Algorithms For Reinforcement Learning Synthesis Lectures On Artificial Intelligence And Machine Learning

Planning with Markov Decision Processes
Applied Machine Learning
An Introduction
Algorithms for Reinforcement Learning
Reinforcement Learning
Privacy and Incentive
An Interface Layer for Artificial Intelligence
Verification, Synthesis, and Learning in Markov Decision Processes
General Video Game Artificial Intelligence
European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II
Algorithms, Compilers, and Processors for Large-Scale Production
Quantum Machine Learning
Modeling and Optimization in Software-Defined Networks
Deep Learning Systems
Efficient Reinforcement Learning Using Gaussian Processes
Reinforcement Learning, second edition
Active Learning
Learning and Decision-Making from Rank Data
Handbook of Simulation Optimization
VLSI and Hardware Implementations using Modern Machine Learning Methods
Control Systems and Reinforcement Learning
Reasoning with Probabilistic and Deterministic Graphical Models
Advanced Classification Techniques for Healthcare Analysis
Markov Logic
Reinforcement Learning and Stochastic Optimization
Machine Learning and Knowledge Discovery in Databases
Adaptive Representations for Reinforcement Learning
Introduction to Graph Neural Networks
What Quantum Computing Means to Data Mining
An AI Perspective
An AI Perspective
Human Computation
A Lyapunov-Based Approach
Federated Learning
Metric Learning
Reinforcement Learning for Optimal Feedback Control
Abstract Dynamic Programming
Second Edition
Planning with Markov Decision Processes

GIADA RUSH

Planning with Markov Decision Processes Morgan & Claypool Publishers
This dissertation studies the applicability of convex optimization to the formal verification and synthesis of systems that exhibit randomness or stochastic uncertainties. These systems can be represented by a general family of uncertain, partially observable, and parametric Markov decision processes (MDPs). These models have found applications in artificial intelligence, planning, autonomy, and control theory and can accurately characterize dynamic, uncertain environments. The synthesis of policies for this family of models has long been regarded theoretically and empirically intractable. The goal of this dissertation is to develop theoretically sound and computationally efficient synthesis algorithms that provably satisfy formal high-level task specifications in temporal logic. The first part is on developing convex-optimization-based techniques to parameter synthesis in parametric Markov decision processes where the values of the transitions are functions over real-valued parameters. The second part builds on the formulations of the first part and utilizes sampling-based methods for verification and optimization in uncertain MDPs that allow the probabilistic transition function to belong to a so-called uncertainty set. The third part develops inverse reinforcement learning algorithms in partially observable MDPs to several limitations of existing techniques that do not take the information asymmetry between the expert and the agent into account. Finally, the fourth part synthesizes policies for uncertain partially observable MDPs that allow both of the probabilistic transition and observation functions to be uncertain. In each part, a unifying theme is, the resulting algorithms approximate the underlying optimization problem as a convex optimization problem. Additionally, by combining techniques from convex optimization and formal methods, the algorithms bring strong performance guarantees with respect to task specifications. The computational efficiency and applicability of the resulting algorithms are demonstrated in numerous domains such as aircraft collision avoidance, spacecraft and unmanned aerial vehicle motion planning, and joint active perception and planning in urban environments
Applied Machine Learning Morgan & Claypool Publishers
Machine learning is a potential solution to resolve bottleneck issues in VLSI via optimizing tasks in the design process. This book aims to provide the latest machine-learning–based methods, algorithms, architectures, and frameworks designed for VLSI design. The focus is on digital, analog, and mixed-signal design techniques, device modeling, physical design, hardware implementation, testability, reconfigurable design, synthesis and verification, and related areas. Chapters include case studies as well as novel research ideas in the given field. Overall, the book provides practical implementations of VLSI design, IC design, and hardware realization using machine learning techniques.
Features: Provides the details of state-of-the-art machine learning methods used in VLSI design Discusses hardware implementation and device modeling pertaining to machine learning algorithms Explores machine learning for various VLSI architectures and reconfigurable computing Illustrates the latest techniques for device size and feature optimization Highlights the latest case studies and reviews of the methods used for hardware implementation This book is aimed at researchers, professionals, and graduate students in VLSI, machine learning, electrical and electronic engineering, computer engineering, and hardware systems.
An Introduction Morgan & Claypool Publishers
This is the first book on synthetic data for deep learning, and its breadth of coverage may render this book as the default reference on synthetic data for years to come. The book can also serve as an introduction to several other important subfields of machine learning that are seldom touched upon in other books. Machine learning as a discipline would not be possible without the inner workings of optimization at hand. The book includes the necessary sinews of optimization though the crux of the discussion centers on the increasingly popular tool for training deep learning models, namely synthetic data. It is expected that the field of synthetic data will undergo exponential growth in the near future. This book serves as a comprehensive survey of the field. In the simplest case, synthetic data refers to computer-generated graphics used to train computer vision models. There are many more facets of synthetic data to consider. In the section on basic computer vision, the book discusses fundamental computer vision problems, both low-level (e.g., optical flow estimation) and high-level (e.g., object detection and semantic segmentation), synthetic environments and datasets for outdoor and urban scenes (autonomous driving), indoor scenes (indoor navigation), aerial navigation, and simulation environments for robotics. Additionally, it touches upon applications of synthetic data outside computer vision (in neural programming, bioinformatics, NLP, and more). It also surveys the work on improving synthetic data development and alternative ways to produce it such as GANs. The book introduces and reviews several different approaches to synthetic data in various domains of machine learning, most notably the following fields: domain adaptation for making synthetic data more realistic and/or adapting the models to be trained on synthetic data and differential privacy for generating synthetic data with privacy guarantees. This discussion is accompanied by an introduction into generative adversarial networks (GAN) and an introduction to differential privacy.
Algorithms for Reinforcement Learning McGraw-Hill Education
Graphical models (e.g., Bayesian and constraint networks, influence diagrams, and Markov decision processes) have become a central paradigm for knowledge representation and reasoning in both artificial intelligence and computer science in general. These models are used to perform many reasoning tasks, such as scheduling, planning and learning, diagnosis and prediction, design, hardware and software verification, and bioinformatics. These problems can be stated as the formal tasks of constraint satisfaction and satisfiability, combinatorial optimization, and probabilistic inference. It is well known that the tasks are computationally hard, but research during the past three decades has yielded a variety of principles and techniques that significantly advanced the state of the art. This book provides comprehensive coverage of the primary exact algorithms for reasoning with such models. The main feature exploited by the algorithms is the model's graph. We present inference-based, message-passing schemes (e.g., variable-elimination) and search-based, conditioning schemes (e.g., cycle-cutset conditioning and AND/OR search). Each class possesses distinguished characteristics and in particular has different time vs. space behavior. We emphasize the dependence of both schemes on few graph parameters such as the treewidth, cycle-cutset, and (the pseudo-tree) height. The new edition includes the notion of influence diagrams, which focus on sequential decision making under uncertainty. We believe the principles outlined in the book would serve well in moving forward to approximation and anytime-based schemes. The target audience of this book is researchers and students in the artificial intelligence and machine learning area, and beyond.
Reinforcement Learning Morgan & Claypool Publishers
Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning,

including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning.

Privacy and Incentive Springer Nature

This book reviews research developments in diverse areas of reinforcement learning such as model-free actor-critic methods, model-based learning and control, information geometry of policy searches, reward design, and exploration in biology and the behavioral sciences. Special emphasis is placed on advanced ideas, algorithms, methods, and applications. The contributed papers gathered here grew out of a lecture course on reinforcement learning held by Prof. Jan Peters in the winter semester 2018/2019 at Technische Universität Darmstadt. The book is intended for reinforcement learning students and researchers with a firm grasp of linear algebra, statistics, and optimization. Nevertheless, all key concepts are introduced in each chapter, making the content self-contained and accessible to a broader audience.

An Interface Layer for Artificial Intelligence MIT Press

Markov Decision Processes (MDPs) are widely popular in Artificial Intelligence for modeling sequential decision-making scenarios with probabilistic dynamics. They are the framework of choice when designing an intelligent agent that needs to act for long periods of time in an environment where its actions could have uncertain outcomes. MDPs are actively researched in two related subareas of AI, probabilistic planning and reinforcement learning. Probabilistic planning assumes known models for the agent's goals and domain dynamics, and focuses on determining how the agent should behave to achieve its objectives. On the other hand, reinforcement learning additionally learns these models based on the feedback the agent gets from the environment. This book provides a concise introduction to the use of MDPs for solving probabilistic planning problems, with an emphasis on the algorithmic perspective. It covers the whole spectrum of the field, from the basics to state-of-the-art optimal and approximation algorithms. We first describe the theoretical foundations of MDPs and the fundamental solution techniques for them. We then discuss modern optimal algorithms based on heuristic search and the use of structured representations. A major focus of the book is on the numerous approximation schemes for MDPs that have been developed in the AI literature. These include determinization-based approaches, sampling techniques, heuristic functions, dimensionality reduction, and hierarchical representations. Finally, we briefly introduce several extensions of the standard MDP classes that model and solve even more complex planning problems. Table of Contents: Introduction / MDPs / Fundamental Algorithms / Heuristic Search Algorithms / Symbolic Algorithms / Approximation Algorithms / Advanced Notes

Verification, Synthesis, and Learning in Markov Decision Processes Morgan & Claypool Publishers

A how-to guide and scientific tutorial covering the universe of reinforcement learning and control theory for online decision making.

General Video Game Artificial Intelligence Morgan & Claypool Publishers

Reinforcement learning is a learning paradigm concerned with learning to control a system so as to maximize a numerical performance measure that expresses a long-term objective. What distinguishes reinforcement learning from supervised learning is that only partial feedback is given to the learner about the learner's predictions. Further, the predictions may have long term effects through influencing the future state of the controlled system. Thus, time plays a special role. The goal in reinforcement learning is to develop efficient learning algorithms, as well as to understand the algorithms' merits and limitations. Reinforcement learning is of great interest because of the large number of practical applications that it can be used to address, ranging from problems in artificial intelligence to operations research or control engineering. In this book, we focus on those algorithms of reinforcement learning that build on the powerful theory of dynamic programming.We give a fairly comprehensive catalog of learning problems, describe the core ideas, note a large number of state of the art algorithms, followed by the discussion of their theoretical properties and limitations.

European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II KIT Scientific Publishing

This text covers all the fundamentals and presents basic theoretical concepts and a wide range of techniques (algorithms) applicable to challenges in our day-to-day lives. The book recognizes that most of the ideas behind machine learning are simple and straightforward. It provides a platform for hands-on experience through self-study machine learning projects. Datasets for some benchmark applications have been explained to encourage the use of algorithms covered in this book. This is a comprehensive text book on machine learning for undergraduates in computer science and all engineering degree programs. Post graduates and research scholars will find it a useful initial exposure to the subject, before they go for highly theoretical depth in the specific areas of their research. For engineers, scientists, business managers and other practitioners, the book will help build the foundations of machine learning.

Algorithms, Compilers, and Processors for Large-Scale Production Morgan & Claypool Publishers

The visual recognition problem is central to computer vision research. From robotics to information retrieval, many desired applications demand the ability to identify and localize categories, places, and objects. This tutorial overviews computer vision algorithms for visual object recognition and image classification. We introduce primary representations and learning approaches, with an emphasis on recent advances in the field. The target audience consists of researchers or students working in AI, robotics, or vision who would like to understand what methods and representations are available for these problems. This lecture summarizes what is and isn't possible to do reliably today, and overviews key concepts that could be employed in systems requiring visual categorization.

Quantum Machine Learning Springer Science & Business Media

Reinforcement learning encompasses both a science of adaptive behavior of rational beings in uncertain environments and a computational methodology for finding optimal behaviors for challenging problems in control, optimization and adaptive behavior of intelligent agents. As a field, reinforcement learning has progressed tremendously in the past decade. The main goal of this book is to present an up-to-date series of survey articles on the main contemporary sub-fields of reinforcement learning. This includes surveys on partially observable environments, hierarchical task decompositions, relational knowledge representation and predictive state representations. Furthermore, topics such as transfer, evolutionary methods and continuous spaces in reinforcement learning are surveyed. In addition, several chapters review reinforcement learning methods in robotics, in games, and in computational neuroscience. In total seventeen different subfields are presented by mostly young experts in those areas, and together they truly represent a state-of-the-art of current reinforcement learning research. Marco Wiering works at the artificial intelligence department of the University of Groningen in the Netherlands. He has published extensively on various reinforcement learning topics. Martijn van Otterlo works in the cognitive artificial intelligence group at the Radboud University Nijmegen in The Netherlands. He has mainly focused on expressive knowledge representation in reinforcement learning settings.

Modeling and Optimization in Software-Defined Networks Springer

Semi-supervised learning is a learning paradigm concerned with the study of how computers and natural systems such as humans learn in the presence of both labeled and unlabeled data. Traditionally, learning has been studied either in the unsupervised paradigm (e.g., clustering, outlier detection) where all the data are unlabeled, or in the supervised paradigm (e.g., classification, regression) where all the data are labeled. The goal of semi-supervised learning is to understand how combining labeled and unlabeled data may change the learning behavior, and design algorithms that take advantage of such a combination. Semi-supervised learning is of great interest in machine learning and data mining because it can use readily available unlabeled data to improve supervised learning tasks when the labeled data are scarce or expensive. Semi-supervised learning also shows potential as a quantitative tool to understand human category learning, where most of the input is self-evidently unlabeled. In this introductory book, we present some popular semi-supervised learning models, including self-training, mixture models, co-training and multiview learning, graph-based methods, and semi-supervised support vector machines. For each model, we discuss its basic mathematical formulation. The success of semi-supervised learning depends critically on some underlying assumptions. We emphasize the assumptions made by each model and give counterexamples when appropriate to demonstrate the limitations of the different models. In addition, we discuss semi-supervised learning for cognitive psychology. Finally, we give a computational learning theoretic perspective on semi-supervised learning, and we conclude the book with a brief discussion of open questions in the field. Table of Contents: Introduction to Statistical Machine Learning / Overview of Semi-Supervised Learning / Mixture Models and EM / Co-Training / Graph-Based Semi-Supervised Learning / Semi-Supervised Support Vector Machines / Human Semi-Supervised Learning / Theory and Outlook

Deep Learning Systems Athena Scientific

The ubiquitous challenge of learning and decision-making from rank data arises in situations where intelligent systems collect preference and behavior data from humans, learn from the data, and then use the data to help humans make efficient, effective, and timely decisions. Often, such data are represented by rankings. This book surveys some recent progress toward addressing the challenge from the considerations of statistics, computation, and socio-economics. We will cover classical statistical models for rank data, including random utility models, distance-based models, and mixture models. We will discuss and compare classical and state of-the-art algorithms, such as algorithms based on Minorize-Majorization (MM), Expectation-Maximization (EM), Generalized Method-of-Moments (GMM), rank breaking, and tensor decomposition. We will also introduce principled Bayesian preference elicitation frameworks for collecting rank data. Finally, we will examine socio-economic aspects of statistically desirable decision-making mechanisms, such as Bayesian estimators. This book can be useful in three ways: (1) for theoreticians in statistics and machine learning to better understand the considerations and caveats of learning from rank data, compared to learning from other types of data, especially cardinal data; (2) for practitioners to apply algorithms covered by the book for sampling, learning, and aggregation; and (3) as a textbook for graduate students or advanced undergraduate students to learn about the field. This book requires that the reader has basic knowledge in probability, statistics, and algorithms. Knowledge in social choice would also help but is not required.

Efficient Reinforcement Learning Using Gaussian Processes Morgan & Claypool Publishers

Medical and information communication technology professionals are working to develop robust classification techniques, especially in healthcare data/image analysis, to ensure quick diagnoses and treatments to patients. Without fast and immediate access to healthcare databases and information, medical professionals' success rates and treatment options become limited and fall to disastrous levels. Advanced Classification Techniques for Healthcare Analysis provides emerging insight into classification techniques in delivering quality, accurate, and affordable healthcare, while also discussing the impact health data has on medical treatments. Featuring coverage on a broad range of topics such as early diagnosis, brain-computer interface, metaheuristic algorithms, clustering techniques, learning schemes, and mobile telemedicine, this book is ideal for medical professionals, healthcare administrators, engineers, researchers, academicians, and technology developers seeking current research on furthering information and communication technology that improves patient care.

Reinforcement Learning, second edition Morgan & Claypool Publishers

This is the 3rd edition of a research monograph providing a synthesis of old research on the foundations of dynamic programming (DP), with the modern theory of approximate DP and new research on semicontractive models. It aims at a unified and economical development of the core theory and algorithms of total cost sequential decision problems, based on the strong connections of the subject with fixed point theory. The analysis focuses on the abstract mapping that underlies DP and defines the mathematical character of the associated problem. The discussion centers on two fundamental properties that this mapping may have: monotonicity and (weighted sup-norm) contraction. It turns out that the nature of the analytical and algorithmic DP theory is determined primarily by the presence or absence of these two properties, and the rest of the problem's structure is largely inconsequential. New research is focused on two areas: 1) The ramifications of these properties in the context of algorithms for approximate DP, and 2) The new class of semicontractive models, exemplified by stochastic shortest path problems, where some but not all policies are contractive. The 3rd edition is very similar to the 2nd edition, except for the addition of a new chapter (Chapter 5), which deals with abstract DP models for sequential minimax problems and zero-sum games, The book is an excellent supplement to several of our books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (Athena Scientific, 2017), Reinforcement Learning and Optimal Control (Athena Scientific, 2019), and Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020).

Active Learning IGI Global

Reinforcement Learning for Optimal Feedback Control develops model-based and data-driven reinforcement learning methods for solving optimal control problems in nonlinear deterministic dynamical systems. In order to achieve learning under uncertainty, data-driven methods for identifying system models in real-time are also developed. The book illustrates the advantages gained from the use of a model and the use of previous experience in the form of recorded data through simulations and experiments. The book's focus on deterministic systems allows for an in-depth Lyapunov-based analysis of the performance of the

methods described during the learning phase and during execution. To yield an approximate optimal controller, the authors focus on theories and methods that fall under the umbrella of actor–critic methods for machine learning. They concentrate on establishing stability during the learning phase and the execution phase, and adaptive model-based and data-driven reinforcement learning, to assist readers in the learning process, which typically relies on instantaneous input-output measurements. This monograph provides academic researchers with backgrounds in diverse disciplines from aerospace engineering to computer science, who are interested in optimal reinforcement learning functional analysis and functional approximation theory, with a good introduction to the use of model-based methods. The thorough treatment of an advanced treatment to control will also interest practitioners working in the chemical-process and power-supply industry.

Learning and Decision-Making from Rank Data Morgan & Claypool Publishers

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

Handbook of Simulation Optimization Morgan & Claypool Publishers

This handbook presents state-of-the-art research in reinforcement learning, focusing on its applications in the control and game theory of dynamic systems and future directions for related research and technology. The contributions gathered in this book deal with challenges faced when using learning and adaptation methods to solve academic and industrial problems, such as optimization in dynamic environments with single and multiple agents, convergence and performance analysis, and online implementation. They explore means by which these difficulties can be solved, and cover a wide range of related topics including: deep learning; artificial intelligence; applications of game theory; mixed modality learning; and multi-agent reinforcement learning. Practicing engineers and scholars in the field of machine learning, game theory, and autonomous control will find the Handbook of Reinforcement Learning and Control to be thought-provoking, instructive and informative.

VLSI and Hardware Implementations using Modern Machine Learning Methods MIT Press

The key idea behind active learning is that a machine learning algorithm can perform better with less training if it is allowed to choose the data from which it learns. An active learner may pose "queries," usually in the form of unlabeled data instances to be labeled by an "oracle" (e.g., a human annotator) that already understands the nature of the problem. This sort of approach is well-motivated in many modern machine learning and data mining applications, where unlabeled data may be abundant or easy to come by, but training labels are difficult, time-consuming, or expensive to obtain. This book is a general introduction to active learning. It outlines several scenarios in which queries might be formulated, and details many query selection algorithms which have been organized into four broad categories, or "query selection frameworks." We also touch on some of the theoretical foundations of active learning, and conclude with an overview of the strengths and weaknesses of these approaches in practice, including a summary of ongoing work to address these open challenges and opportunities. Table of Contents: Automating Inquiry / Uncertainty Sampling / Searching Through the Hypothesis Space / Minimizing Expected Error and Variance / Exploiting Structure in Data / Theory / Practical Considerations

Related with Algorithms For Reinforcement Learning Synthesis Lectures On Artificial Intelligence And Machine Learning:

• Andre Drummond Contract History : click here