


The book was found

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



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

Series: Adaptive Computation and Machine Learning series

Hardcover: 322 pages

Publisher: A Bradford Book; 1st Edition edition (March 1, 1998)

Language: English

ISBN-10: 9780262193986

ISBN-13: 978-0262193986

ASIN: 0262193981

Product Dimensions: 7 x 0.8 x 9 inches

Shipping Weight: 1.8 pounds (View shipping rates and policies)

Average Customer Review: 4.5 out of 5 stars 24 customer reviews

Best Sellers Rank: #57,688 in Books (See Top 100 in Books) #23 in [Books > Textbooks >](#)

[Computer Science > Artificial Intelligence](#) #58 in [Books > Computers & Technology >](#)

[Computer Science > AI & Machine Learning > Intelligence & Semantics](#) #4444 in [Books >](#)

[Education & Teaching > Schools & Teaching](#)

Customer Reviews

This is a highly intuitive and accessible introduction to the recent major developments in reinforcement learning, written by two of the field's pioneering contributors. (Dimitri P. Bertsekas and John N. Tsitsiklis, Professors, Department of Electrical Engineering and Computer Science,

Massachusetts Institute of Technology) This book not only provides an introduction to learning theory but also serves as a tremendous source of ideas for further development and applications in the real world. (Toshio Fukuda, Nagoya University, Japan; President, IEEE Robotics and Automantion Society) Reinforcement learning has always been important in the understanding of the driving force behind biological systems, but in the last two decades it has become increasingly important, owing to the development of mathematical algorithms. Barto and Sutton were the prime movers in leading the development of these algorithms and have described them with wonderful clarity in this new text. I predict it will be the standard text. (Dana Ballard, Professor of Computer Science, University of Rochester) The widely acclaimed work of Sutton and Barto on reinforcement learning applies some essentials of animal learning, in clever ways, to artificial learning systems. This is a very readable and comprehensive account of the background, algorithms, applications, and future directions of this pioneering and far-reaching work. (Wolfram Schultz, University of Fribourg, Switzerland)

Richard S. Sutton is Senior Research Scientist, Department of Computer Science, University of Massachusetts. Andrew G. Barto is Professor of Computer Science at the University of Massachusetts.

A great introduction introduction to RL. A bit date at points, but remarkably accessible.

Excellent text from founders of the field

A great book to get started on reinforcement learning.

As a subfield of artificial intelligence, reinforcement learning has shown great success from both a theoretical and practical viewpoint. Taking the form of numerous applications in finance, network engineering, robot toys, and games, it is clear that his learning paradigm shows even greater promise for future developments. The authors summarize the foundations of reinforcement learning, some of this coming from their own work over the last decade. The authors define reinforcement learning as learning how to map situations to actions so as to maximize a numerical reward. The machine that is indulging in reinforcement learning discovers on its own which actions will optimize the reward by trying out these actions. It is the ability of such a machine to learn from experience that distinguishes it from one that is indulging in supervised learning, for in the latter examples are

needed to guide the machine to the proper concept or knowledge. The authors emphasize the "exploration-exploitation" tradeoffs that reinforcement-learning machines have to deal with as they interact with the environment. For the authors, a reinforcement learning system consists of a 'policy', a 'reward function', a 'value function', and a 'model' of the environment. A policy is a mapping from the states of the environment that are perceived by the machine to the actions that are to be taken by the machine when in those states. The reward function maps each perceived state of the environment to a number (the reward). A value function specifies what is the good for the machine over the long run. A model, as the name implies, is a representation of the behavior of the environment. The authors emphasize that all of the reinforcement learning methods that are discussed in the book are concerned with the estimation of value functions, but they point out that other techniques are available for solving reinforcement learning problems, such as genetic algorithms and simulated annealing. The authors use dynamic programming, Monte Carlo simulation, and temporal-difference learning to solve the reinforcement learning problem, but they emphasize that each of these methods will not give a free-lunch. An entire chapter is devoted to each of these methods however, giving the reader a good overview of the weaknesses and strengths of each of these approaches. The differences between them usual boil down to issues of performance rather than accuracy in the generated solutions. Temporal difference learning in fact is viewed in the book as a combination of Monte Carlo and dynamic programming techniques, and in the opinion of this reviewer, has resulted in some of the most impressive successes for applications based on reinforcement learning. One of these is TD-Gammon, developed to play backgammon, and which is also discussed in the book. The authors emphasize that these three main strategies for solving reinforcement learning problems are not mutually exclusive. Instead each of them could be used simultaneously with the others, and they devote a few chapters in the book illustrating how this "unified" approach can be advantageous for reinforcement learning problems. They do this by using explicit algorithms and not just philosophical discussion. These discussions are very interesting and illustrate beautifully the idea that there is no "free lunch" in any of the different algorithms involved in reinforcement learning. In the last chapter of the book the authors overview some of the more successful applications of reinforcement learning, one of them already mentioned. Another one discussed is the 'acrobot', which is a two-link, underactuated robot, which models to some extent the motion of a gymnast on a high bar. The motion of the acrobot is to be controlled by swinging its tip above the first joint, with appropriate rewards given until this goal is reached. The authors use the 'Sarsa' learning algorithm, developed earlier in the book, for solving this reinforcement learning problem. The acrobot is an example of the current intense interest in machine learning of physical

motion and intelligent control theory. Another example discussed in this chapter deals with the problem of elevator dispatching, which the authors include as an example of a problem that cannot be dealt with efficiently by dynamic programming. This problem is studied with Q-learning and via the use of a neural network trained by back propagation. The authors also treat a problem of great importance in the cellular phone industry, namely that of dynamic channel allocation. This problem is formulated as a semi-Markov decision problem, and reinforcement learning techniques were used to minimize the probability of blocking a call. Reinforcement learning has become very important in the communications industry of late, as well as in queuing networks.

A must read if you're taking your first steps into RL. It covers the basics in an easy-to-understand language. Covers MDP, Q Learning, Sarsa, TD-Lambda etc.

I am keep reading the book, so it is not bad.

No wonder this is the most widely adopted textbook on RL.

perfect

[Download to continue reading...](#)

Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Deep Learning (Adaptive Computation and Machine Learning series) Semi-Supervised Learning (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) Algorithms for Reinforcement Learning (Synthesis Lectures on Artificial Intelligence and Machine Learning) IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning) TensorFlow for Deep Learning: From Linear Regression to Reinforcement Learning BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Machine Learning: An Algorithmic Perspective, Second Edition

(Chapman & Hall/Crc Machine Learning & Pattern Recognition) Lose Weight (Self Hypnosis and Subliminal Reinforcement) Glencoe Biology: The Dynamics of Life, Reinforcement and Study Guide, Student Edition (BIOLOGY DYNAMICS OF LIFE) The Sound Reinforcement Handbook Live Sound Reinforcement (Mix Pro Audio Series) JBL Audio Engineering for Sound Reinforcement Bringing Out the Best in People: How to Apply the Astonishing Power of Positive Reinforcement, Third Edition (Business Books)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)