



GLOBAL EDITION



Stuart
Russell
Peter
Norvig

Artificial Intelligence

A Modern Approach

Fourth Edition



Artificial Intelligence

A Modern Approach

Fourth Edition

Global Edition

Artificial Intelligence: A Modern Approach, Global Edition

Table of Contents

Cover

Half Title

AI Pearson Series in Artificial Intelligence

Title Page

Copyright

Dedication

Preface

About the Authors

Contents

I: Artificial Intelligence

Chapter 1: Introduction

1.1 What Is AI?

1.2 The Foundations of Artificial Intelligence

1.3 The History of Artificial Intelligence

1.4 The State of the Art

1.5 Risks and Benefits of AI

Summary

Bibliographical and Historical Notes

Chapter 2: Intelligent Agents

2.1 Agents and Environments

2.2 Good Behavior: The Concept of Rationality

2.3 The Nature of Environments

2.4 The Structure of Agents

Summary

Bibliographical and Historical Notes

Table of Contents

II: Problem-solving

Chapter 3: Solving Problems by Searching

- 3.1 Problem-Solving Agents
- 3.2 Example Problems
- 3.3 Search Algorithms
- 3.4 Uninformed Search Strategies
- 3.5 Informed (Heuristic) Search Strategies
- 3.6 Heuristic Functions

Summary

Bibliographical and Historical Notes

Chapter 4: Search in Complex Environments

- 4.1 Local Search and Optimization Problems
- 4.2 Local Search in Continuous Spaces
- 4.3 Search with Nondeterministic Actions
- 4.4 Search in Partially Observable Environments
- 4.5 Online Search Agents and Unknown Environments

Summary

Bibliographical and Historical Notes

Chapter 5: Constraint Satisfaction Problems

- 5.1 Defining Constraint Satisfaction Problems
- 5.2 Constraint Propagation: Inference in CSPs
- 5.3 Backtracking Search for CSPs
- 5.4 Local Search for CSPs
- 5.5 The Structure of Problems

Summary

Bibliographical and Historical Notes

Chapter 6: Adversarial Search and Games

- 6.1 Game Theory
- 6.2 Optimal Decisions in Games
- 6.3 Heuristic AlphaBeta Tree Search
- 6.4 Monte Carlo Tree Search
- 6.5 Stochastic Games
- 6.6 Partially Observable Games

Table of Contents

6.7 Limitations of Game Search Algorithms

Summary

Bibliographical and Historical Notes

III: Knowledge, reasoning, and planning

Chapter 7: Logical Agents

7.1 Knowledge-Based Agents

7.2 The Wumpus World

7.3 Logic

7.4 Propositional Logic: A Very Simple Logic

7.5 Propositional Theorem Proving

7.6 Effective Propositional Model Checking

7.7 Agents Based on Propositional Logic

Summary

Bibliographical and Historical Notes

Chapter 8: First-Order Logic

8.1 Representation Revisited

8.2 Syntax and Semantics of First-Order Logic

8.3 Using First-Order Logic

8.4 Knowledge Engineering in First-Order Logic

Summary

Bibliographical and Historical Notes

Chapter 9: Inference in First-Order Logic

9.1 Propositional vs. First-Order Inference

9.2 Unification and First-Order Inference

9.3 Forward Chaining

9.4 Backward Chaining

9.5 Resolution

Summary

Bibliographical and Historical Notes

Chapter 10: Knowledge Representation

10.1 Ontological Engineering

10.2 Categories and Objects

10.3 Events

Table of Contents

10.4 Mental Objects and Modal Logic
10.5 Reasoning Systems for Categories
10.6 Reasoning with Default Information
Summary

Bibliographical and Historical Notes

Chapter 11: Automated Planning

11.1 Definition of Classical Planning
11.2 Algorithms for Classical Planning
11.3 Heuristics for Planning
11.4 Hierarchical Planning
11.5 Planning and Acting in Nondeterministic Domains
11.6 Time, Schedules, and Resources
11.7 Analysis of Planning Approaches
Summary

Bibliographical and Historical Notes

IV: Uncertain knowledge and reasoning

Chapter 12: Quantifying Uncertainty

12.1 Acting under Uncertainty
12.2 Basic Probability Notation
12.3 Inference Using Full Joint Distributions
12.4 Independence
12.5 Bayes Rule and Its Use
12.6 Naive Bayes Models
12.7 The Wumpus World Revisited

Summary

Bibliographical and Historical Notes

Chapter 13: Probabilistic Reasoning

13.1 Representing Knowledge in an Uncertain Domain
13.2 The Semantics of Bayesian Networks
13.3 Exact Inference in Bayesian Networks
13.4 Approximate Inference for Bayesian Networks
13.5 Causal Networks

Summary

Table of Contents

Bibliographical and Historical Notes

Chapter 14: Probabilistic Reasoning over Time

14.1 Time and Uncertainty

14.2 Inference in Temporal Models

14.3 Hidden Markov Models

14.4 Kalman Filters

14.5 Dynamic Bayesian Networks

Summary

Bibliographical and Historical Notes

Chapter 15: Making Simple Decisions

15.1 Combining Beliefs and Desires under Uncertainty

15.2 The Basis of Utility Theory

15.3 Utility Functions

15.4 Multiattribute Utility Functions

15.5 Decision Networks

15.6 The Value of Information

15.7 Unknown Preferences

Summary

Bibliographical and Historical Notes

Chapter 16: Making Complex Decisions

16.1 Sequential Decision Problems

16.2 Algorithms for MDPs

16.3 Bandit Problems

16.4 Partially Observable MDPs

16.5 Algorithms for Solving POMDPs

Summary

Bibliographical and Historical Notes

Chapter 17: Multiagent Decision Making

17.1 Properties of Multiagent Environments

17.2 Non-Cooperative Game Theory

17.3 Cooperative Game Theory

17.4 Making Collective Decisions

Summary

Table of Contents

Bibliographical and Historical Notes

Chapter 18: Probabilistic Programming

18.1 Relational Probability Models

18.2 Open-Universe Probability Models

18.3 Keeping Track of a Complex World

18.4 Programs as Probability Models

Summary

Bibliographical and Historical Notes

V: Machine Learning

Chapter 19: Learning from Examples

19.1 Forms of Learning

19.2 Supervised Learning

19.3 Learning Decision Trees

19.4 Model Selection and Optimization

19.5 The Theory of Learning

19.6 Linear Regression and Classification

19.7 Nonparametric Models

19.8 Ensemble Learning

19.9 Developing Machine Learning Systems

Summary

Bibliographical and Historical Notes

Chapter 20: Knowledge in Learning

20.1 A Logical Formulation of Learning

20.2 Knowledge in Learning

20.3 Explanation-Based Learning

20.4 Learning Using Relevance Information

20.5 Inductive Logic Programming

Summary

Bibliographical and Historical Notes

Chapter 21: Learning Probabilistic Models

21.1 Statistical Learning

21.2 Learning with Complete Data

21.3 Learning with Hidden Variables: The EM Algorithm

Table of Contents

Summary

Bibliographical and Historical Notes

Chapter 22: Deep Learning

22.1 Simple Feedforward Networks

22.2 Computation Graphs for Deep Learning

22.3 Convolutional Networks

22.4 Learning Algorithms

22.5 Generalization

22.6 Recurrent Neural Networks

22.7 Unsupervised Learning and Transfer Learning

22.8 Applications

Summary

Bibliographical and Historical Notes

Chapter 23: Reinforcement Learning

23.1 Learning from Rewards

23.2 Passive Reinforcement Learning

23.3 Active Reinforcement Learning

23.4 Generalization in Reinforcement Learning

23.5 Policy Search

23.6 Apprenticeship and Inverse Reinforcement Learning

23.7 Applications of Reinforcement Learning

Summary

Bibliographical and Historical Notes

VI: Communicating, perceiving, and acting

Chapter 24: Natural Language Processing

24.1 Language Models

24.2 Grammar

24.3 Parsing

24.4 Augmented Grammars

24.5 Complications of Real Natural Language

24.6 Natural Language Tasks

Summary

Bibliographical and Historical Notes

Table of Contents

Chapter 25: Deep Learning for Natural Language Processing

- 25.1 Word Embeddings
- 25.2 Recurrent Neural Networks for NLP
- 25.3 Sequence-to-Sequence Models
- 25.4 The Transformer Architecture
- 25.5 Pretraining and Transfer Learning
- 25.6 State of the art
- Summary
- Bibliographical and Historical Notes

Chapter 26: Robotics

- 26.1 Robots
- 26.2 Robot Hardware
- 26.3 What kind of problem is robotics solving?
- 26.4 Robotic Perception
- 26.5 Planning and Control
- 26.6 Planning Uncertain Movements
- 26.7 Reinforcement Learning in Robotics
- 26.8 Humans and Robots
- 26.9 Alternative Robotic Frameworks
- 26.10 Application Domains
- Summary
- Bibliographical and Historical Notes

Chapter 27: Computer Vision

- 27.1 Introduction
- 27.2 Image Formation
- 27.3 Simple Image Features
- 27.4 Classifying Images
- 27.5 Detecting Objects
- 27.6 The 3D World
- 27.7 Using Computer Vision
- Summary
- Bibliographical and Historical Notes

VII: Conclusions

Table of Contents

Chapter 28: Philosophy, Ethics, and Safety of AI

28.1 The Limits of AI

28.2 Can Machines Really Think?

28.3 The Ethics of AI

Summary

Bibliographical and Historical Notes

Chapter 29: The Future of AI

29.1 AI Components

29.2 AI Architectures

Appendixes

Appendix A: Mathematical Background

A.1 Complexity Analysis and $O()$ Notation

A.2 Vectors, Matrices, and Linear Algebra

A.3 Probability Distributions

Bibliographical and Historical Notes

Appendix B: Notes on Languages and Algorithms

B.1 Defining Languages with BackusNaur Form (BNF)

B.2 Describing Algorithms with Pseudocode

B.3 Online Supplemental Material

Bibliography

Index