene functions.

Chapter 17: Data-intensive seismology: research horizons Seismology has moved from focusing on events to analyzing continuous streams of data obtained from thousands of seismometers. This is fundamental to understanding the inner structure and processes of the Earth and this chapter investigates the data-intensive architecture necessary to enable the analysis of large-scale distributed seismic waveforms.

Part V: Data-intensive beacons of success This part introduces a group of challenging, sophisticated data-intensive applications, which are starting to shape and promote a new generation of knowledge discovery technology. The chapters show that science, engineering, and society are fertile lands for data-intensive

research. This part is targeted at novel application developers who like to include visionary aspects in their research.

Chapter 18: Data-intensive methods in astronomy Astronomy has been at the forefront of the digital revolution as it pioneered faster and more sensitive digital cameras, and established a new *modus operandi* for sharing and integrating data globally. These are yielding floods of data, and opening up new approaches to exploring the cosmos and testing the physical models that underpin it. This chapter describes two examples that exemplify the data-intensive science now underway.

Chapter 19: Interactive interpretation of environmental data A crucial step in any science is to explore the data available; this will often stimulate new insights and hypotheses. As the volumes of data grow and diverse formats are encountered, the effort of handling data inhibits exploration. This chapter shows how these inhibitory difficulties can be overcome, so that oceanographers and atmospheric scientists can easily select, visualize and explore compositions of their data.

Chapter 20: Data-driven research in the humanities Researchers in the arts and humanities are using digitization to see new aspects of the many artifacts and phenomena they study. Digital resources allow statistical methods and computational matching to be employed, as well as the full panoply of text processing and collaborative annotation. In this chapter, researchers show their plans for a Europe-wide data infrastructure to facilitate this new research.

Chapter 21: Analysis of engineering and transport data Analysis of vibration data from aero-engines, turbines, and locomotives, ‘listening to the engine’, can reveal incipient problems and trigger appropriate remedial responses. To do this safely for large numbers of operational engines is a major data-intensive challenge. This chapter reports on ten years’ progress and its spin-offs into analyzing medical time series.

Chapter 22: Determining the patterns of bird species occurrence In this chapter, the ornithologists describe the challenge of estimating the populations of birds as they migrate and of inferring the factors affecting species numbers. It takes a great deal of sophisticated data analysis to extract and visualize the relevant information and much ingenuity to discover and use the data required from other disciplines.

Part VI: The data-intensive future This part presents a summary of the state of the industry and research, the observed trends and the current ‘hot-spots’ of data-intensive innovation. It provides a framework for reviewing the current activity and anticipated changes it will bring about. It offers a rich set of pointers to the literature and Web sites, built over the 15 months of the data-intensive research theme at the e-Science Institute. This should help readers select and find highly relevant further reading.

Chapter 23: Data-intensive trends This chapter summarizes the learning about data-intensive methods and their potential power. It then analyzes some of the categories of data-intensive challenge and assesses how they will develop.

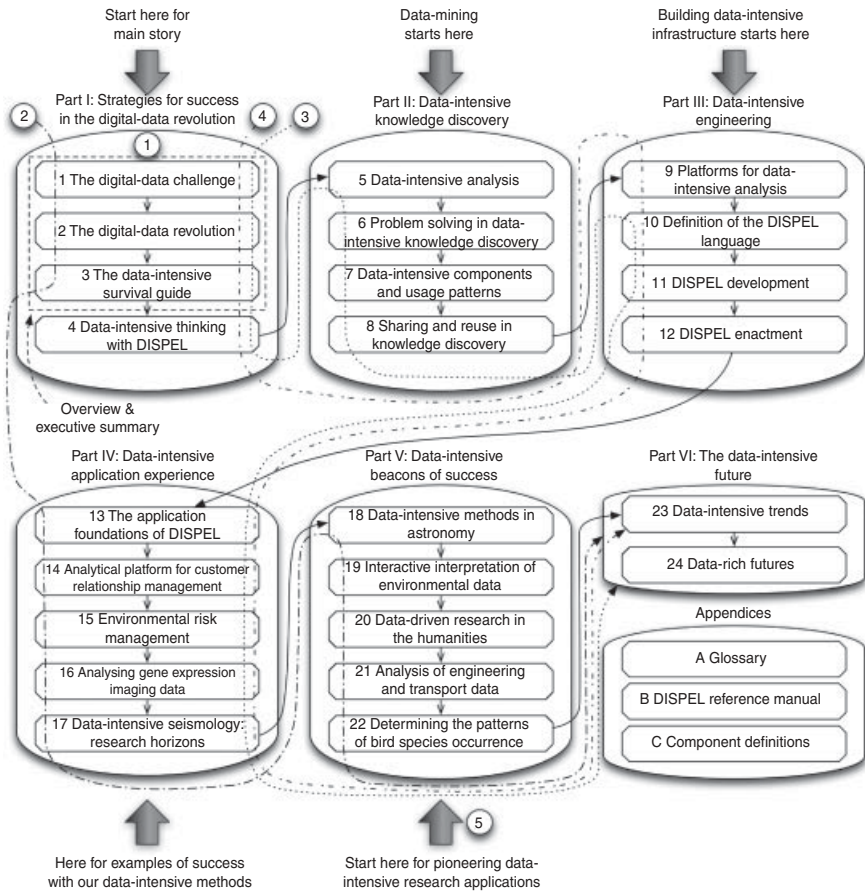
Chapter 24: Data-rich futures This chapter dangerously attempts some ‘crystal gazing’. It looks first at technological factors and current research that should be observed by those who wish to further develop data-intensive strategies. It explores some of the economic factors that will shape a data-rich future and concludes with a view on the social issues that will emerge. We call on those with influence to strive for professional standards in handling data, from their collection to the actions based on their evidence.

How to Read this Book

There are two pieces of information that we wish to offer readers in order to help them get the most benefit from this book: the choice of routes through the book and conventions used in the book. We consider each of these in turn.

1. The primary story line of the book develops in the conventional way by reading the parts and the chapters in order, up to the end of Part IV. It begins with scene setting to establish a conceptual framework. It then addresses the methods and engineering needed to put the ideas into practice before successful applications in a wide range of domains. Many readers will want a selected-reading route that matches their needs; the following map and suggested routes will encourage them to plan their reading.
2. Research strategists, scientists and managers, who may be less interested in technical detail should follow Chapters 1 to 3→Part IV→Part V→Part VI.
3. Data-analysis experts should try Part I→Part II→Chapter 10→Chapter 11→Part IV→Part VI.
4. Those building data-intensive systems should read Part I→Part III→Part IV→Part VI.
5. Domain experts, who are looking for ideas that may pay off with their specialist data could read the applications first, Part V→Part IV→Chapters 1 to 3. We hope there are enough signposts that you can adjust your route easily to follow new interests.

The conventions concern the structure and the representation of certain items. As the map shows, there are six parts; each one begins with a preamble intended to orient readers who start in that part, and to relate that part to the other parts of the book. Appendix A provides a glossary, where we hope you will find useful definitions of terms used frequently in the book. Terms used infrequently may be traced via the index. Each chapter concludes with the references cited in that chapter. References to Web sites are shown as URLs in line with the text if they are short and in footnotes if they are long. The `http://` prefix is omitted.



Readers' map of the book.

There are many programming examples, mostly represented in DISPEL. A consistent highlighting convention has been used for these, which we believe helps legibility. These are also available on the Web site (see following text), so that you can view them in your favorite editor. There are often corresponding diagrams, showing the data-flow graph that the program would generate. These pick up the same color conventions. The language DISPEL is introduced in Chapters 3, 4 and 10; the language definition is in Appendix B. The components that are used in the examples are provided from standard libraries of components; these are described in Appendix C.

Web Site with Additional Material

A Web site at www.dispel.lang.org holds material intended to help readers of this book; these include pages covering the following topics.

- An overview and table of contents of the book.
- A collection of success stories.
- Teaching material including presentations to be used in conjunction with chapters in the book.
- Copies of the program examples from the book, so that they may be used by readers.
- The libraries of components used in the book, with corresponding descriptions in a local registry.
- The DISPEL reference manual.
- Links to other sites including many of those referenced from the book, the associated open-source project and new developments.

This Web site will be updated and contributions from others will be welcome.

Acknowledgments

Many people have contributed to this book through discussions and visits over the past 6 years; we thank them all and limit the explicit list to those who worked directly on the book.

Funders The primary funder of the editorial process was the European Commission's Framework Program 7's support for the ADMIRE project (www.admire-project.eu), grant number 215024. The tour of the USA by Atkinson and De Roure, to study their use of data (bit.ly/c0G2rn), and Atkinson's work on the book was funded by the UK's Engineering and Physical Sciences Research Council (EPSRC) Senior Research Fellowship to undertake the UK e-Science Envoy role, grant EP/D079829/1. The initial workshop on data-intensive research (bit.ly/bQpu5h), where many of the commitments for contributions to the book were made, and the one-year data-intensive research theme (bit.ly/cygimA), were funded by the e-Science Institute, EPSRC grant EP/D056314/1. The work to further develop this book, the technology and the methodology described, is partially funded in the University of Edinburgh by the EPSRC NeSC Research Platform (research.nesc.ac.uk) grant EP/F057695/1, and in Universidad Politècnica de Madrid by the Spanish Ministry of Science and Technology under grant TIN2010-17060. Baxter would also like to acknowledge the direct support of EPCC at the University of Edinburgh, and support from the UK's Software Sustainability Institute under EPSRC grant number EP/H043160/1.

Helpers A big thank you to Jo Newman, who was a continuous support to the editors of the book, impeccably arranging many meetings and tele-communications, as well as thorough proof reading, checking copyright information, and communicating with all of the contributing authors. Kathy Humphry read nearly every chapter and gave good advice on each one of them.

In addition to being contributing authors, we must thank Ivan Janciak, in the University of Vienna, who made life much easier for all of us by setting up convenient macros for building the book; Paul Martin and Gagarine Yaikhom, of the Universities of Edinburgh and Cardiff respectively, for setting up the system for typesetting DISPEL highlighting its structure; and Amrey Krause and Chee Sun Liew of the University of Edinburgh, who set up the system for validating DISPEL text used in the book. Chee Sun Liew also did a great deal of L^AT_EX wrangling to shape the book into its final form. Ivan Janciak, Alexander Wöhrer and Marek Lenart of the University of Vienna reviewed several book chapters and helped improve the quality of figures and the formatting of the text.

We would also like to thank our colleagues Martin Šeleng and Peter Krammer from the Institute of Informatics of the Slovak Academy of Sciences for their excellent work on the data mining scenarios of the environmental risk management application, and our dear friends at the Slovak Hydro-Meteorological Institute and the Slovak Water Enterprise for their invaluable help with designing the pilot scenarios and providing real input data for them. A big thank you must go to the EPCC software engineering team for making large swathes of the prototype data-intensive platform work: Ally Hume, Malcolm Illingworth, Amrey Krause, Adrian Mouat and David Scott, with a special mention for our integration, test and build-meister Radek Ostrowski.

We heartily thank all of the open-source developers on whose work we built; all those who helped and confirm that all the remaining errors are the responsibility of the editors, led by myself.

MALCOLM ATKINSON

Edinburgh, UK
April 2012