

Markov decision process (MDP) models are widely used for modeling sequential decision-making problems that arise in engineering, economics, computer science, and the social sciences. It is well-known that many real-world problems modeled by MDPs have huge state and/or action spaces, leading to the notorious curse of dimensionality that makes practical solution of the resulting models intractable. In other cases, the system of interest is complex enough that it is not feasible to specify some of the MDP model parameters explicitly, but simulation samples are readily available (e.g., for random transitions and costs). For these settings, various sampling and population-based numerical algorithms have been developed recently to overcome the difficulties of computing an optimal solution in terms of a policy and/or value function. Specific approaches include:

- multi-stage adaptive sampling;
- evolutionary policy iteration;
- evolutionary random policy search; and
- model reference adaptive search.

Simulation-based Algorithms for Markov Decision Processes brings this state-of-the-art research together for the first time and presents it in a manner that makes it accessible to researchers with varying interests and backgrounds. In addition to providing numerous specific algorithms, the exposition includes both illustrative numerical examples and rigorous theoretical convergence results. The algorithms developed and analyzed differ from the successful computational methods for solving MDPs based on neurodynamic programming or reinforcement learning and will complement work in those areas. Furthermore, the authors show how to combine the various algorithms introduced with approximate dynamic programming methods that reduce the size of the state space and ameliorate the effects of dimensionality.

The self-contained approach of this book will appeal not only to researchers in MDPs, stochastic modeling, and control, and simulation but will be a valuable source of instruction and reference for students of control and operations research.

The **Communications and Control Engineering** series reports major technological advances which have potential for great impact in the fields of communication and control. It reflects research in industrial and academic institutions around the world so that the readership can exploit new possibilities as they become available.

ISBN 978-1-84628-689-6



springer.com

Chang · Fu
Hu · Marcus



Simulation-based Algorithms for Markov Decision Processes

Hyeong Soo Chang · Michael C. Fu
Jiaqiao Hu · Steven I. Marcus



Simulation-based Algorithms for Markov Decision Processes

 Springer