



Contents

Preface ..... *v*

Acknowledgments ..... *vii*

PART I Modern Foundations of Artificial Intelligence

Chapter 1 Evolution of Modern Artificial Intelligence **3**

1.1 From Classical Artificial Intelligence to Modern Artificial Intelligence 3
1.1.1 Characteristics of Classical Artificial Intelligence 4
1.1.2 Limitations of Classical Artificial Intelligence 4
1.1.3 Transition to Modern Artificial Intelligence 5
1.2 Limitations of Symbolic and Rule-Based Approaches 6
1.2.1 Knowledge Acquisition Bottleneck 6
1.2.2 Scalability Issues 7
1.2.3 Lack of Adaptability 7
1.2.4 Difficulty with Uncertainty and Incomplete Information 8
1.3 Rise of Data-Driven Intelligence 9
1.3.1 Role of Data 9
1.3.2 Data as Knowledge 10
1.4 Emergence of Machine Learning and Deep Learning 11
1.4.1 Machine Learning 11
1.4.2 Deep Learning 12
1.4.3 Key Differences 12
1.5 Artificial Intelligence as an End-to-End System 14
1.5.1 Components of an Artificial Intelligence System 14
1.5.2 System-Level Thinking 15
1.6 Contemporary Scope of Advanced Artificial Intelligence 15
1.6.1 Application Areas 15
1.6.2 Emerging Capabilities 16

1.7 Artificial Intelligence Engineering Mindset for Real-World Systems 16
1.7.1 Key Principles 16
1.7.2 From Models to Products 17
Student Assignment 18

Chapter 2 Data-Centric Artificial Intelligence 20

2.1 Importance of Data in Modern Artificial Intelligence 20
2.1.1 Why Data Matters 21
2.2 Types of Data in Artificial Intelligence 22
2.2.1 Structured Data 22
2.2.2 Semi-Structured Data 23
2.2.3 Unstructured Data 24
2.2.4 Streaming and Sensor Data 25
2.2.5 Integration of Multiple Data Types 26
2.3 Data Collection and Annotation 27
2.3.1 Data Collection Sources 27
2.3.2 Data Annotation 28
2.4 Data Preprocessing and Cleaning 30
2.4.1 Key Preprocessing Steps 30
2.5 Feature Representation: A Conceptual View 31
2.5.1 What is a Feature? 31
2.5.2 Types of Features 31
2.5.3 Importance of Feature Representation 32
2.6 Data Quality, Robustness, and Reliability 32
2.6.1 Dimensions of Data Quality 33
2.6.2 Robustness 33
2.6.3 Reliability 33
2.7 Data Bias, Fairness, and Representation 34
2.7.1 Types of Bias 34
2.7.2 Fairness in Artificial Intelligence 34
2.7.3 Representation in Data 35
Student Assignment 35

Chapter 3 Machine Learning Systems and Lifecycle 38

3.1 Role of Machine Learning in Advanced Artificial Intelligence 38
3.1.1 Machine Learning as the Core of Advanced Artificial Intelligence 39
3.1.2 Why Machine Learning is Essential in Advanced Artificial Intelligence 39
3.1.3 From Data to Intelligence 40
3.2 Learning Paradigms: A Brief Engineering View 41
3.2.1 Supervised Learning 42
3.2.2 Unsupervised Learning 44
3.2.3 Semi-Supervised Learning 47
3.2.4 Reinforcement Learning 49

3.3 Training and Inference 50
3.3.1 Training Phase 50
3.3.2 Inference Phase 51
3.3.3 Spam Classifier: Example of Training and Inference 51
3.4 Generalization, Overfitting, and Underfitting 51
3.4.1 Overfitting 51
3.4.2 Underfitting 51
3.4.3 Bias–Variance Trade-off 51
3.5 Model Evaluation Metrics 52
3.5.1 Classification Metrics 52
3.5.2 Regression Metrics 54
3.6 Validation, Selection, and Benchmarking 55
3.6.1 Validation 55
3.6.2 Model Selection 55
3.6.3 Benchmarking 56
3.7 Deployment-Oriented Machine Learning 56
3.7.1 Key Considerations 56
3.8 Drifts in Machine Learning 57
3.8.1 Model Drift 57
3.9 Introduction to Machine Learning Operations and Lifecycle Management 60
3.9.1 Core Principles of Machine Learning Operations 60
3.9.2 Lifecycle Stages 61
3.9.3 Machine Learning Operations Pipeline 61
3.9.4 Goals of Machine Learning Operations 62
Student Assignment 63

PART II Deep Learning and Language Intelligence

Chapter 4 Deep Learning Fundamentals 67
4.1 Why Deep Learning Matters 67
4.1.1 Key Motivations 68
4.2 From Perceptron to Deep Networks 68
4.2.1 The Perceptron: The Fundamental Computational Unit 69
4.2.2 Single Layer Perceptron 69
4.2.3 Multilayer Perceptron 70
4.2.4 From Multilayer Perceptron to Deep Networks 72
4.3 The Need for Backpropagation 72
4.4 Artificial Neural Networks as Function Approximators 73
4.4.1 Mathematical Representation 73
4.4.2 Understanding Function Approximation 73
4.4.3 Layers, Parameters, and Activation Functions 74

4.4.4 Forward Computation and Learning Intuition 75
4.4.5 Optimization in Deep Networks 78
4.4.6 Learning as Function Approximation 79
4.5 Challenges in Deep Learning 80
4.5.1 Vanishing and Exploding Gradients 80
4.5.2 Overfitting and Regularization 80
4.5.3 Data and Compute Bottlenecks 81
4.6. Hardware Acceleration for Deep Learning 82
4.6.1 Hardware Types 82
4.6.2 Why Graphics Processing Unit? 82
4.7 Edge and Embedded Deep Learning 82
4.7.1 Characteristics 83
4.7.2 Techniques for Edge Deployment 83
Student Assignment 83

Chapter 5 Deep Sequential and Transformer Architectures 86

5.1 Evolution of Sequence Modeling 86
5.1.1 Classical Approaches 88
5.2 Word Representations and Embeddings 92
5.2.1 One-Hot Encoding 92
5.2.2 Dense Embeddings 92
5.3 Sequence Modeling with Neural Networks 93
5.3.1 Recurrent Neural Networks 93
5.3.2 Long Short-Term Memory Networks 98
5.4 Attention Mechanism and Context Modeling 104
5.5 Transformer Architecture 106
5.5.1 Advantages 106
5.5.2 Self-Attention 106
5.5.3 Positional Encoding 107
5.5.4 Encoder and Decoder Structures 110
5.6 Large Language Models 112
5.7 Capabilities, Limitations, and Engineering Implications 113
5.7.1 Capabilities 113
5.7.2 Limitations 113
5.7.3 Engineering Implications 113
Student Assignment 113

Chapter 6 Natural Language Processing in the Era of Large Language Models 116

6.1 Language as Data, Knowledge, and Interface 117
6.1.1 Language as Data 117
6.1.2 Language as Knowledge 117
6.1.3 Language as Interface 118
6.2 Modern Natural Language Processing Pipeline 118

6.3 Text Representation in Contemporary Natural Language Processing 119
6.3.1 Dense Embeddings 120
6.3.2 Contextual Representations 122
6.4 Core Natural Language Processing Tasks in Modern Systems 125
6.4.1 Text Generation and Content Creation 125
6.4.2 Information Extraction 126
6.4.3 Summarization 126
6.4.4 Retrieval-Augmented Generation 127
6.5 Conversational Artificial Intelligence and Dialogue Systems 128
6.6 Prompting and Human–Model Interaction 129
6.7 Evaluation of Natural Language Processing Systems 130
6.8 Ethical and Societal Issues in Natural Language Processing 132
6.8.1 Key Ethical Concerns 132
6.8.2 Broader Societal Implications 133
6.8.3 Approaches for Responsible Natural Language Processing 133
Student Assignment 133

PART III Perception and Intelligent Autonomous Systems

Chapter 7 Computer Vision 139

7.1 Vision as an Artificial Intelligence Problem 139
7.1.1 Why Vision is Challenging 140
7.2 Digital Images as Data 141
7.2.1 Pixel Representation 141
7.2.2 Intensity and Channels 141
7.2.3 Spatial Structure and Sensitivity to Transformations 141
7.3 Visual Representation and Feature Learning 142
7.3.1 Traditional Feature Engineering 142
7.3.2 Learning-Based Features 144
7.4 Architecture of Convolutional Neural Networks 147
7.4.1 Architectural Overview 147
7.4.2 Core Components 148
7.5 Vision Transformer 152
7.5.1 Patch-Based Image Representation 152
7.5.2 Class Token and Positional Encoding 153
7.5.3 Transformer Encoder 153
7.5.4 Classification Head 153
7.5.5 Key Characteristics 153
7.6 Image Classification 154
7.6.1 Problem Definition 154
7.6.2 Mathematical Interpretation 154

7.7 Object Detection 155
7.7.1 Detection as a Regression Problem 155
7.7.2 Bounding Box Representation 156
7.7.3 Class Probability Prediction 156
7.7.4 Postprocessing Nonmaximum Suppression 156
7.7.5 Recognition versus Detection 156
7.8 Image Segmentation 156
7.8.1 Problem Definition 156
7.8.2 Types of Image Segmentation 157
7.8.3 Mathematical Formulation 157
7.8.4 Learning Objective 157
7.9 Face Detection 158
7.9.1 Problem Definition 158
7.9.2 Detection Pipeline 159
7.9.3 Mathematical Formulation 159
7.10 Face Landmark Detection 159
7.10.1 Problem Definition 159
7.10.2 Geometric Interpretation 159
7.10.3 Learning Objective 160
7.11 Pose Landmark Detection 160
7.11.1 Problem Definition 160
7.11.2 Geometric Interpretation 161
7.11.3 Mathematical Formulation 161
Student Assignment 162

Chapter 8 Robotics and Autonomous Systems 165

8.1 Role of Artificial Intelligence in Robotics 165
8.2 Perception–Planning–Action Cycle 166
8.2.1 Perception 167
8.2.2 Planning 167
8.2.3 Action 167
8.3 Sensors and Actuators 168
8.3.1 Sensors 168
8.3.2 Actuators 169
8.4 Environmental Awareness and Localization 169
8.4.1 Mapping 169
8.4.2 Localization 171
8.4.3 Simultaneous Localization and Mapping 172
8.5 Decision-Making in Autonomous Systems 173
8.5.1 Rule-Based Systems 173
8.5.2 Probabilistic Decision-Making 173
8.5.3 Learning-Based Decision-Making 174

8.6 Industrial Robotics and Intelligent Automation 175
8.6.1 Traditional Automation 175
8.6.2 Artificial Intelligence-Driven Automation 176
8.6.3 Applications 176
8.7 Human–Robot Interaction 177
8.7.1 Modes of Interaction 177
8.7.2 Safety and Trust 177
8.8 Emerging Paradigms in Robotics 178
8.8.1 Cloud and Fog Robotics 179
8.8.2 Cognitive Robotics 179
8.8.3 Multiagent Robotic Systems 179
8.8.4 Telerobotics 180
8.8.5 Self-Reconfiguring Modular Robots 181
8.9 Challenges in Robotic Intelligence 182
8.9.1 Challenges in Self-Reconfiguring Robots 183
8.9.2 Ethical and Governance Challenges 183
Student Assignment 183

PART IV Optimization and Soft Computing

Chapter 9 Evolutionary Algorithms 189
9.1 Biological Inspiration and Evolutionary Thinking 189
9.2 Search and Optimization in Complex Spaces 190
9.3 Genetic Algorithms: Representation and Workflow 191
9.3.1 Basic Workflow 192
9.3.2 Workflow of a Genetic Algorithm 192
9.4 Population Initialization and Fitness Function 193
9.4.1 Population Initialization 193
9.4.2 Fitness Function 194
9.5 Selection Mechanisms 195
9.6 Crossover and Recombination 197
9.7 Mutation and Diversity Maintenance 199
9.8 Elitism, Convergence, and Termination Criteria 200
9.8.1 Elitism 200
9.8.2 Convergence 201
9.8.3 Termination Criteria 201
9.9 Engineering Applications of the Evolutionary Algorithms 203
9.10 Strengths and Limitations of Evolutionary Methods 204
9.10.1 Strengths 204
9.10.2 Limitations 204
Student Assignment 205

Chapter 10 Fuzzy Logic and Intelligent Decision Systems 207

10.1 Need for Approximate Reasoning 207
10.2 Classical Sets and Fuzzy Sets 208
10.3 Membership Functions 210
10.3.1 Triangular Membership Function 210
10.3.2 Trapezoidal Membership Function 211
10.3.3 Gaussian Membership Function 212
10.3.4 Sigmoidal Membership Function 212
10.3.5 Comparison and Selection 213
10.4 Operations on Fuzzy Sets 213
10.4.1 Standard Fuzzy Set Operations 214
10.4.2 Generalized Framework: t-norms, t-conorms, and Negators 216
10.4.3 Fuzzy Complements: Mathematical Properties 216
10.4.4 Fuzzy Intersections (t-norms) 217
10.4.5 Fuzzy Unions (t-conorms) 217
10.4.6 Aggregation Operations 218
10.4.7 Additional Useful Operations 218
10.5 Fuzzy Relations and Fuzzy Rules 219
10.5.1 Fuzzy Relations 219
10.5.2 Fuzzy Rules 219
10.6 Fuzzy Inference Systems 220
10.6.1 Mamdani Inference 220
10.6.2 Sugeno Inference 220
10.6.3 Comparison of Mamdani and Sugeno 221
10.7 Defuzzification Techniques 221
10.7.1 Need for Defuzzification 222
10.7.2 Common Defuzzification Methods 222
10.7.3 Fuzzy Control Systems 224
10.7.4 Decision Making with Fuzzy Logic 225
10.7.5 Applications of Fuzzy Systems 225
Student Assignment 226

PART V Generative, Applied, and Responsible Artificial Intelligence

Chapter 11 Generative Artificial Intelligence 231

11.1 Foundations of Generative Artificial Intelligence 232
11.1.1 What is Generative Modeling? 232
11.1.2 Types of Generative Models 233
11.2 Discriminative and Generative Modeling 234
11.2.1 Discriminative Models 234

11.2.2 Generative Models 235
11.2.3 Comparison 236
11.3 Generative Adversarial Networks 237
11.3.1 Core Components of a Generative Adversarial Network 237
11.3.2 Adversarial Learning Objective 238
11.3.3 How Generative Adversarial Network Training Works 238
11.3.4 Practical Strengths of Generative Adversarial Networks 238
11.3.5 Training Challenges and Limitations 239
11.4 Variational Autoencoders 239
11.4.1 Core Idea 240
11.4.2 Architecture Components 240
11.4.3 Probabilistic Formulation 240
11.4.4 Evidence Lower Bound 240
11.4.5 Reparameterization Trick 240
11.4.6 Working Principle 241
11.4.7 Applications of Variational Autoencoders 241
11.5 Diffusion Models 242
11.5.1 Forward Process 242
11.5.2 Reverse Process 243
11.5.3 Noise Prediction Perspective 243
11.5.4 Training Objective 243
11.5.5 Advantages of Diffusion Models 244
11.5.6 Limitations 244
11.6 Text Generation with Large Language Models 245
11.6.1 Autoregressive Modeling 245
11.6.2 Transformer Architecture: Conceptual View 245
11.6.3 Text Generation Process and Capabilities 246
11.7 Applications of Generative Artificial Intelligence 247
11.7.1 Content Creation 247
11.7.2 Healthcare and Drug Discovery 247
11.7.3 Education 247
11.7.4 Design and Art 247
11.7.5 Simulation and Gaming 247
11.8 Risks, Limitations, and Failure Modes 248
11.8.1 Key Risks 248
11.8.2 Failure Modes 248
Student Assignment 249

Chapter 12 Artificial Intelligence Platforms, Tools, and Deployment Ecosystems 252

12.1 Artificial Intelligence Development Ecosystem 253
12.1.1 Core Components of the Artificial Intelligence Ecosystem 253
12.1.2 Role of Frameworks and Libraries 253

12.2 Data Science and Model Development Platforms ... 256
12.2.1 Notebook-Based Environments ... 256
12.2.2 Experiment Tracking and Reproducibility ... 257
12.3 Cloud-Based Artificial Intelligence Services ... 257
12.3.1 Key Features of Cloud Artificial Intelligence ... 258
12.3.2 Major Artificial Intelligence Cloud Platforms ... 258
12.4 No-Code and Low-Code Artificial Intelligence Platforms ... 259
12.4.1 What is No-Code Artificial Intelligence? ... 259
12.4.2 Motivation and Importance ... 259
12.4.3 Key Advantages ... 260
12.4.4 Limitations and Challenges ... 260
12.4.5 Popular No-Code Artificial Intelligence Tools ... 260
12.5 Generative Artificial Intelligence Tools and Application Programming Interfaces ... 261
12.5.1 Capabilities ... 261
12.5.2 Generative Artificial Intelligence Application Programming Interfaces and Platforms ... 262
12.5.3 Role of Prompts in Generative Artificial Intelligence ... 262
12.6 Edge Artificial Intelligence and On-Device Inference Tools ... 263
12.6.1 Motivation for Edge Artificial Intelligence ... 263
12.6.2 Challenges in Edge Artificial Intelligence ... 263
12.6.3 Model Optimization Techniques for Edge Artificial Intelligence ... 263
12.6.4 Edge Artificial Intelligence Frameworks and Tools ... 265
12.6.5 Local Large Language Model and Offline Edge Artificial Intelligence Tools ... 266
12.7 Model Integration and Service Architectures ... 267
12.7.1 Microservices Architecture ... 267
12.7.2 Microservice-Based Artificial Intelligence Flow Architecture ... 267
12.8 Real-Time and Distributed Artificial Intelligence Systems ... 268
12.8.1 Real-Time Processing ... 268
12.8.2 Distributed Systems ... 268
12.9 Deployment Challenges in Practice ... 269
12.9.1 Common Challenges ... 269
12.9.2 Engineering Considerations ... 269
Student Assignment ... 270

Chapter 13 Responsible and Explainable Artificial Intelligence ... 273

13.1 Why Responsible Artificial Intelligence Matters ... 274
13.1.1 Artificial Intelligence in High-Impact Decision Making ... 274
13.1.2 Societal and Organizational Impact ... 274
13.1.3 Framework of Responsible Artificial Intelligence ... 274
13.1.4 Business Case for Responsible Artificial Intelligence ... 276
13.1.5 Need for Ethical and Trustworthy Artificial Intelligence ... 276
13.1.6 From Performance to Responsibility ... 277

13.2 Bias, Fairness, and Inclusivity ... 277
13.2.1 Sources and Types of Bias ... 277
13.2.2 Fairness Definitions ... 278
13.2.3 Inclusivity in Artificial Intelligence ... 278
13.3 Privacy and Security in Artificial Intelligence Systems ... 279
13.3.1 Privacy Concerns ... 279
13.3.2 Security Threats ... 280
13.4 Explainable Artificial Intelligence ... 280
13.4.1 Interpretability and Explainability ... 281
13.4.2 Model Transparency ... 281
13.4.3 Local and Global Explanations ... 281
13.4.4 Techniques for Explainability ... 282
13.4.5 Trust and Human Understanding ... 284
13.5 Trustworthy Artificial Intelligence ... 285
13.5.1 Core Principles ... 285
13.6 Human-in-the-Loop Systems ... 286
13.6.1 Role of Humans ... 286
13.7 Artificial Intelligence Safety and Risk Mitigation ... 288
13.7.1 Mitigation Strategies ... 288
13.8 Governance, Regulation, and Standards ... 290
13.8.1 Need for Regulation ... 290
13.8.2 Major Regulatory and Policy Frameworks ... 290
13.8.3 Global Standards ... 291
13.8.4 Organizational Governance ... 292
Student Assignment ... 292

Chapter 14 Emerging Paradigms in Artificial Intelligence ... 295

14.1 Foundation Models ... 296
14.1.1 Concept of Foundation Models ... 296
14.1.2 Training Paradigm ... 296
14.1.3 Adaptation Mechanisms ... 296
14.1.4 Types of Foundation Models ... 297
14.1.5 Examples of Foundation Models ... 297
14.2 Scaling Laws in Artificial Intelligence ... 297
14.2.1 A Simple Power-Law View ... 298
14.2.2 Scaling with Model Size ... 298
14.2.3 Scaling with Dataset Size ... 299
14.2.4 Scaling with Compute Budget ... 299
14.2.5 Joint Scaling of Parameters and Data ... 299
14.2.6 Optimal Scaling Under Fixed Compute ... 300
14.2.7 Meaning of Power-Law Behavior ... 300

14.3 Multimodal Artificial Intelligence 300
14.3.1 Need for Multimodal Systems 300
14.3.2 Representation of Multimodal Learning 301
14.3.3 Integration Mechanisms 301
14.3.4 Challenges in Multimodal Learning 302
14.4 Edge Artificial Intelligence and On-Device Intelligence 302
14.4.1 Motivation 302
14.4.2 System Model of Edge Artificial Intelligence 302
14.4.3 Trade-Off between Accuracy and Efficiency 303
14.5 Neuromorphic and Brain-Inspired Computing 303
14.5.1 Biological Inspiration and Computational Model 304
14.5.2 Core Principles 304
14.5.3 Advantages 305
14.5.4 Neuromorphic Hardware and Sensors 306
14.6 Quantum Perspectives in Artificial Intelligence 308
14.6.1 Quantum Concepts in Artificial Intelligence 308
14.6.2 Quantum Machine Learning Framework 309
14.6.3 Potential Applications 310
14.7 Autonomous Agents and Agentic Systems 310
14.7.1 Characteristics of Autonomous Agents 310
14.7.2 Agent Architectures 311
14.7.3 Agentic Artificial Intelligence in the Current Scenario 311
14.8 Open Challenges in Advanced Artificial Intelligence 312
14.8.1 Technical Challenges 312
14.8.2 Societal Challenges 313
14.9 Future Skills for Artificial Intelligence Engineers 313
14.9.1 Core Technical Skills 313
14.9.2 Emerging Skill Areas 314
Student Assignment 314

Index *317*