

# Exploiting Linked Data and Knowledge Graphs in Large Organizations

Jeff Z. Pan · Guido Vetere  
Jose Manuel Gomez-Perez  
Honghan Wu  
Editors

# Exploiting Linked Data and Knowledge Graphs in Large Organizations

 Springer

*Editors*

Jeff Z. Pan  
University of Aberdeen  
Aberdeen  
UK

Jose Manuel Gomez-Perez  
iSOCO Lab  
Madrid  
Spain

Guido Vetere  
IBM Italia  
Rome  
Italy

Honghan Wu  
University of Aberdeen  
Aberdeen  
UK

ISBN 978-3-319-45652-2

ISBN 978-3-319-45654-6 (eBook)

DOI 10.1007/978-3-319-45654-6

Library of Congress Control Number: 2016949103

© Springer International Publishing Switzerland 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

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

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature

The registered company is Springer International Publishing AG

The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Foreword

When I began my research career as a graduate student at Rensselaer Polytechnic Institute in 1989, the phrase “knowledge graph” was not in use. The use of graphs, however, as a notation for “knowledge representation” (KR) was quite common. CLASSIC, the first real implemented description logic, was just being introduced from Bell Labs, and although it had a linear syntax, the community was still in the habit of drawing graphs that depicted the knowledge that was being represented.

This habit traced its history at least as far as M. Ross Quillian’s work on *Semantic Networks*, and subsequent researchers imagined knowledge to be intrinsic in the design of Artificial Intelligence (AI) systems, universally sketching the role of knowledge in a graphical form. By the late 1980s the community had more or less taken up the call for formalisation proposed by Bill Woods and later his student, Ron Brachman; graph formalisms were perhaps the central focus of AI at the time, and stayed that way for another decade.

Despite this attention and focus, by the time I moved from academia to industrial research at IBM’s Watson Research Centre in 2002, the knowledge representation community had never really solved any problems other than our own. Knowledge representation and reasoning evolved, or perhaps devolved, into a form of mathematics, in which researchers posed difficult-to-solve puzzles that arose more from syntactic properties of various formalisms than consideration of anyone else’s actual use cases. Even though we tended to use the words, “semantic” and “knowledge”, there was nothing particularly semantic about any of it, and indeed the co-opting by the KR community of terms like semantics, ontology, epistemology, etc. to refer to our largely algorithmic work, reliably confused the hell out of people who actually knew what those terms meant.

In my 12-year career at IBM, I found myself shifting with the times as a revolution was happening in AI. Many researchers roundly rejected the assumptions of the KR field, finding the focus on computation rather than data to be problematic. A new generation of data scientists who wanted to instrument and measure everything began to take over. I spent a lot of my time at IBM trying to convince others that the KR technology was useful, and even helping them use it. It was a

losing battle, and like the field in general I began to become enamoured of the influential power of empirical evidence—it made me feel like a scientist. Still, however, my allegiance to the KR vision, that knowledge was intrinsic to the design of AI systems, could not be completely dispelled.

In 2007, a group of 12 researchers at IBM began working on a top secret moonshot project which we code-named “BlueJ”—building a natural language question answering system capable of the speed and accuracy necessary to achieve expert human-level performance on the TV quiz show, *Jeopardy!* It was the most compelling and interesting project I have ever worked on, and it gave me an opportunity to prove that knowledge—human created and curated knowledge—is a valuable tool. At the start of the project, Dave Ferrucci, the team leader, challenged us all to “make bets” on what we thought would work and commit to being measured on how well our bets impacted the ability to find the right answer as well as to *understand if the answer is correct*. I bet on KR, and for the first year, working alone on this particular bet, I failed, much as the KR community had failed more broadly to have any impact on any real problems other people had. But in the following year, Ferrucci agreed to put a few more people on it (partly because of my persuasive arguments, but mostly because he believed in the KR vision, too) and with the diversity of ideas and perspectives that naturally comes from having more people, we started to show impact. After our widely publicised and viewed victory over the two greatest *Jeopardy!* players in history, my team published the results of our experiments that demonstrated more than 10 % of Watson’s winning performance (again, in terms of both finding answers and determining if they were correct) came from represented knowledge.

### **Knowledge is not the destination**

In order to make this contribution to IBM’s Watson, my team and I had to abandon our traditional notion of KR and adopt a new one, that I later came to call, “Knowledge is not the destination”. The abject failure of KR to have any measurable impact on anything up to that point in time was due, I claim, to a subtle shift in that research community, sometime in the 1980s, from knowledge representation and reasoning as an integral part of some larger system, to KR&R as the ultimate engine of AI. This is where we were when I came into the field, and this was tacit in how I approached AI when I was working in Digital Libraries, Web Systems, and my early efforts at IBM in natural language question answering.

The most ambitious KR&R activity before that time was Cyc, which prided itself on being able to conclude, “If you leave a snowman outside in the sun it will melt”. But Cyc could never possibly answer any of the myriad possible questions that might get asked about snowmen melting, because it would need a person to find the relevant Cyc micro-theory, look up the actual names and labels used in the axioms, type them in the correct and rather peculiar syntax, debug the reasoner and find the right set of heuristics that would make it give an answer, and even with all that it still probably could not answer a question like, “If your snowman starts to do *this*, turn on the air conditioner”, Watson might actually have had a shot at answering something like this, but only because it knew from large language corpora that

‘snowman starts to melt’ is a common n-gram, not because it understands thermodynamics.

Working with people from Cycorp, or with anyone in the KR&R world, we became so enamoured of our elegant logic that, without a doubt, the knowledge became our focus. We—and I can say this with total confidence—we absolutely believed that getting the right answer was a trivial matter as long as you had the knowledge and reasoning right. The knowledge was the point.

“Knowledge is not the destination” refers to the epiphany that I had while working on Watson. The knowledge was important, but it wasn't the point—the point was to get answers right and to have confidence in them. If knowledge could not help with this, then it really was useless. But what kind of knowledge would help? Axioms about all the most general possible things in the world? Naïve physics? Expert Physics? Deep Aristotelean theories? No.

What mattered for Watson was having millions of simple “propositional” facts available at very high speed. Recognising entities by their names, knowing some basic type of information, knowing about very simple geospatial relationships like capitals and borders, where famous people were born and when, and much much more. Knowing all this was useful not because we looked up answers this way—*Jeopardy!* never asked about a person's age—but because these little facts could be stitched together with many other pieces of evidence from other sources to understand how confident we were in each answer.

This knowledge, a giant collection of subject-property-object triples, can be viewed as a graph. A very simple one, especially by KR&R standards, but this knowledge graph was not itself the goal of the project. The goal—the destination—of the project was winning *Jeopardy!* So, in fact, we made absolutely no effort to improve the knowledge we used from DBpedia and freebase. We needed to understand how well it worked for our problem in the general case, because there was no way to know what actual questions would be asked in the ultimate test in front of 50 million people.

### **Knowledge Graphs are Everywhere!**

As of the publication of this book, most major IT companies—more accurately, most major information companies—including Bloomberg, NY Times, Microsoft, Facebook, Twitter and many more, have significant knowledge graphs like Watson did, and have invested in their curation. Not because any of these graphs is their business, but because using this knowledge helps them in their business.

After Watson I moved to Google Research, where freebase lives on in our own humongous knowledge graph. And while Google invests a lot in its curation and maintenance, Google's purpose is not to build the greatest and most comprehensive knowledge graph on Earth, but to make a search, email, youtube, personal assistants and all the rest of our Web-scale services, better. That's our destination.

Many believe that the success of this kind of simplistic, propositional, knowledge graph proves that the original KR&R vision was a misguided mistake, but an outspoken few have gone so far as to claim it was a 40+ year waste of some great minds. As much as I appreciate being described as a great mind, I prefer a different

explanation: the work in KR for the past 40 years was not a waste of time, it was just the wrong place to start. It was solving a problem no one yet had, because no one had yet built systems that used this much explicit and declared knowledge.

Now, *knowledge graphs are everywhere*. Now industry is investing in the knowledge that drives their core systems. The editors of this volume, Jeff Pan, Guido Vetere, José Manuel Gómez Pérez and Honghan Wu, all themselves experts in this old yet burgeoning area of research, have gone to great lengths to put together research that matters today, in this world of large-scale graphs representing knowledge that makes a difference in the systems we use on the Web, on our phones, at work and at home.

The editorial team members have unique backgrounds, yet have worked together before, such as in the EU Marie Curie *K-Drive* project, and this book is a natural extension of their recent work on studying the properties of knowledge graphs. Jeff started at Manchester and has done a widely published work in formal reasoning systems, and moved to Aberdeen where his portfolio broadened considerably to include Machine Learning, large data analysis, and others, although he never strayed too far from practical reasoning, such as *approximate reasoning*, and querying for knowledge graphs. Guido has run several successful schema management projects on large data systems at IBM, and was part of the team that worked to bring Watson to Italy. Jose has done important research in the area of distributed systems, semantic data management and NLP, making knowledge easier to understand, access and consume by real users, and Honghan has been doing research in the area of medical knowledge systems.

After you finish this book, try to find a faded red copy of *Readings in Knowledge Representation* lest we forget and reinvent the Semantic Network.

May 2016

Dr. Christopher Welty  
Google Research NYC

# Preface

A few years after Google announced that their ‘Knowledge Graph’ would have allowed searching for *things, not strings*,<sup>1</sup> knowledge graphs start entering information retrieval, databases, Semantic Web, artificial intelligence, social media and enterprise information systems. But what exactly is Knowledge Graph? Where did it come from? What are the major differences between knowledge graphs for enterprise information management and those for Web search? What are the key components in a knowledge graph architecture? How can knowledge graphs help in enterprise information management? How can you build good quality knowledge graphs and utilise them to achieve your goals?

The main purpose of this book is to provide answers to these questions in a systematic way. Specifically, this book is for academic researchers, knowledge engineers and IT professionals who are interested in acquiring industrial experiences in using knowledge graphs for enterprises and large organisations. The book provides readers with an updated view on methods and technologies related to knowledge graphs, including illustrative corporate use cases.

In the last four years, we have been working hard and closely in the K-Drive—Knowledge Driven Data Exploitation—project (286348), which was funded by EU FP7/Marie Curie Industry-Academia Partnerships and Pathways schema/PEOPLE Work Programme. The main purpose of this project was to apply and extend advanced knowledge techniques to solve real-world problems, such as those in corporate knowledge management, healthcare and cultural heritage. Most of the challenges we encountered and techniques we dug into are highly related knowledge graph techniques. This book is a natural outcome of the K-Drive project that reflects and concludes the understanding we accumulated from the past four years of work, the lessons we have learned and the experiences we gained.

Contentwise, we will focus on the key technologies for constructing, understanding and consuming knowledge graphs, which constitute the three parts of this book, respectively. **Part I** introduces some background knowledge and technologies,

---

<sup>1</sup>Introducing the Knowledge Graph: things, not strings, [googleblog.blogspot.com](http://googleblog.blogspot.com) May 16, 2012



and then presents a simple architecture in order to help you to understand the main phases and tasks required during the lifecycle of knowledge graphs. **Part II** is the main technical part that starts with the state-of-the-art Knowledge Graph construction approaches, then focuses on exploration and exploitation techniques and finishes with advanced topics of Question Answering over/using knowledge graphs. Finally, **Part III** demonstrates successful stories of knowledge graph applications in Media Industry, Healthcare and Cultural Heritage; and ends with conclusions and future visions.

It is true that there is no *gold standard* definition of Knowledge Graph (KG). While working on the book, the editors and chapter contributors have debated lively on *what constitutes KG?, how is it related to relevant techniques like Semantic Web and Linked Data techniques?* and *what are its key features?* Fortunately, most, if not all, arguments have been settled and the conclusions and agreements have been put into the book, e.g. into the last two sections of Chap. 2. Even luckier, when finalising the book, editors have got the opportunity to collect opinions on *visions, barriers and next steps of Knowledge Graph* from key figures in the community including outstanding researchers, practitioners in leading organisations and start-ups, and representative users of various domains. Such valuable opinions have also been compiled into this book as part of its conclusion and future vision.

We would like to thank all of the chapter contributors as well as all members of the K-Drive project, who have given so much of their time and efforts for this book, in particular Dr. Yuting Zhao, who offered much helpful advice on the organisation of the book.

We had great pleasure in having Chris Welty write a touching Foreword for this book, sharing with us his rich experience and epiphany he had during the compelling BlueJ project, as well as his opinions on the motivation (*'Knowledge Graphs are Everywhere!'*) and the importance of this book.

We would also like to acknowledge the IBM DeepQA research team for allowing us to use their architecture diagram marked as Fig. 7.1 in the book.

We are grateful to the following experts in the field for sharing with us their visions, barriers and next steps of Knowledge Graph in our concluding chapter: Sören Auer, Riccardo Bellazzi, Oscar Corcho, Richard Dobson, Junlan Feng, Aldo Gangemi, Alfio M. Gliozzo, Tom Heath, Juanzi Li, Peter Mika, Fabrizio Renzi, Marco Varone, Denny Vrandečić and Haofen Wang.

Aberdeen, UK  
 Rome, Italy  
 Madrid, Spain  
 Aberdeen, UK  
 June 2016

Jeff Z. Pan  
 Guido Vetere  
 Jose Manuel Gomez-Perez  
 Honghan Wu

# Contents

<b>1 Enterprise Knowledge Graph: An Introduction</b> . . . . .	1
Jose Manuel Gomez-Perez, Jeff Z. Pan, Guido Vetere and Honghan Wu	
1.1 A Brief History of Knowledge Graph. . . . .	2
1.1.1 The Arrival of Semantic Networks . . . . .	2
1.1.2 From Semantic Networks to Linked Data . . . . .	3
1.1.3 Knowledge Graphs: An Entity-Centric View of Linked Data . . . . .	4
1.2 Knowledge Graph Technologies in a Nutshell . . . . .	5
1.3 Applications of Knowledge Graphs for Enterprise . . . . .	5
1.4 How to Read This Book. . . . .	10
1.4.1 Structure of This Book . . . . .	10
1.4.2 Who This Book Is For . . . . .	12
1.4.3 How to Use This Book . . . . .	12
<b>Part I Knowledge Graph Foundations &amp; Architecture</b>	
<b>2 Knowledge Graph Foundations</b> . . . . .	17
Boris Villazon-Terrazas, Nuria Garcia-Santa, Yuan Ren, Alessandro Faraotti, Honghan Wu, Yuting Zhao, Guido Vetere and Jeff Z. Pan	
2.1 Knowledge Representation and Query Languages . . . . .	18
2.1.1 RDF and RDFS . . . . .	18
2.1.2 OWL. . . . .	26
2.1.3 SPARQL. . . . .	30
2.2 Ontologies and Vocabularies . . . . .	33
2.2.1 Some Standard Vocabularies. . . . .	34
2.2.2 schema.org . . . . .	38
2.3 Data Lifting Standards . . . . .	39
2.3.1 RDB2RDF . . . . .	39
2.3.2 GRDDL . . . . .	48

- 2.4 Knowledge Graph Versus Linked Data . . . . . 51
- 2.5 Knowledge Graph for Web Searching and Knowledge Graph for Enterprise . . . . . 52
- 3 Knowledge Architecture for Organisations . . . . . 57**  
 Ronald Denaux, Yuan Ren, Boris Villazon-Terrazas,  
 Panos Alexopoulos, Alessandro Faraotti and Honghan Wu
- 3.1 Architecture Overview . . . . . 58
- 3.2 Acquisition and Integration Layer. . . . . 60
  - 3.2.1 Ontology Development . . . . . 60
  - 3.2.2 Ontologisation of Non-Ontological Resources. . . . . 61
  - 3.2.3 Text Integration via Named Entity and Thematic Scope Resolution . . . . . 62
  - 3.2.4 Ontology Learning . . . . . 64
- 3.3 Knowledge Storing and Accessing Layer . . . . . 65
  - 3.3.1 Ontology-Based Data Access . . . . . 65
  - 3.3.2 RDF Stores . . . . . 69
  - 3.3.3 Property Graph-Based Stores . . . . . 73
  - 3.3.4 Conclusion: Storing Knowledge Graphs Versus Relational Databases . . . . . 75
- 3.4 Knowledge Consumption Layer . . . . . 77
  - 3.4.1 Semantic Search . . . . . 77
  - 3.4.2 Summarisation. . . . . 78
  - 3.4.3 Query Generation . . . . . 80
  - 3.4.4 Question Answering . . . . . 81
  - 3.4.5 Conclusion . . . . . 83

**Part II Constructing, Understanding and Consuming Knowledge Graphs**

- 4 Construction of Enterprise Knowledge Graphs (I) . . . . . 87**  
 Boris Villazon-Terrazas, Nuria Garcia-Santa, Yuan Ren,  
 Kavitha Srinivas, Mariano Rodriguez-Muro, Panos Alexopoulos  
 and Jeff Z. Pan
- 4.1 Knowledge Construction and Maintenance Lifecycle . . . . . 88
- 4.2 Ontology Authoring: A Competency Question-Driven Approach . . . . . 93
  - 4.2.1 Competency Questions . . . . . 94
  - 4.2.2 Formulation of Competency Questions . . . . . 96
  - 4.2.3 Ontology Authoring Workflow . . . . . 98
- 4.3 Semi-automated Linking of Enterprise Data for Virtual Knowledge Graphs . . . . . 104
  - 4.3.1 Virtual Knowledge Graph for Knowledge Discovery . . . . . 104
  - 4.3.2 Semantic Tagging and Data Interlinking . . . . . 106

4.3.3	Usage Scenarios . . . . .	109
4.3.4	Conclusion . . . . .	116
<b>5</b>	<b>Construction of Enterprise Knowledge Graphs (II)*</b> . . . . .	<b>117</b>
	Panos Alexopoulos, Yuting Zhao, Jeff Z. Pan and Man Zhu	
5.1	Scenario-Driven Named Entity and Thematic Scope	
	Resolution of Unstructured Data* . . . . .	117
5.1.1	Framework Description . . . . .	118
5.1.2	Framework Application Evaluation . . . . .	123
5.2	Open-World Schema Learning for Knowledge Graphs* . . . . .	126
5.2.1	Motivation. . . . .	126
5.2.2	BelNet <sup>+</sup> . . . . .	128
5.2.3	TBox Learning as Inference . . . . .	133
5.2.4	A Novel Evaluation Framework . . . . .	133
5.2.5	Experiments . . . . .	136
5.2.6	Experimental Results. . . . .	139
5.2.7	Related Work and Summary . . . . .	144
5.2.8	Conclusion . . . . .	145
<b>6</b>	<b>Understanding Knowledge Graphs</b> . . . . .	<b>147</b>
	Honghan Wu, Ronald Denaux, Panos Alexopoulos, Yuan Ren and Jeff Z. Pan	
6.1	Understanding Things in KGs: The Summary of Individual Entities . . . . .	149
6.1.1	Entity Data Organisation . . . . .	150
6.1.2	Summarisation of Entity Data . . . . .	156
6.1.3	Conclusion . . . . .	157
6.2	Exploring KGs: The Summary of Entity Description Patterns . . . . .	157
6.2.1	What Is the Entity Description Pattern? . . . . .	159
6.2.2	How Can the Entity Description Pattern Help in Knowledge Exploitations? . . . . .	161
6.2.3	Conclusion . . . . .	163
6.3	Profiling KGs: A Goal-Driven Summarisation . . . . .	164
6.3.1	Motivating Scenario and Problem Definition. . . . .	164
6.3.2	Framework Description . . . . .	164
6.3.3	Implementation . . . . .	166
6.3.4	Application Example . . . . .	169
6.3.5	Conclusion . . . . .	171
6.4	Revealing Insights from KGs: A Query Generation Approach* . . . . .	171
6.4.1	Candidate Insightful Queries . . . . .	174
6.4.2	Query Generation Framework . . . . .	177
6.4.3	Evaluation of the Query Generation Method . . . . .	178
6.4.4	Conclusion and Future Work . . . . .	180

**7 Question Answering and Knowledge Graphs** . . . . . 181  
Alessandro Moschitti, Kateryna Tymoshenko,  
Panos Alexopoulos, Andrew Walker, Massimo Nicosia,  
Guido Vetere, Alessandro Faraotti, Marco Monti, Jeff Z. Pan,  
Honghan Wu and Yuting Zhao

7.1 Question Answering over Text Documents. . . . . 182  
7.1.1 Realising a QA System: Approaches and Key Steps. . . . . 182

7.2 Question Answering over Knowledge Graphs . . . . . 189  
7.2.1 State-of-the-Art Approaches for Question Answering  
Over Knowledge Graphs. . . . . 190  
7.2.2 Question Answering in the Enterprise. . . . . 193

7.3 Knowledge Graph and Watson DeepQA . . . . . 196  
7.3.1 What Is Watson DeepQA? . . . . . 196  
7.3.2 What Are the Knowledge Graphs Used in Watson  
DeepQA? . . . . . 198  
7.3.3 How Knowledge Graphs Are Used in Watson  
DeepQA? . . . . . 199  
7.3.4 Lessons Learnt from Watson DeepQA . . . . . 201

7.4 Using Knowledge Graphs for Improving Textual Question  
Answering\* . . . . . 202  
7.4.1 A Flexible QA Pipeline. . . . . 202  
7.4.2 Exploiting External Knowledge (Graphs)  
for Re-ranking. . . . . 204  
7.4.3 Evaluation: Impact of Knowledge Graphs  
in Semantic Structures. . . . . 209  
7.4.4 Conclusion . . . . . 212

**Part III Industrial Applications and Successful Stories**

**8 Success Stories**. . . . . 215  
Marco Monti, Fernanda Perego, Yuting Zhao, Guido Vetere,  
Jose Manuel Gomez-Perez, Panos Alexopoulos, Hai Nguyen,  
Gemma Webster, Boris Villazon-Terrazas, Nuria Garcia-Santa  
and Jeff Z. Pan

8.1 A Knowledge Graph for Innovation in the Media Industry . . . . . 216  
8.1.1 The Business Problem. . . . . 216  
8.1.2 The HAVAS 18 Knowledge Graph. . . . . 217  
8.1.3 Value Proposition . . . . . 218  
8.1.4 Challenges. . . . . 219

8.2 Applying Knowledge Graphs in Cultural Heritage . . . . . 219  
8.2.1 Digital Cultural Heritage and Linked Data . . . . . 219  
8.2.2 The Challenges . . . . . 220  
8.2.3 The CURIOS Project . . . . . 221  
8.2.4 Constructing the Knowledge Graph . . . . . 223

- 8.2.5 CURIOS—A Linked Data Adaptor for Content Management Systems . . . . . 224
- 8.2.6 Presenting and Visualising Cultural Heritage Knowledge Graphs . . . . . 225
- 8.2.7 Collaborative Construction and Maintenance of Cultural Heritage Knowledge Graphs . . . . . 228
- 8.3 Applying Knowledge Graphs in Healthcare . . . . . 228
  - 8.3.1 The Problem in Clinical Practice Guidelines . . . . . 229
  - 8.3.2 Preparing the Data and Building the Knowledge Graphs . . . . . 230
  - 8.3.3 Services Based on the Knowledge Graphs . . . . . 233
  - 8.3.4 Contributions to Healthcare Practice . . . . . 235
- 9 Enterprise Knowledge Graph: Looking into the Future.** . . . . . 237
  - Jeff Z. Pan, Jose Manuel Gomez-Perez, Guido Vetere, Honghan Wu, Yuting Zhao and Marco Monti
  - 9.1 Conclusion . . . . . 238
  - 9.2 Get Started with Knowledge Graphs. . . . . 239
    - 9.2.1 A Small but Powerful Knowledge Graph . . . . . 240
    - 9.2.2 Troubleshooting . . . . . 242
    - 9.2.3 Variations . . . . . 243
  - 9.3 What is Next: Experts’ Predictions into the Future of Knowledge Graph . . . . . 243
    - 9.3.1 Future Visions. . . . . 244
    - 9.3.2 Foreseeable Obstacles . . . . . 246
    - 9.3.3 Suggestions on Next Steps . . . . . 248
- References** . . . . . 251
- Index** . . . . . 265

# Contributors

- Panos Alexopoulos** Expert System, Madrid, Spain
- Ronald Denaux** Expert System, Madrid, Spain
- Alessandro Faraotti** IBM Italia, Rome, Italy
- Nuria Garcia-Santa** Expert System, Madrid, Spain
- Jose Manuel Gomez-Perez** Expert System, Madrid, Spain
- Marco Monti** IBM Italia, Milan, Italy
- Alessandro Moschitti** University of Trento, Trento, Italy
- Hai Nguyen** University of Aberdeen, King's College, Aberdeen, UK
- Massimo Nicosia** University of Trento, Trento, Italy
- Jeff Z. Pan** University of Aberdeen, King's College, Aberdeen, UK
- Fernanda Perego** IBM Italia, Milan, Italy
- Yuan Ren** University of Aberdeen, King's College, Aberdeen, UK
- Mariano Rodriguez-Muro** IBM USA, Thomas J. Watson Research Center, Yorktown Heights, NY, USA
- Kavitha Srinivas** IBM USA, Thomas J. Watson Research Center, Yorktown Heights, NY, USA
- Kateryna Tymoshenko** Trento RISE, Povo di Trento, Trento, Italy
- Guido Vetere** IBM Italia, Rome, Italy
- Boris Villazon-Terrazas** Expert System, Madrid, Spain
- Andrew Walker** University of Aberdeen, King's College, Aberdeen, UK
- Gemma Webster** University of Aberdeen, King's College, Aberdeen, UK

**Honghan Wu** King's College London, London, UK

**Yuting Zhao** IBM Italia, Milan, Italy

**Man Zhu** Southeast University, Nanjing, China