

Contents

Acknowledgment — V

List of chapter authors — XV

Emmanuel Gillain

1 Preface and introduction — 1

- 1.1 Fast progress in artificial intelligence (AI) — **1**
- 1.2 Convergence of factors — **5**
 - 1.2.1 The globalization and democratization of access to computing and AI resources — **6**
 - 1.2.2 Dematerialization, digitization, and exponential growth of data — **7**
 - 1.2.3 AI accelerators: AI techniques that speed up the adoption by the industry — **8**
 - 1.2.4 The dynamics of software, amplified by open-source: reuse, iterative improvement and innovation — **9**
- 1.3 Opportunities and challenges for the industry — **10**
- 1.4 Purpose and target audience of the book: enlighten the business practitioners — **15**

Emmanuel Gillain

2 A holistic view of AI techniques, their limitations and complementarities — 19

- 2.1 What is artificial intelligence (AI)? — **19**
- 2.2 Decision by search and how an AI agent can construct a plan of actions — **21**
 - 2.2.1 Overview and key concepts — **21**
 - 2.2.2 Industry applications and examples in the book — **21**
 - 2.2.3 Moving to constraint search — **22**
 - 2.2.4 Constraint satisfaction propagation: applications and examples in the book — **23**
 - 2.2.5 Limitations of search-based techniques — **23**
- 2.3 Represent knowledge and reason with symbolic AI — **24**
 - 2.3.1 Overview and key concepts — **25**
 - 2.3.2 Industry applications and examples in the book — **27**
 - 2.3.3 Limitations of symbolic AI — **28**
- 2.4 Learning with data-driven AI — **30**
 - 2.4.1 Overview — **30**
 - 2.4.2 Supervised learning key concepts, industry applications and examples in the book — **31**
 - 2.4.3 Unsupervised learning – key concepts, industry applications, and examples in the book — **32**

2.4.4	Anomaly detection – key concepts, industry applications, and examples in the book —	34
2.4.5	Reinforcement learning: key concepts, industry applications, and examples in the book —	35
2.4.6	What about deep learning? —	37
2.4.7	Limitations of data-driven learning —	41
2.5	Learning in a logic framework —	42
2.5.1	Overview and key concepts —	42
2.5.2	Industry applications and examples in the book —	43
2.5.3	Limitations of inductive logic programming —	44
2.6	Probabilistic reasoning when there is uncertainty —	44
2.6.1	Key concepts—uncertain facts and outcomes with probabilistic reasoning —	45
2.6.2	Industry application and examples from the book —	48
2.6.3	Probabilistic reasoning over time, or sequence to sequence —	49
2.6.4	Industry applications and examples from the book —	50
2.6.5	Enrich the abstraction—unify probabilistic models and formal logic —	50
2.6.6	Uncertain states or outcomes in the context of a decision —	51
2.6.7	Limitations of probabilistic reasoning —	52
2.7	Interaction with machines in natural language —	52
2.7.1	Overview of some key concepts —	53
2.7.2	Industry applications and examples in the book —	58
2.7.3	Limitations of natural language processing —	59
2.8	The importance of an ethical approach to AI —	60

Yves Deville

3	Solve problems by searching, including with constraints, a fundamental pillar —	63
3.1	Why is solving problems by search important within the broader artificial intelligence (AI) domain? —	63
3.2	Search algorithms —	64
3.2.1	What category of problems do search algorithms solve? —	64
3.2.2	Solving problems by search without heuristics —	65
3.2.3	Solving problems by search with heuristics —	70
3.2.4	What are the limitations of search algorithms? —	73
3.2.5	Industry example. The fastest route with search and heuristics, using Bing Maps Routes API —	75
3.3	Constraint satisfaction problems —	80
3.3.1	What category of problems does constraint satisfaction solve? —	80
3.3.2	Solving optimization problems by local search —	84
3.3.3	Solving problems by constraint programming —	88

- 3.3.4 What are the limitations of constraint satisfaction problems? — **93**
- 3.3.5 Industry example: containers configuration using CSP at Maersk Container Industry A/S — **94**
- 3.4 SAT solvers — **97**
 - 3.4.1 What category of problems does SAT solvers solve? — **97**
 - 3.4.2 Solving problems with SAT solvers — **98**
 - 3.4.3 What are the limitations of SAT solvers? — **103**
 - 3.4.4 Industry example: containers configuration using SAT, at Maersk Container Industry A/S — **104**

Bart Bogaerts

- 4 Reasoning with first-order logic — 107**
 - 4.1 Why is reasoning with first-order logic important within the broader artificial intelligence (AI) domain? — **107**
 - 4.2 What category of problems does reasoning with first-order logic solve? — **109**
 - 4.3 How are those problems solved? — **111**
 - 4.3.1 Representing knowledge in first-order logic — **112**
 - 4.3.2 Reasoning with first-order logic — **120**
 - 4.4 What are the limitations of reasoning with first-order logic — **126**
 - 4.5 Industry examples — **127**
 - 4.5.1 Automated design-driven diagnostics for lithography machines at ASML — **127**
 - 4.5.2 Modeling and verifying simple vehicle controller, such as the Triton unmanned aircraft systems of the US Navy: using Imandra system and first-order logic — **136**

Isabelle Linden

- 5 Knowledge representation and engineering with ontologies — 149**
 - 5.1 Why are knowledge representation and engineering with ontologies important within the broader AI domain? — **149**
 - 5.2 What category of problems do knowledge representation and engineering with ontologies solve? — **151**
 - 5.3 How are those problems solved? — **153**
 - 5.3.1 \mathcal{ALC} description logic formal definition — **154**
 - 5.3.2 Queries and reasoning problems — **158**
 - 5.4 What are the limitation of knowledge representation and engineering with ontologies? — **165**
 - 5.4.1 Expressiveness and complexity of DLs — **165**
 - 5.4.2 DLs expressiveness limitations and challenges — **167**
 - 5.4.3 Knowledge engineering — **168**
 - 5.4.4 OWL and the semantic web — **168**

- 5.4.5 Knowledge graphs — 169
- 5.4.6 Conclusion — 170
- 5.5 Industry example: drug development using ontologies and semantic search, the ONTOFORCE example — 171
 - 5.5.1 Context — 171
 - 5.5.2 Data is ubiquitous but siloed — 171
 - 5.5.3 ONTOFORCE — 172
 - 5.5.4 Use case introduction: semantics and ontologies in basic research — 174
 - 5.5.5 Use case scenario: semantics and linked data at work for knowledge discovery in early-stage drug research — 174
 - 5.5.6 Application of early-drug research knowledge for translational research use case — 176

Aleksandra Pižurica

- 6 Probabilistic reasoning: When the environment is uncertain — 181**
 - 6.1 Why is probabilistic reasoning important within the broader AI domain? — 181
 - 6.2 What category of problems does probabilistic reasoning solve? — 183
 - 6.3 How probabilistic reasoning problems are solved? — 185
 - 6.3.1 Modeling causal problems with Bayesian networks — 185
 - 6.3.2 Modeling noncausal problems with Markov random fields — 188
 - 6.3.3 Modeling noncausal problems with conditional random fields — 193
 - 6.3.4 A unifying representation: factor graphs — 196
 - 6.3.5 Approximate inference algorithms — 198
 - 6.3.6 Probabilistic reasoning over time — 204
 - 6.4 What are the limitations of probabilistic reasoning? — 209
 - 6.5 Industry examples — 211
 - 6.5.1 Hunting for anomalous sessions using Markov chain model in Microsoft Sentinel — 211
 - 6.5.2 Crop forecasting with Bayesian inference for agriculture producers — 218
 - 6.5.3 Estimating players' skills with XBOX TrueSkill™ — 224
 - 6.6 Useful reminders of probability theory — 233
 - 6.6.1 Basic concepts in probability — 233
 - 6.6.2 Maximum likelihood estimation and Bayesian estimation — 236

Hendrik Blockeel

- 7 Learning from data — 241**
 - 7.1 Why is learning important within the broader AI domain? — 241
 - 7.2 What category of problems does machine learning solve? — 243
 - 7.2.1 Learning predictive functions — 243
 - 7.2.2 Learning descriptive models — 247

7.3	How are learning problems solved? — 250
7.3.1	From data to model: an overview — 250
7.3.2	Data collection and preparation — 250
7.3.3	Format of the data — 251
7.3.4	Paradigms for learning — 256
7.3.5	Instance-based methods — 258
7.3.6	Decision trees — 259
7.3.7	Ensembles of decision trees — 262
7.3.8	If-then-rules — 265
7.3.9	Statistical model fitting — 269
7.3.10	Support vector machines — 273
7.3.11	Neural networks — 277
7.3.12	Dimensionality reduction — 287
7.3.13	Matrix factorization and tensor decomposition — 289
7.3.14	Probabilistic graphical models — 291
7.3.15	Clustering and density estimation — 297
7.3.16	Automata and hidden Markov models — 299
7.3.17	Reinforcement learning — 301
7.3.18	Other aspects of learning — 306
7.4	Evaluating the results — 309
7.5	What are the limitations of learning? — 312
7.6	Industry examples — 314
7.6.1	Predicting the metallization rate in iron production using ensemble methods at ArcelorMittal — 314
7.6.2	Estimating the value of real estate with supervised regression models, at KBC Group — 319
7.6.3	Fraud detection in insurance claims with unsupervised outlier detection, at KBC Group — 322
7.6.4	Defect detection in textile using CNN, <i>k</i> -means clustering, and unsupervised anomaly detection, at Veranneman Technical Textiles — 329
7.6.5	Extracting information from forms using clustering techniques with Microsoft Form Recognizer — 333
7.6.6	Recommending pages in Microsoft News with reinforcement learning — 338

Walter Daelemans

8	Between language and knowledge — 343
8.1	Why is natural language processing important within the broader AI domain? — 343
8.1.1	A short history of NLP — 344
8.2	What category of problems does natural language processing solve? — 346

8.2.1	NLP tasks —	348
8.2.2	Speech to text processing —	348
8.2.3	Image to text processing —	349
8.2.4	Text to text processing —	350
8.3	How are natural language processing problems solved? —	352
8.3.1	Static word embeddings —	352
8.3.2	Language models —	354
8.3.3	Transformers —	357
8.3.4	Contextual word embeddings —	359
8.3.5	From text to knowledge and back —	362
8.4	What are the limitations of natural language processing? —	366
8.5	Industry examples —	367
8.5.1	Automate question answering using NLP with Microsoft QnAMaker —	368
8.5.2	Extract information from contracts using NLP, the case of Daimler —	376
8.5.3	Analyze legal content for lawyers at AXA's legal protection unit —	384
8.6	Addendum —	392

Erik Mannens

9	Some words about ethics. The angles of fairness and transparency —	395
9.1	The challenges of ethical AI —	395
9.2	Initial framing of fairness —	396
9.2.1	Generic definition on bias —	399
9.2.2	What is the problem to solve? —	399
9.2.3	How do we mitigate it? —	404
9.2.4	When do we apply these techniques? —	406
9.2.5	What are the limitations? —	407
9.3	Initial framing of Interpretability —	409
9.3.1	What is the problem to solve? —	410
9.3.2	How do we mitigate it? —	411
9.3.3	When do we apply these techniques? —	414
9.3.4	What are the limitations? —	416
9.4	Industry examples —	417
9.4.1	Fair loan adjudication models with Fairlearn at EY —	417
9.4.2	Detect and reduce fraud for loyalty services using InterpretML to respect ethical AI principles —	424
9.5	Final conclusions on AI ethics —	428

Oussama Chelly and Hendrik Blockeel

10	Industry examples where different AI techniques are combined —	431
10.1	An AI-enabled chatbot for the Casco insurance industry, an example from KBC Group —	431

- 10.2 An automated engineering assistant that uses a mix of learning and reasoning techniques in manufacturing — **436**
- 10.2.1 The problem setting — **437**
- 10.2.2 How are those problems solved ? A mix of techniques from Chapters 4, 5, and 7 — **438**
- 10.2.3 How well does it work? — **443**

Emmanuel Gillain

- 11 Conclusion – Moving forward — 445**
- 11.1 So far... — **445**
- 11.2 Moving forward — **445**
- 11.2.1 Hybrid AI systems to merge the best of both worlds — **446**
- 11.2.2 Cause-effect and causal AI — **447**
- 11.2.3 Generative AI — **449**
- 11.3 Final word — **451**

Trademarks — 453

Index — 455