

# Contents

|   |    |
|---|----|
| <b>1 Enterprise Knowledge Graph: An Introduction</b> . . . . .  | 1  |
| Jose Manuel Gomez-Perez, Jeff Z. Pan, Guido Vetere<br>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<br>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 |
| <br>  |    |
| <b>Part I Knowledge Graph Foundations &amp; Architecture</b>  |    |
| <b>2 Knowledge Graph Foundations</b> . . . . .  | 17 |
| Boris Villazon-Terrazas, Nuria Garcia-Santa, Yuan Ren,<br>Alessandro Faraotti, Honghan Wu, Yuting Zhao,<br>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,<br>Yuan Ren and Jeff Z. Pan            |            |
| 6.1      | Understanding Things in KGs: The Summary of Individual<br>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<br>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<br>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



<http://www.springer.com/978-3-319-45652-2>

Exploiting Linked Data and Knowledge Graphs in Large Organisations

Pan, J.Z.; Vetere, G.; Gomez-Perez, J.M.; Wu, H. (Eds.)

2017, XVIII, 266 p. 59 illus., 44 illus. in color., Hardcover

ISBN: 978-3-319-45652-2