

The Handbook of Data Science and AI

Generate Value from Data with Machine
Learning and Data Analytics

DAS INHALTS- VERZEICHNIS

» Hier geht's
direkt
zum Buch

Table of Contents

Preface	XXIX
Acknowledgments	XXXI
1 Introduction	1
<i>Stefan Papp</i>	
1.1 About this Book	1
1.2 The Halford Group	2
1.2.1 Alice Halford – Chairwoman	3
1.2.2 Analysts	4
1.2.3 “CDO”	6
1.2.4 Sales	7
1.2.5 IT	8
1.2.6 Security	9
1.2.7 Production Leader	10
1.2.8 Customer Service	11
1.2.9 HR	12
1.2.10 CEO	13
1.3 In a Nutshell	15
2 The Alpha and Omega of AI	17
<i>Stefan Papp</i>	
2.1 The Data Use Cases	18
2.1.1 Bias	18
2.1.2 Data Literacy	21

2.2	Culture Shock	23
2.3	Ideation	26
2.4	Design Process Models	28
2.4.1	Design Thinking	29
2.4.2	Double Diamond	29
2.4.3	Conducting Workshops	31
2.5	In a Nutshell	38
3	Cloud Services	39
	<i>Stefan Papp</i>	
3.1	Introduction	39
3.2	Cloud Essentials	40
3.2.1	XaaS	42
3.2.2	Cloud Providers	44
3.2.3	Native Cloud Services	45
3.2.4	Cloud-native Paradigms	48
3.3	Infrastructure as a Service	50
3.3.1	Hardware	50
3.3.2	Distributed Systems	52
3.3.3	Linux Essentials for Data Professionals	56
3.3.4	Infrastructure as Code	63
3.4	Platform as a Service	67
3.4.1	Cloud Native PaaS Solutions	68
3.4.2	External Solutions	72
3.5	Software as a Service	75
3.6	In a Nutshell	76
4	Data Architecture	77
	<i>Zoltan C. Toth and Sean McIntyre</i>	
4.1	Overview	77
4.1.1	Maslow's Hierarchy of Needs for Data	78
4.1.2	Data Architecture Requirements	80
4.1.3	The Structure of a Typical Data Architecture	80
4.1.4	ETL (Extract, Transform, Load)	85
4.1.5	ELT (Extract, Load, Transform)	85
4.1.6	ETLT	86

4.2	Data Ingestion and Integration	87
4.2.1	Data Sources	87
4.2.2	Traditional File Formats.....	89
4.2.3	Modern File Formats.....	91
4.2.4	Which Storage Option to Choose?	93
4.3	Data Warehouses, Data Lakes, and Lakehouses	94
4.3.1	Data Warehouses	94
4.3.2	Data Lakes and Cloud Data Platforms.....	98
4.4	Data Transformation	102
4.4.1	SQL	103
4.4.2	Big Data & Apache Spark	113
4.4.3	Cloud Data Platforms for Apache Spark.....	121
4.5	Workflow Orchestration	123
4.5.1	Dagster and the Modern Data Stack	125
4.6	A Data Architecture Use Case.....	127
4.7	In a Nutshell.....	132
5	Data Engineering	133
	<i>Stefan Papp</i>	
5.1	Differentiating from Software Engineering	134
5.2	Programming Languages	135
5.2.1	Code or No Code?.....	136
5.2.2	Language Ecosystem.....	137
5.2.3	Python.....	139
5.2.4	Scala	142
5.3	Software Engineering Processes for Data	145
5.3.1	Configuration Management	145
5.3.2	CI/CD	146
5.4	Data Pipelines	148
5.4.1	Common Characteristics of a Data Pipeline	149
5.4.2	Data Pipelines in the Unified Data Architecture	149
5.5	Storage Options	153
5.5.1	File Era.....	153
5.5.2	Database Era	154
5.5.3	Data Lake Era	156
5.5.4	Serverless Era.....	157

5.5.5	Polyglot Storage	157
5.5.6	Data Mesh Era	159
5.6	Tooling	160
5.6.1	Batch: Airflow	160
5.6.2	Streaming: Kafka	162
5.6.3	Transformation: Databricks Notebooks	167
5.7	Common Challenges	168
5.7.1	Data Quality and Different Standards	168
5.7.2	Skewed Data	170
5.7.3	Stressed Operational Systems	171
5.7.4	Legacy Operational Systems	172
5.7.5	Platform and Information Security	172
5.8	In a Nutshell	173
6	Data Governance	175
	<i>Victoria Rugli, Mario Meir-Huber</i>	
6.1	Why Do We Need Data Governance?	175
6.1.1	Sample 1: Achieving Clarity with Data Governance	176
6.1.2	Sample 2: The (Negative) Impact of Poor Data Governance	177
6.2	The Building Blocks of Data Governance	178
6.2.1	Data Governance Explained	180
6.3	People	181
6.3.1	Data Ownership	182
6.3.2	Data Stewardship	185
6.3.3	Data Governance Board	186
6.3.4	Change Management	188
6.4	Process	189
6.4.1	Metadata Management	190
6.4.2	Data Quality Management	194
6.4.3	Data Security and Privacy	198
6.4.4	Master Data Management	201
6.4.5	Data Access and Search	205
6.5	Technology (Data Governance Tools)	207
6.5.1	Open-Source Tools	208
6.5.2	Cloud-based Data Governance Tools	214
6.6	In a Nutshell	218

7	Machine Learning Operations (ML Ops)	219
	<i>Zoltan C. Toth, György Móra</i>	
7.1	Overview	219
7.1.1	Scope of MLOps	220
7.1.2	Data Collection and Exploration	221
7.1.3	Feature Engineering	221
7.1.4	Model Training	221
7.1.5	Models Deployed to Production	222
7.1.6	Model Evaluation	223
7.1.7	Model Understanding	223
7.1.8	Model Versioning	224
7.1.9	Model Monitoring	224
7.2	MLOps in an Organization	224
7.2.1	Main Benefits of MLOps	225
7.2.2	Capabilities Needed for MLOps	225
7.3	Several Common Scenarios in the MLOps Space	226
7.3.1	Integrating Notebooks	226
7.3.2	Features in Production	227
7.3.3	Model Deployment	230
7.3.4	Model Formats	230
7.4	MLOps Tooling and MLflow	231
7.4.1	MLflow	232
7.5	In a Nutshell	235
8	Machine Learning Security	237
	<i>Manuel Pasieka</i>	
8.1	Introduction to Cybersecurity	238
8.2	Attack Surface	239
8.3	Attack Methods	241
8.3.1	Model Stealing	241
8.3.2	Data Extraction	243
8.3.3	Data Poisoning	245
8.3.4	Adversarial Attack	248
8.3.5	Backdoor Attack	251

8.4	Machine Learning Security of Large Language Models	254
8.4.1	Data Extraction	254
8.4.2	Jailbreaking	255
8.4.3	Prompt Injection	257
8.5	AI Threat Modelling	260
8.6	Regulations	261
8.7	Where to go from here	263
8.8	Conclusion	265
8.9	In a Nutshell	266
9	Mathematics	267
	<i>Annalisa Cadonna</i>	
9.1	Linear Algebra	268
9.1.1	Vectors and Matrices	268
9.1.2	Operations between Vectors and Matrices	272
9.1.3	Linear Transformations	275
9.1.4	Eigenvalues, Eigenvectors, and Eigendecomposition	276
9.1.5	Other Matrix Decompositions	277
9.2	Calculus and Optimization	279
9.2.1	Derivatives	279
9.2.2	Gradient and Hessian	282
9.2.3	Gradient Descent	283
9.2.4	Constrained Optimization	285
9.3	Probability Theory	286
9.3.1	Discrete and Continuous Random Variables	287
9.3.2	Expected Value, Variance, and Covariance	291
9.3.3	Independence, Conditional Distributions, and Bayes' Theorem	293
9.4	In a Nutshell	295
10	Statistics – Basics	297
	<i>Rania Wazir, Georg Langs, Annalisa Cadonna</i>	
10.1	Data	298
10.2	Simple Linear Regression	300
10.3	Multiple Linear Regression	307
10.4	Logistic Regression	310
10.5	How Good is Our Model?	317
10.6	In a Nutshell	319

11	Business Intelligence (BI)	321
	<i>Christian Mata</i>	
11.1	Introduction to Business Intelligence	323
11.1.1	Definition of Business Intelligence	324
11.1.2	Role in Organizations	324
11.1.3	Development of Business Intelligence	325
11.1.4	Data Science and AI in the Context of BI	327
11.1.5	Data for Decision-Making	330
11.1.6	Understanding Business Context	331
11.1.7	Business Intelligence Activities	334
11.2	Data Management Fundamentals	336
11.2.1	What is Data Management, Data Integration and Data Warehousing?	336
11.2.2	Data Load Processes – The Case of ETL or ELT	338
11.2.3	Data Modeling	339
11.3	Reporting and Data Analysis	346
11.3.1	Reporting	347
11.3.2	Types of Reports	350
11.3.3	Data Analysis	351
11.3.4	Visual Analysis	353
11.3.5	Significant Trends	355
11.3.6	Relevant BI Technologies	356
11.3.7	BI Tool Examples	360
11.4	BI and Data Science: Complementary Disciplines	363
11.4.1	Differences	363
11.4.2	Similarities	364
11.4.3	Interdependencies	364
11.5	Outlook for Business Intelligence	365
11.5.1	Expectations for the Evolution of BI	366
11.6	In a Nutshell	368
12	Machine Learning	369
	<i>Georg Langs, Katherine Munro, Rania Wazir</i>	
12.1	Introduction	369
12.2	Basics: Feature Spaces	371

12.3	Classification Models	374
12.3.1	K-Nearest-Neighbor-Classifier	375
12.3.2	Support Vector Machine	376
12.3.3	Decision Trees	377
12.4	Ensemble Methods	379
12.4.1	Bias and Variance	379
12.4.2	Bagging: Random Forests	381
12.4.3	Boosting: AdaBoost	385
12.4.4	The Limitations of Feature Construction and Selection	385
12.5	Unsupervised learning: Learning without labels	386
12.5.1	Clustering	386
12.5.2	Manifold Learning	387
12.5.3	Generative Models	388
12.6	Artificial Neural Networks and Deep Learning	389
12.6.1	The Perceptron	389
12.6.2	Artificial Neural Networks	390
12.6.3	Deep Learning	393
12.6.4	Convolutional Neural Networks	393
12.6.5	Training Convolutional Neural Networks	394
12.6.6	Recurrent Neural Networks	396
12.6.7	Long Short-Term Memory	398
12.6.8	Autoencoders and U-Nets	399
12.6.9	Adversarial Training Approaches	401
12.6.10	Generative Adversarial Networks	402
12.6.11	Cycle GANs and Style GANs	404
12.7	Transformers and Attention Mechanisms	405
12.7.1	The Transformer Architecture	406
12.7.2	What the Attention Mechanism Accomplishes	407
12.7.3	Applications of Transformer Models	408
12.8	Reinforcement Learning	409
12.9	Other Architectures and Learning Strategies	412
12.10	Validation Strategies for Machine Learning Techniques	413
12.11	Conclusion	414
12.12	In a Nutshell	415

13	Building Great Artificial Intelligence	417
	<i>Danko Nikolić</i>	
13.1	How AI Relates to Data Science and Machine Learning	418
13.2	A Brief History of AI	422
13.3	Five Recommendations for Designing an AI Solution	424
13.3.1	Recommendation No. 1: Be Pragmatic	424
13.3.2	Recommendation No. 2: Make it Easier for Machines to Learn – Create Inductive Biases	426
13.3.3	Recommendation No. 3: Perform Analytics	432
13.3.4	Recommendation No. 4: Beware of the Scaling Trap	435
13.3.5	Recommendation No. 5: Beware of the Generality Trap (there is no such a thing as free lunch)	445
13.4	Human-level Intelligence	450
13.5	In a Nutshell	453
14	Signal Processing	455
	<i>Jana Eder</i>	
14.1	Introduction	456
14.2	Sampling and Quantization	457
14.3	Frequency Domain Analysis	460
14.3.1	Fourier Transform	461
14.4	Noise Reduction and Filtering Techniques	467
14.4.1	Denoising Using a Gaussian Low-pass Filter	468
14.5	Time Domain Analysis	470
14.5.1	Signal Normalization and Standardization	470
14.5.2	Signal Transformation and Feature Extraction	471
14.5.3	Time Series Decomposition Techniques	474
14.5.4	Autocorrelation: Understanding Signal Similarity over Time	477
14.6	Time-Frequency Domain Analysis	480
14.6.1	Short Term Fourier Transform and Spectrogram	480
14.6.2	Discrete Wavelet Transform	481
14.6.3	Gramian Angular Field	482
14.7	The Relationship of Signal Processing and Machine Learning	484
14.7.1	Techniques for Feature Engineering	485
14.7.2	Preparing for Machine Learning	485

- 14.8 Practical Applications 486
- 14.9 In a Nutshell 488
- 15 Foundation Models 489**
 - Danko Nikolić*
 - 15.1 The Idea of a Foundation Model 489
 - 15.2 How to Train a Foundation Model? 493
 - 15.3 How Do we Use Foundation Models? 495
 - 15.4 A Breakthrough: There is no End to Learning 503
 - 15.5 In a Nutshell 504
- 16 Generative AI and Large Language Models 507**
 - Katherine Munro, Gerald Hahn, Danko Nikolić*
 - 16.1 Introduction to “Gen AI” 507
 - 16.2 Generative AI Modalities 508
 - 16.2.1 Methods for Training Generative Models 510
 - 16.3 Large Language Models 510
 - 16.3.1 What are “LLMs”? 510
 - 16.3.2 How is Something like ChatGPT Trained? 512
 - 16.3.3 Methods for Using LLMs Directly 514
 - 16.3.4 Methods for Customizing an LLM 525
 - 16.4 Vulnerabilities and Limitations of Gen AI Models 534
 - 16.4.1 Introduction 534
 - 16.4.2 Prompt Injection and Jailbreaking Attacks 535
 - 16.4.3 Hallucinations, Confabulations, and Reasoning Errors 538
 - 16.4.4 Copyright Concerns 540
 - 16.4.5 Bias 544
 - 16.5 Building Robust, Effective Gen AI Applications 546
 - 16.5.1 Control Strategies Throughout Development and Use 546
 - 16.5.2 Guardrails 548
 - 16.5.3 Using Generative AI Safely and Successfully 549
 - 16.6 In a Nutshell 550

17	AI Agents	555
	<i>Stefan Papp, Zoltan Toth</i>	
17.1	The Case for AI Agents	555
17.2	Agentic Design Patterns	557
17.2.1	Reflection	558
17.2.2	Tool Use	559
17.2.3	Planning	560
17.2.4	Multi-Agent Pattern	561
17.2.5	Scaling the Complexity	563
17.3	Implementing Agents	564
17.3.1	Frameworks and Platforms	565
17.3.2	Augmented Data Generation	567
17.3.3	Perspectives	570
17.4	Case Study	571
17.4.1	A Simple Agent Using OpenAI Agents	571
17.4.2	Tool Calling	572
17.4.3	MCP in OpenAI Agents	572
17.4.4	Multiple Agents	572
17.5	The Future of Agents	573
17.6	Summary	574
18	Vibe Coding	575
	<i>Stefan Papp</i>	
18.1	What Is Vibe Coding	575
18.2	Vibe Coding Use Cases	577
18.3	Tooling	579
18.3.1	Cursor	580
18.3.2	Claude Code	583
18.3.3	Bolt	586
18.3.4	What Shall I Use	587
18.4	Our Task	588
18.4.1	Solving with Cursor	589
18.4.2	Google Colab	595
18.5	Consequences	598
18.6	Summary	599

19	Natural Language Processing (NLP)	601
	<i>Katherine Munro</i>	
19.1	What is NLP and Why is it so Valuable?	601
19.2	Why Learn “Traditional” NLP in the “Age of Large Language Models”?...	603
19.3	NLP Data Preparation Techniques	605
19.3.1	The NLP Pipeline	605
19.3.2	Converting the Input Format for Machine Learning	612
19.4	NLP Tasks and Methods.....	614
19.4.1	Rule-Based (Symbolic) NLP	615
19.4.2	Statistical Machine Learning Approaches.....	618
19.4.3	Neural NLP	627
19.4.4	Transfer Learning	633
19.5	In a Nutshell.....	646
20	Computer Vision	649
	<i>Roxane Licandro</i>	
20.1	What is Computer Vision?.....	649
20.2	A Picture Paints a Thousand Words	651
20.2.1	The Human Eye	651
20.2.2	Image Acquisition Principle	653
20.2.3	Digital File Formats	659
20.2.4	Image Compression.....	661
20.3	I Spy With My Little Eye Something That Is	662
20.3.1	Computational Photography and Image Manipulation	665
20.4	Computer Vision Applications & Future Directions	668
20.4.1	Image Retrieval Systems	669
20.4.2	Object Detection, Classification and Tracking.....	672
20.4.3	Medical Computer Vision.....	673
20.5	Making Humans See.....	677
20.6	In a Nutshell.....	679
21	Modelling and Simulation – Create your own Models	681
	<i>Günther Zauner, Wolfgang Weidinger, Dominik Brunmeir, Benedikt Spiegel</i>	
21.1	Introduction	682
21.2	General Considerations during Modeling.....	683
21.3	Modelling to Answer Questions.....	684

21.4	Reproducibility and Model Lifecycle	686
21.4.1	The Lifecycle of a Modelling and Simulation Question	687
21.4.2	Parameter and Output Definition	689
21.4.3	Documentation	692
21.4.4	Verification and Validation	693
21.5	Methods	697
21.5.1	Ordinary Differential Equations (ODEs)	698
21.5.2	System Dynamics (SD)	699
21.5.3	Discrete Event Simulation	702
21.5.4	Agent-based Modelling	706
21.6	Modelling and Simulation Examples	709
21.6.1	Dynamic Modelling of Railway Networks for Optimal Pathfinding Using Agent-based Methods and Reinforcement Learning	709
21.6.2	Agent-Based Covid Modelling Strategies	712
21.6.3	Deep Reinforcement Learning Approach for Optimal Replenishment Policy in a VMI Setting	718
21.6.4	Finding Feasible Solutions for a Resource-constrained Project Scheduling Problem with Reinforcement Learning and Implementing a Dynamic Planing Scheme with Discrete Event Simulation	720
21.7	Summary and Lessons Learned	725
21.8	In a Nutshell	726
22	Data Visualization	729
	<i>Barbora Antosova Vesela</i>	
22.1	History	730
22.2	Which Tools to Use	736
22.3	Types of Data Visualizations	739
22.3.1	Scatter Plot	739
22.3.2	Line Chart	740
22.3.3	Column and Bar Charts	741
22.3.4	Histogram	742
22.3.5	Pie Chart	743
22.3.6	Box Plot	743
22.3.7	Heat Map	744
22.3.8	Tree Diagram	745
22.3.9	Other Types of Visualizations	746

22.4	Select the right Data Visualization	746
22.5	Tips and Tricks.....	749
22.6	Presentation of Data Visualization	754
22.7	In a Nutshell.....	755
23	Data Driven Enterprises	757
	<i>Mario Meir-Huber, Stefan Papp</i>	
23.1	The three Levels of a Data Driven Enterprise	758
23.2	Culture.....	758
	23.2.1 Corporate Strategy for Data.....	759
	23.2.2 The Current State Analysis	762
	23.2.3 Culture and Organization of a Successful Data Organisation.....	764
	23.2.4 Core Problem: The Skills Gap	772
23.3	Technology	774
	23.3.1 The Impact of Open Source	774
	23.3.2 Cloud.....	775
	23.3.3 Vendor Selection	775
	23.3.4 Data Lake from a Business Perspective	776
	23.3.5 The Role of IT	777
	23.3.6 Data Science Labs	777
	23.3.7 Revolution in Architecture: The Data Mesh	778
23.4	Business.....	780
	23.4.1 Buy and Share Data.....	780
	23.4.2 Analytical Use Case Implementation	781
	23.4.3 Self-service Analytics	782
23.5	In a Nutshell.....	782
24	Creating High-Performing Teams	783
	<i>Stefan Papp</i>	
24.1	Forming.....	784
24.2	Storming	784
	24.2.1 Scenario: 50 Shades of Red	784
	24.2.2 Scenario: Retrospective	788
24.3	Norming	791
	24.3.1 Change Management and Transition	791
	24.3.2 RACI Matrix.....	794

- 24.3.3 SMART 795
- 24.3.4 Agile Processes 797
- 24.3.5 Communication Culture 799
- 24.3.6 DataOps 800
- 24.4 Performing 805
 - 24.4.1 Scenario: A new Dawn 805
 - 24.4.2 Growth Mindsets 806
- 24.5 In a Nutshell 809
- 25 Artificial Intelligence Act 811**
Jeannette Gorzala, Karin Bruckmüller
- 25.1 Introduction 812
- 25.2 Definition of AI Systems 814
- 25.3 Scope and Purpose of the AI Act 815
 - 25.3.1 The Risk-Based Approach 817
 - 25.3.2 Unacceptable Risk and Prohibited AI Practices 818
 - 25.3.3 High-Risk AI Systems and Compliance 820
 - 25.3.4 Medium Risk and Transparency Obligations 822
 - 25.3.5 Minimal Risk and Voluntary Commitments 823
- 25.4 General Purpose AI Models 824
- 25.5 Timeline and Applicability 826
- 25.6 Penalties 827
- 25.7 AI and Civil Liability 827
- 25.8 AI and Criminal Liability 828
- 25.9 In a Nutshell 831
- 26 AI in Different Industries 833**
Stefan Papp, Mario Meir-Huber, Wolfgang Weidinger, Thomas Treml
- 26.1 Automotive 837
 - 26.1.1 Vision 838
 - 26.1.2 Data 839
 - 26.1.3 Use Cases 839
 - 26.1.4 Challenges 841
- 26.2 Aviation 842
 - 26.2.1 Vision 842
 - 26.2.2 Data 843

26.2.3	Use Cases	843
26.2.4	Challenges	844
26.3	Energy	845
26.3.1	Vision	845
26.3.2	Data	846
26.3.3	Use Cases	847
26.3.4	Challenges	848
26.4	Finance	848
26.4.1	Vision	849
26.4.2	Data	849
26.4.3	Use Cases	849
26.4.4	Challenges	851
26.5	Health	852
26.5.1	Vision	852
26.5.2	Data	853
26.5.3	Use Cases	853
26.5.4	Challenges	854
26.6	Government	855
26.6.1	Vision	855
26.6.2	Data	855
26.6.3	Use Cases	856
26.6.4	Challenges	859
26.7	Art	860
26.7.1	Vision	860
26.7.2	Data	861
26.7.3	Use cases	861
26.7.4	Challenges	862
26.8	Manufacturing	862
26.8.1	Vision	862
26.8.2	Data	863
26.8.3	Use Cases	863
26.8.4	Challenges	864
26.9	Oil and Gas	864
26.9.1	Vision	865
26.9.2	Data	865

- 26.9.3 Use Cases 866
- 26.9.4 Challenges 868
- 26.10 Retail 868
 - 26.10.1 Vision 868
 - 26.10.2 Data 869
 - 26.10.3 Use Cases 869
 - 26.10.4 Challenges 870
- 26.11 Telecommunications Provider 871
 - 26.11.1 Vision 871
 - 26.11.2 Data 872
 - 26.11.3 Use Cases 872
 - 26.11.4 Challenges 874
- 26.12 Transport 874
 - 26.12.1 Vision 874
 - 26.12.2 Data 875
 - 26.12.3 Use Cases 875
 - 26.12.4 Challenges 876
- 26.13 Teaching and Training 876
 - 26.13.1 Vision 877
 - 26.13.2 Data 878
 - 26.13.3 Use Cases 878
 - 26.13.4 Challenges 879
- 26.14 The Digital Society 879
- 26.15 In a Nutshell 882

- 27 Climate Change and AI 883**
Stefan Papp
 - 27.1 Introduction 883
 - 27.2 AI – a Climate Saver? 885
 - 27.3 Measuring and Reducing Emissions 886
 - 27.3.1 Baseline 886
 - 27.3.2 Data Use Cases 888
 - 27.4 Sequestration 889
 - 27.4.1 Biological Sequestration 891
 - 27.4.2 Geological Sequestration 892

27.5	Prepare for Impact	894
27.6	Geoengineering	895
27.7	Greenwashing	896
27.8	Outlook	898
27.9	In a Nutshell	900
28	Mindset and Community	901
	<i>Stefan Papp</i>	
28.1	Data-Driven Mindset	902
28.2	Data Science Culture	904
	28.2.1 Start-up or Consulting Firm?	904
	28.2.2 Labs Instead of Corporate Policy	905
	28.2.3 Keiretsu Instead of Lone Wolf	906
	28.2.4 Agile Software Development	907
	28.2.5 Company and Work Culture	908
28.3	Antipatterns	911
	28.3.1 Devaluation of Domain Expertise	911
	28.3.2 IT Will Take Care of It	912
	28.3.3 Resistance to Change	913
	28.3.4 Know-it-all Mentality	913
	28.3.5 Doom and Gloom	914
	28.3.6 Penny-pinching	915
	28.3.7 Fear Culture	915
	28.3.8 Control over Resources	916
	28.3.9 Blind Faith in Resources	916
	28.3.10 The Swiss Army Knife	917
	28.3.11 Over-Engineering	918
28.4	In a Nutshell	918
29	Trustworthy AI	919
	<i>Rania Wazir</i>	
29.1	Legal and Soft-Law Framework	920
	29.1.1 Standards	922
	29.1.2 Regulations	923
29.2	AI Stakeholders	925

- 29.3 Fairness in AI 926
 - 29.3.1 Bias 928
 - 29.3.2 Fairness Metrics 930
 - 29.3.3 Mitigating Unwanted Bias in AI Systems 934
- 29.4 Transparency of AI Systems 935
 - 29.4.1 Documenting the Data 937
 - 29.4.2 Documenting the Model 938
 - 29.4.3 Explainability 939
- 29.5 Conclusion 941
- 29.6 In a Nutshell 942
- 30 Epilogue 943**
 - Stefan Papp*
 - 30.1 Halford 2.0 943
 - 30.1.1 Environmental, Social and Governance 944
 - 30.1.2 HR 945
 - 30.1.3 Customer Satisfaction 946
 - 30.1.4 Production 948
 - 30.1.5 IT 949
 - 30.1.6 Strategy 951
 - 30.2 Final Words 953
 - 30.3 In a Nutshell 953
- 31 The Authors 955**
- Index 963**