

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

Verification, Synthesis, and Learning in Markov Decision Processes
 Second Edition
 Planning with Markov Decision Processes
 State-of-the-Art
 Convex Optimization Meets Formal Methods
 Privacy and Incentive
 Applied Machine Learning
 Learning and Decision-Making from Rank Data
 Reinforcement Learning
 Handbook of Reinforcement Learning and Control
 Compiling Algorithms for Heterogeneous Systems
 Foundations of Machine Learning, second edition
 Lifelong Machine Learning
 An AI Perspective
 Handbook of Simulation Optimization
 Reinforcement Learning Algorithms: Analysis and Applications
 3rd Edition
 An AI Perspective
 Markov Logic
 Advanced Classification Techniques for Healthcare Analysis
 Insights in Reinforcement Learning
 Reasoning with Probabilistic and Deterministic Graphical Models
 Algorithms for Reinforcement Learning
 Human Computation
 Machine Learning and Knowledge Discovery in Databases
 Reinforcement Learning for Optimal Feedback Control
 An Introduction
 Algorithms, Compilers, and Processors for Large-Scale Production
 What Quantum Computing Means to Data Mining
 Reinforcement Learning and Stochastic Optimization
 A Lyapunov-Based Approach
 Reinforcement Learning, second edition
 Quantum Machine Learning
 Federated Learning
 Active Learning
 Graph Representation Learning
 Introduction to Semi-supervised Learning
 Deep Learning Systems
 Active Learning

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

Downloaded from usabutt.onpol.l.co.uk by guest

COPELAND EZRA

Verification, Synthesis, and Learning in Markov Decision Processes Cambridge University Press
 Algorithms for Reinforcement Learning Morgan & Claypool Publishers

Second Edition Springer

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

Planning with Markov Decision Processes 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.

State-of-the-Art 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 *Convex Optimization Meets Formal Methods* Morgan & Claypool Publishers

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.

Privacy and Incentive 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.

Applied Machine Learning Springer Nature

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 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.

Reinforcement Learning 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.

Handbook of Reinforcement Learning and Control Springer Nature

This two-volume set LNAI 7523 and LNAI 7524 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2012, held in Bristol, UK, in September 2012. The 105 revised research papers presented together with 5 invited talks were carefully reviewed and selected from 443 submissions. The final

sections of the proceedings are devoted to Demo and Nectar papers. The Demo track includes 10 papers (from 19 submissions) and the Nectar track includes 4 papers (from 14 submissions). The papers grouped in topical sections on association rules and frequent patterns; Bayesian learning and graphical models; classification; dimensionality reduction, feature selection and extraction; distance-based methods and kernels; ensemble methods; graph and tree mining; large-scale, distributed and parallel mining and learning; multi-relational mining and learning; multi-task learning; natural language processing; online learning and data streams; privacy and security; rankings and recommendations; reinforcement learning and planning; rule mining and subgroup discovery; semi-supervised and transductive learning; sensor data; sequence and string mining; social network mining; spatial and geographical data mining; statistical methods and evaluation; time series and temporal data mining; and transfer learning.

Compiling Algorithms for Heterogeneous Systems Hado van Hasselt

Most emerging applications in imaging and machine learning must perform immense amounts of computation while holding to strict limits on energy and power. To meet these goals, architects are building increasingly specialized compute engines tailored for these specific tasks. The resulting computer systems are heterogeneous, containing multiple processing cores with wildly different execution models. Unfortunately, the cost of producing this specialized hardware—and the software to control it—is astronomical. Moreover, the task of porting algorithms to these heterogeneous machines typically requires that the algorithm be partitioned across the machine and rewritten for each specific architecture, which is time consuming and prone to error. Over the last several years, the authors have approached this problem using domain-specific languages (DSLs): high-level programming languages customized for specific domains, such as database manipulation, machine learning, or image processing. By giving up generality, these languages are able to provide high-level abstractions to the developer while producing high performance output. The purpose of this book is to spur the adoption and the creation of domain-specific languages, especially for the task of creating hardware designs. In the first chapter, a short historical journey explains the forces driving computer architecture today. Chapter 2 describes the various methods for producing designs for accelerators, outlining the push for more abstraction and the tools that enable designers to work at a higher conceptual level. From there, Chapter 3 provides a brief introduction to image processing algorithms and hardware design patterns for implementing them. Chapters 4 and 5 describe and compare Darkroom and Halide, two domain-specific languages created for image processing that produce high-performance designs for both FPGAs and CPUs from the same source code, enabling rapid design cycles and quick porting of algorithms. The final section describes how the DSL approach also simplifies the problem of interfacing between application code and the accelerator by generating the driver stack in addition to the accelerator configuration. This book should serve as a useful introduction to domain-specialized computing for computer architecture students and as a primer on domain-specific languages and image processing hardware for those with more experience in the field. *Foundations of Machine Learning, second edition* Morgan & Claypool Publishers

This book provides a comprehensive and self-contained introduction to federated learning, ranging from the basic knowledge and theories to various key applications. Privacy and incentive issues are the focus of this book. It is timely as federated learning is becoming popular after the release of the General Data Protection Regulation (GDPR). Since federated learning aims to enable a machine model to be collaboratively trained without each party exposing private data to others. This setting adheres to regulatory requirements of data privacy protection such as GDPR. This book contains three main parts. Firstly, it introduces different privacy-preserving methods for protecting a federated learning model against different types of attacks such as data leakage and/or data poisoning. Secondly, the book presents incentive mechanisms which aim to encourage individuals to participate in the federated learning ecosystems. Last but not least, this book also describes how federated learning can be applied in industry and business to address data silo and privacy-preserving problems. The book is intended for readers from both the academia and the industry, who would like to learn about federated learning, practice its implementation, and apply it in their own business. Readers are expected to have some basic understanding of linear algebra, calculus, and neural network. Additionally, domain knowledge in FinTech and marketing would be helpful."

Lifelong Machine Learning KIT Scientific Publishing
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.

An AI Perspective Springer Nature

Lifelong Machine Learning, Second Edition is an introduction to an advanced machine learning paradigm that continuously learns by accumulating past knowledge that it then uses in future learning and problem solving. In contrast, the current dominant machine learning paradigm learns in isolation: given a training dataset, it runs a machine learning algorithm on the dataset to produce a model that is then used in its intended application. It makes no attempt to retain the learned knowledge and use it in subsequent learning. Unlike this isolated system, humans learn effectively with only a few examples precisely because our learning is very knowledge-driven: the knowledge learned in the past helps us learn new things with little data or effort. Lifelong learning aims to emulate this capability, because without it, an AI system cannot be considered truly intelligent. Research in lifelong learning has developed significantly in the relatively short time since the first edition of this book was published. The purpose of this second edition is to expand the definition of lifelong learning, update the content of several chapters, and add a new chapter about continual learning in deep neural networks—which has been actively researched over the past two or three years. A few chapters have also been reorganized to make each of them more coherent for the reader. Moreover, the authors want to propose a unified framework for the research area. Currently, there are several research topics in machine learning that are closely related to lifelong learning—most notably, multi-task learning, transfer learning, and meta-learning—because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. Its goal is to introduce this emerging machine learning paradigm and present a comprehensive survey and review of the important research results and latest ideas in the area. This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book for courses in any of these related fields. *Handbook of Simulation Optimization* Springer

Similarity between objects plays an important role in both human cognitive processes and artificial systems for recognition and categorization. How to appropriately measure such similarities for a given task is crucial to the performance of many machine learning, pattern recognition and data mining methods. This book is devoted to metric learning, a set of techniques to automatically learn similarity and distance functions from data that has attracted a lot of interest in machine learning and related fields in the past ten years. In this book, we provide a thorough review of the metric learning literature that covers algorithms, theory and applications for both numerical and structured data. We first introduce relevant definitions and classic metric functions, as well as examples of their use in machine learning and data mining. We then review a wide range of metric learning algorithms, starting with the simple setting of linear distance and similarity learning. We show how one may scale-up these methods to very large amounts of training data. To go beyond the linear case, we discuss methods that learn nonlinear metrics or multiple linear metrics throughout the feature space, and review methods for more complex settings such as multi-task and semi-supervised learning. Although most of the existing work has focused on numerical data, we cover the literature on metric learning for structured data like strings, trees, graphs and time series. In the more technical part of the book, we present some recent statistical frameworks for analyzing the generalization performance in metric learning and derive results for some of the algorithms presented earlier. Finally, we illustrate the relevance of metric learning in real-world problems through a series of successful applications to computer vision, bioinformatics and information retrieval.

Reinforcement Learning Algorithms: Analysis and Applications MIT Press

Human computation is a new and evolving research area that

centers around harnessing human intelligence to solve computational problems that are beyond the scope of existing Artificial Intelligence (AI) algorithms. With the growth of the Web, human computation systems can now leverage the abilities of an unprecedented number of people via the Web to perform complex computation. There are various genres of human computation applications that exist today. Games with a purpose (e.g., the ESP Game) specifically target online gamers who generate useful data (e.g., image tags) while playing an enjoyable game.

Crowdsourcing marketplaces (e.g., Amazon Mechanical Turk) are human computation systems that coordinate workers to perform tasks in exchange for monetary rewards. In identity verification tasks, users perform computation in order to gain access to some online content; an example is reCAPTCHA, which leverages millions of users who solve CAPTCHAs every day to correct words in books that optical character recognition (OCR) programs fail to recognize with certainty. This book is aimed at achieving four goals: (1) defining human computation as a research area; (2) providing a comprehensive review of existing work; (3) drawing connections to a wide variety of disciplines, including AI, Machine Learning, HCI, Mechanism/Market Design and Psychology, and capturing their unique perspectives on the core research questions in human computation; and (4) suggesting promising research directions for the future. Table of Contents: Introduction / Human Computation Algorithms / Aggregating Outputs / Task Routing / Understanding Workers and Requesters / The Art of Asking Questions / The Future of Human Computation
3rd Edition Morgan & Claypool Publishers

Graphs are useful data structures in complex real-life applications such as modeling physical systems, learning molecular fingerprints, controlling traffic networks, and recommending friends in social networks. However, these tasks require dealing with non-Euclidean graph data that contains rich relational information between elements and cannot be well handled by traditional deep learning models (e.g., convolutional neural networks (CNNs) or recurrent neural networks (RNNs)). Nodes in graphs usually contain useful feature information that cannot be well addressed in most unsupervised representation learning methods (e.g., network embedding methods). Graph neural networks (GNNs) are proposed to combine the feature information and the graph structure to learn better representations on graphs via feature propagation and aggregation. Due to its convincing performance and high interpretability, GNN has recently become a widely applied graph analysis tool. This book provides a

comprehensive introduction to the basic concepts, models, and applications of graph neural networks. It starts with the introduction of the vanilla GNN model. Then several variants of the vanilla model are introduced such as graph convolutional networks, graph recurrent networks, graph attention networks, graph residual networks, and several general frameworks. Variants for different graph types and advanced training methods are also included. As for the applications of GNNs, the book categorizes them into structural, non-structural, and other scenarios, and then it introduces several typical models on solving these tasks. Finally, the closing chapters provide GNN open resources and the outlook of several future directions.

An AI Perspective 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.

Markov Logic IGI Global

This book provides a quick reference and insights into modeling and optimization of software-defined networks (SDNs). It covers various algorithms and approaches that have been developed for optimizations related to the control plane, the considerable research related to data plane optimization, and topics that have significant potential for research and advances to the state-of-the-art in SDN. Over the past ten years, network programmability has transitioned from research concepts to more mainstream technology through the advent of technologies amenable to programmability such as service chaining, virtual network functions, and programmability of the data plane. However, the rapid development in SDN technologies has been the key driver

behind its evolution. The logically centralized abstraction of network states enabled by SDN facilitates programmability and use of sophisticated optimization and control algorithms for enhancing network performance, policy management, and security. Furthermore, the centralized aggregation of network telemetry facilitates use of data-driven machine learning-based methods. To fully unleash the power of this new SDN paradigm, though, various architectural design, deployment, and operations questions need to be addressed. Associated with these are various modeling, resource allocation, and optimization opportunities. The book covers these opportunities and associated challenges, which represent a "call to arms" for the SDN community to develop new modeling and optimization methods that will complement or improve on the current norms.

Advanced Classification Techniques for Healthcare Analysis

Springer Science & Business Media

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).

Best Sellers - Books :

- [Taylor Swift: A Little Golden Book Biography By Wendy Loggia](#)
- [Girl In Pieces](#)
- [If He Had Been With Me](#)
- [Killers Of The Flower Moon: The Osage Murders And The Birth Of The Fbi By David Grann](#)
- [The Very Hungry Caterpillar](#)
- [Mad Honey: A Novel](#)
- [My Butt Is So Christmassy! By Dawn Mcmillan](#)
- [How To Catch A Mermaid](#)
- [You Will Own Nothing: Your War With A New Financial World Order And How To Fight Back By Carol Roth](#)
- [Stop Overthinking: 23 Techniques To Relieve Stress, Stop Negative Spirals, Declutter Your Mind, And Focus On The Present \(the Path To Calm\) By Nick Trenton](#)