

<<人工智能>>

图书基本信息

书名：<<人工智能>>

13位ISBN编号：9787302252955

10位ISBN编号：7302252955

出版时间：2011-7

出版时间：清华大学出版社

作者：Stuart J. Russell, Peter Norvig

页数：1132

版权说明：本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问：<http://www.tushu007.com>

<<人工智能>>

内容概要

《人工智能(一种现代的方法第3版影印版)》(作者拉塞尔、诺维格)是最权威、最经典的人工智能教材,已被全世界100多个国家的1200多所大学用作教材。

《人工智能(一种现代的方法第3版影印版)》的最新版全面而系统地介绍了人工智能的理论和实践,阐述了人工智能领域的核心内容,并深入介绍了各个主要的研究方向。全书仍分为八大部分:第一部分“人工智能”,第二部分“问题求解”,第三部分“知识与推理”,第四部分“规划”,第五部分“不确定知识与推理”,第六部分“学习”,第七部分“通信、感知与行动”,第八部分“结论”。

《人工智能(一种现代的方法第3版影印版)》既详细介绍了人工智能的基本概念、思想和算法,还描述了其各个研究方向最前沿的进展,同时收集整理了详实的历史文献与事件。

另外,《人工智能(一种现代的方法第3版影印版)》的配套网址为教师和学生提供了大量教学和学习资料。

《人工智能(一种现代的方法第3版影印版)》适合于不同层次和领域的研究人员及学生,是高等院校本科生和研究生人工智能课的首选教材,也是相关领域的科研与工程技术人员的重要参考书。

<<人工智能>>

作者简介

作者：（美国）拉塞尔（Stuart J.Russell）（美国）诺维格（Peter Norvig）

书籍目录

- I Artificial Intelligence
- 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 Summary, Bibliographical and Historical Notes, Exercises
- 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
 - 2.5 Summary, Bibliographical and Historical Notes, Exercises
- II Problem-solving
- 3 Solving Problems by Searching
 - 3.1 Problem-Solving Agents
 - 3.2 Example Problems
 - 3.3 Searching for Solutions
 - 3.4 Uninformed Search Strategies
 - 3.5 Informed (Heuristic) Search Strategies
 - 3.6 Heuristic Functions
 - 3.7 Summary, Bibliographical and Historical Notes, Exercises
- 4 Beyond Classical Search
 - 4.1 Local Search Algorithms and Optimization Problems
 - 4.2 Local Search in Continuous Spaces
 - 4.3 Searching with Nondeterministic Actions
 - 4.4 Searching with Partial Observations
 - 4.5 Online Search Agents and Unknown Environments
 - 4.6 Summary, Bibliographical and Historical Notes, Exercises
- 5 Adversarial Search
 - 5.1 Games
 - 5.2 Optimal Decisions in Games
 - 5.3 Alpha-Beta Pruning
 - 5.4 Imperfect Real-Time Decisions
 - 5.5 Stochastic Games
 - 5.6 Partially Observable Games
 - 5.7 State-of-the-Art Game Programs
 - 5.8 Alternative Approaches
 - 5.9 Summary, Bibliographical and Historical Notes, Exercises
- 6 Constraint Satisfaction Problems
 - 6.1 Defining Constraint Satisfaction Problems
 - 6.2 Constraint Propagation: Inference in CSPs
 - 6.3 Backtracking Search for CSPs
 - 6.4 Local Search for CSPs
 - 6.5 The Structure of Problems

<<人工智能>>

- 6.6 Summary, Bibliographical and Historical Notes, Exercises
- III Knowledge, reasoning, and planning
- 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
 - 7.8 Summary, Bibliographical and Historical Notes, Exercises
- 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
 - 8.5 Summary, Bibliographical and Historical Notes, Exercises
- 9 Inference in First-Order Logic
 - 9.1 Propositional vs. First-Order Inference
 - 9.2 Unification and Lifting
 - 9.3 Forward Chaining
 - 9.4 Backward Chaining
 - 9.5 Resolution
 - 9.6 Summary, Bibliographical and Historical Notes, Exercises
- 10 Classical Planning
 - 10.1 Definition of Classical Planning
 - 10.2 Algorithms for Planning as State-Space Search
 - 10.3 Planning Graphs
 - 10.4 Other Classical Planning Approaches
 - 10.5 Analysis of Planning Approaches
 - 10.6 Summary, Bibliographical and Historical Notes, Exercises
- 11 Planning and Acting in the Real World
 - 11.1 Time, Schedules, and Resources
 - 11.2 Hierarchical Planning
 - 11.3 Planning and Acting in Nondeterministic Domains
 - 11.4 Multiagent Planning
 - 11.5 Summary, Bibliographical and Historical Notes, Exercises
- 12 Knowledge Representation
 - 12.1 Ontological Engineering
 - 12.2 Categories and Objects
 - 12.3 Events
 - 12.4 Mental Events and Mental Objects
 - 12.5 Reasoning Systems for Categories
 - 12.6 Reasoning with Default Information
 - 12.7 The Internet Shopping World
 - 12.8 Summary, Bibliographical and Historical Notes, Exercises
- IV Uncertain knowledge and reasoning

<<人工智能>>

- 13 Quantifying Uncertainty
 - 13.1 Acting under Uncertainty
 - 13.2 Basic Probability Notation
 - 13.3 Inference Using Full Joint Distributions
 - 13.4 Independence
 - 13.5 Bayes' Rule and Its Use
 - 13.6 The Wumpus World Revisited
 - 13.7 Summary, Bibliographical and Historical Notes, Exercises
- 14 Probabilistic Reasoning
 - 14.1 Representing Knowledge in an Uncertain Domain
 - 14.2 The Semantics of Bayesian Networks
 - 14.3 Efficient Representation of Conditional Distributions
 - 14.4 Exact Inference in Bayesian Networks
 - 14.5 Approximate Inference in Bayesian Networks
 - 14.6 Relational and First-Order Probability Models
 - 14.7 Other Approaches to Uncertain Reasoning
 - 14.8 Summary, Bibliographical and Historical Notes, Exercises
- 15 Probabilistic Reasoning over Time
 - 15.1 Time and Uncertainty
 - 15.2 Inference in Temporal Models
 - 15.3 Hidden Markov Models
 - 15.4 Kalman Filters
 - 15.5 Dynamic Bayesian Networks
 - 15.6 Keeping Track of Many Objects
 - 15.7 Summary, Bibliographical and Historical Notes, Exercises
- 16 Making Simple Decisions
 - 16.1 Combining Beliefs and Desires under Uncertainty
 - 16.2 The Basis of Utility Theory
 - 16.3 Utility Functions
 - 16.4 Multiattribute Utility Functions
 - 16.5 Decision Networks
 - 16.6 The Value of Information
 - 16.7 Decision-Theoretic Expert Systems
 - 16.8 Summary, Bibliographical and Historical Notes, Exercises
- 17 Making Complex Decisions
 - 17.1 Sequential Decision Problems
 - 17.2 Value Iteration
 - 17.3 Policy Iteration
 - 17.4 Partially Observable MDPs
 - 17.5 Decisions with Multiple Agents: Game Theory
 - 17.6 Mechanism Design
 - 17.7 Summary, Bibliographical and Historical Notes, Exercises
- V Learning
 - 18 Learning from Examples
 - 18.1 Forms of Learning
 - 18.2 Supervised Learning
 - 18.3 Learning Decision Trees

<<人工智能>>

- 18.4 Evaluating and Choosing the Best Hypothesis
- 18.5 The Theory of Learning
- 18.6 Regression and Classification with Linear Models
- 18.7 Artificial Neural Networks
- 18.8 Nonparametric Models
- 18.9 Support Vector Machines
- 18.10 Ensemble Learning
- 18.11 Practical Machine Learning
- 18.12 Summary, Bibliographical and Historical Notes, Exercises
- 19 Knowledge in Learning
 - 19.1 A Logical Formulation of Learning
 - 19.2 Knowledge in Learning
 - 19.3 Explanation-Based Learning
 - 19.4 Learning Using Relevance Information
 - 19.5 Inductive Logic Programming
 - 19.6 Summary, Bibliographical and Historical Notes, Exercises
- 20 Learning Probabilistic Models
 - 20.1 Statistical Learning
 - 20.2 Learning with Complete Data
 - 20.3 Learning with Hidden Variables: The EM Algorithm
 - 20.4 Summary, Bibliographical and Historical Notes, Exercises
- 21 Reinforcement Learning
 - 21.1 Introduction
 - 21.2 Passive Reinforcement Learning
 - 21.3 Active Reinforcement Learning
 - 21.4 Generalization in Reinforcement Learning
 - 21.5 Policy Search
 - 21.6 Applications of Reinforcement Learning
 - 21.7 Summary, Bibliographical and Historical Notes, Exercises
- VI Communicating, perceiving, and acting
- 22 Natural Language Processing
 - 22.1 Language Models
 - 22.2 Text Classification
 - 22.3 Information Retrieval
 - 22.4 Information Extraction
 - 22.5 Summary, Bibliographical and Historical Notes, Exercises
- 23 Natural Language for Communication
 - 23.1 Phrase Structure Grammars
 - 23.2 Syntactic Analysis (Parsing)
 - 23.3 Augmented Grammars and Semantic Interpretation
 - 23.4 Machine Translation
 - 23.5 Speech Recognition
 - 23.6 Summary, Bibliographical and Historical Notes, Exercises
- 24 Perception
 - 24.1 Image Formation
 - 24.2 Early Image-Processing Operations
 - 24.3 Object Recognition by Appearance

- 24.4 Reconstructing the 3D World
- 24.5 Object Recognition from Structural Information
- 24.6 Using Vision
- 24.7 Summary, Bibliographical and Historical Notes, Exercises
- 25 Robotics
 - 25.1 Introduction
 - 25.2 Robot Hardware
 - 25.3 Robotic Perception
 - 25.4 Planning to Move
 - 25.5 Planning Uncertain Movements
 - 25.6 Moving
 - 25.7 Robotic Software Architectures
 - 25.8 Application Domains
 - 25.9 Summary, Bibliographical and Historical Notes, Exercises
- VII Conclusions
- 26 Philosophical Foundations
 - 26.1 Weak AI: Can Machines Act Intelligently?
 - 26.2 Strong AI: Can Machines Really Think?
 - 26.3 The Ethics and Risks of Developing Artificial Intelligence
 - 26.4 Summary, Bibliographical and Historical Notes, Exercises
- 27 AI: The Present and Future
 - 27.1 Agent Components
 - 27.2 Agent Architectures
 - 27.3 Are We Going in the Right Direction?
 - 27.4 What If AI Does Succeed?
- A Mathematical background
 - A.1 Complexity Analysis and $O()$ Notation
 - A.2 Vectors, Matrices, and Linear Algebra
 - A.3 Probability Distributions
- B Notes on Languages and Algorithms
 - B.1 Defining Languages with Backus-Naur Form (BNF)
 - B.2 Describing Algorithms with Pseudocode
 - B.3 Online Help
- Bibliography
- Index

章节摘录

版权页：插图：The last component of the learning agent is the problem generator. It is responsible for suggesting actions that will lead to new and informative experiences. The point is that if the performance element had its way, it would keep doing the actions that are best, given what it knows. But if the agent is willing to explore a little and do some perhaps suboptimal actions in the short run, it might discover much better actions for the long run. The problem generator's job is to suggest these exploratory actions. This is what scientists do when they carry out experiments. Galileo did not think that dropping rocks from the top of a tower in Pisa was valuable in itself. He was not trying to break the rocks or to modify the brains of unfortunate passers-by. His aim was to modify his own brain by identifying a better theory of the motion of objects. To make the overall design more concrete, let us return to the automated taxi example. The performance element consists of whatever collection of knowledge and procedures the taxi has for selecting its driving actions. The taxi goes out on the road and drives, using this performance element. The critic observes the world and passes information along to the learning element. For example, after the taxi makes a quick left turn across three lanes of traffic, the critic observes the shocking language used by other drivers. From this experience, the learning element is able to formulate a rule saying this was a bad action, and the performance element is modified by installation of the new rule. The problem generator might identify certain areas of behavior in need of improvement and suggest experiments, such as trying out the brakes on different road surfaces under different conditions. The learning element can make changes to any of the "knowledge" components shown in the agent diagrams (Figures 2.9, 2.11, 2.13, and 2.14). The simplest cases involve learning directly from the percept sequence. Observation of pairs of successive states of the environment can allow the agent to learn "How the world evolves," and observation of the results of its actions can allow the agent to learn "What my actions do." For example, if the taxi exerts a certain braking pressure when driving on a wet road, then it will soon find out how much deceleration is actually achieved. Clearly, these two learning tasks are more difficult if the environment is only partially observable.

<<人工智能>>

编辑推荐

《人工智能:一种现代的方法(第3版)(影印版)》为大学计算机教育国外著名教材系列之一。

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:<http://www.tushu007.com>