

Simulation-Based Algorithms for Markov Decision Processes

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. Many real-world problems modeled by MDPs have huge state and/or action spaces, giving an opening to the curse of dimensionality and so making practical solution of the resulting models intractable. In other cases, the system of interest is too complex to allow explicit specification of some of the MDP model parameters, but simulation samples are readily available (e.g., for random transitions and costs). For these settings, various sampling and population-based algorithms have been developed to overcome the difficulties of computing an optimal solution in terms of a policy and/or value function. Specific approaches include adaptive sampling, evolutionary policy iteration, evolutionary random policy search, and model reference adaptive search. This substantially enlarged new edition reflects the latest developments in novel algorithms and their underpinning theories, and presents an updated account of the topics that have emerged since the publication of the first edition. Includes: innovative material on MDPs, both in constrained settings and with uncertain transition properties;


game-theoretic method for solving MDPs;

theories for developing roll-out based algorithms; and

details of approximation stochastic annealing, a population-based on-line simulation-based algorithm.

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 tuition and reference for students of control and operations research.

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. Many real-world problems modeled by MDPs have huge state and/or action spaces, giving an opening to the curse of dimensionality and so making practical solution of the resulting models intractable. In other cases, the system of interest is too complex to allow explicit specification of some of the MDP model parameters, but simulation samples are readily available (e.g., for random transitions and costs). For these settings, various sampling and population-based algorithms have been developed to overcome the difficulties of computing an optimal solution in terms of a policy and/or value function. Specific approaches include adaptive sampling, evolutionary policy iteration, evolutionary random policy search, and model reference adaptive search. This substantially enlarged new edition reflects the latest developments in novel algorithms and their underpinning theories, and presents an updated account of the topics that have emerged since the publication of the first edition. Includes: innovative material on MDPs, both in constrained settings and with uncertain transition properties; game-theoretic method for solving MDPs; theories for developing roll-out based algorithms; and details of approximation stochastic annealing, a population-based on-line simulation-based algorithm. 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 tuition and reference for students of control and operations research.



106,99 €

99,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9781447159902

Medium: Buch

ISBN: 978-1-4471-5990-2

Verlag: Springer

Erscheinungstermin: 07.03.2015

Sprache(n): Englisch

Auflage: 2. Auflage 2013

Serie: Communications and Control Engineering

Produktform: Kartoniert

Gewicht: 382 g

Seiten: 229

Format (B x H): 155 x 235 mm

