The book was found

Reinforcement Learning: An Introduction (Adaptive Computation And Machine Learning Series)

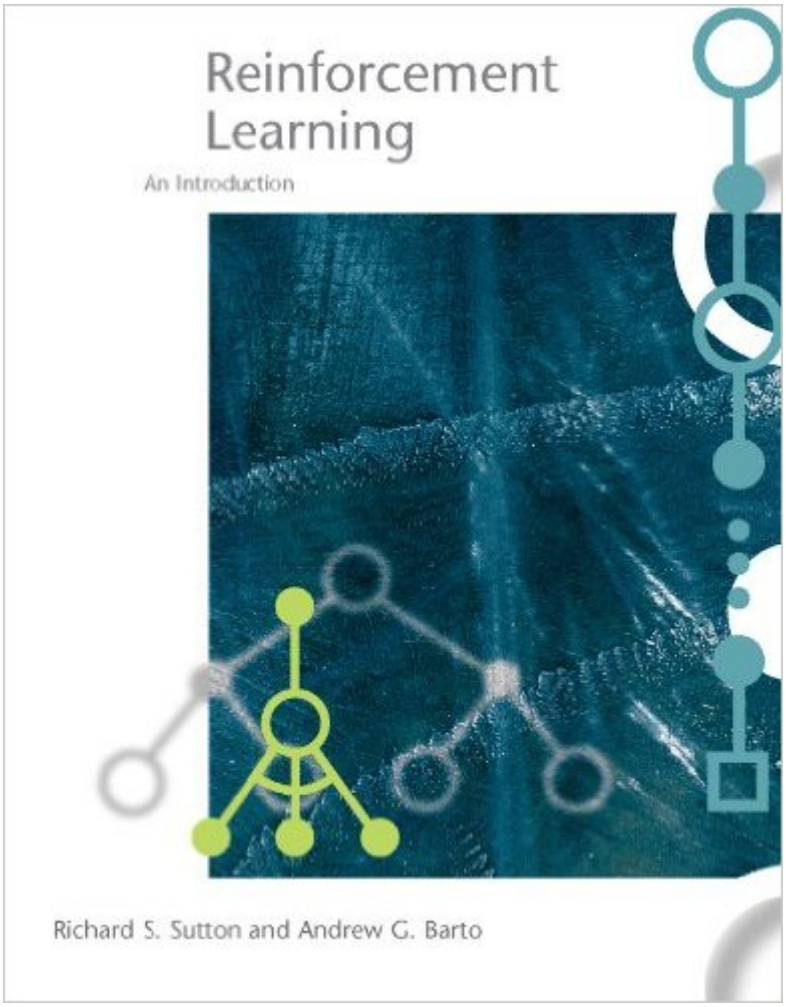

PDF
Adobe
DOWNLOAD EBOOK

Synopsis

Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives when interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the key ideas and algorithms of reinforcement learning. Their discussion ranges from the history of the field's intellectual foundations to the most recent developments and applications. The only necessary mathematical background is familiarity with elementary concepts of probability.The book is divided into three parts. Part I defines the reinforcement learning problem in terms of Markov decision processes. Part II provides basic solution methods: dynamic programming, Monte Carlo methods, and temporal-difference learning. Part III presents a unified view of the solution methods and incorporates artificial neural networks, eligibility traces, and planning; the two final chapters present case studies and consider the future of reinforcement learning.

Book Information

File Size: 8287 KB

Print Length: 344 pages

Publisher: A Bradford Book (February 26, 1998)

Publication Date: February 26, 1998

Sold by:Â Digital Services LLC

Language: English

ASIN: B008H5Q8VA

Text-to-Speech: Not enabled

X-Ray: Not Enabled

Word Wise: Not Enabled

Lending: Not Enabled

Enhanced Typesetting: Not Enabled

Best Sellers Rank: #485,062 Paid in Kindle Store (See Top 100 Paid in Kindle Store) #491 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics #7793 inÂ Kindle Store > Kindle eBooks > Computers & Technology #171059 inÂ Kindle Store > Kindle eBooks > Nonfiction

Customer Reviews

Reinforcement Learning is an exceptionally clear introduction to a field that also goes under names

such as approximate dynamic programming, adaptive dynamic programming and neuro-dynamic programming. The book is written entirely from the perspective of computer science, where problems tend to have discrete states (although potentially large state spaces) and (typically) small action spaces.The book provides numerous step-by-step algorithms that makes it relatively easy to get started writing algorithms. The presentation uses minimal mathematics, and avoids the difficult theory supporting the convergence proofs, making it a nice introduction for undergraduates and graduates alike. But throughout the presentation is evidence of extensive experience with applying these methods to a range of classical problems in artificial intelligence.Students interested in a stronger theoretical foundation should look at Neuro-Dynamic Programming (Optimization and Neural Computation Series, 3). My recent book, Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics), puts far more emphasis on mathematical modeling, and presents the field more from the perspective of the operations research community.

I have this book more than a year now and I am going through it for the second time, so I think I have a pretty good picture about it.The book consists of three parts. In the first part, "The Problem", the authors define the scope of issues reinfocement learning is dealing with and they give some interesting introductory examples. Then, they move on to the concept of evaluative feedback and, eventually, define the reinforcement learning problem formally.The second part, "Elementary Solution Methods" consists of three more-less independent subparts: Dynamic Programming, Monte Carlo Methods and Temporal Difference Learning. All three fundamental reinforcement learning methods are presented in an interesting way and using good examples. Personally, I liked the TD-Learning part best and I agree that this method is indeed the central method and an original contribution of reinforecement learning to the field of machine learning.The third part, "A Unified View" present more advanced techniques. The last chapter gives the most important case studies in reinforcement learning including Samuel’s Checkers Player and Thesauro’s TD-Gammon.The book is very readable and every chapter ends with illustrative exercises (many of them actually are real programming projects!), always useful summary and very valuable bibliographical and historical remarks. Some subchapters are more advanced and therefore marked with ’*’. I really recommend first two parts to any student ofd computer science or anyone interested in machine learning and fuzzy computing. The third part is much more advanced but it would be definitely interesting for advanced computer scientists and graduate students.

Download to continue reading...

Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series) Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and Machine Learning) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear Regression and Logistic Regression in Python (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine (Statistics for Biology and Health) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms Training the Best Dog Ever: A 5-Week Program Using the Power of Positive Reinforcement Parrot Tricks: Teaching Parrots with Positive Reinforcement

Dmca