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

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 fro 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.

REINFORCEMENT LEARNING AND STOCHASTIC OPTIMIZATION Clearing the jungle of stochastic optimization Sequential decision problems, which consist of “decision, information, decision, information,” are ubiquitous, spanning virtually every human activity ranging from business applications, health (personal and public health, and medical decision making), energy, the sciences, all fields of engineering, finance, and e-commerce. The diversity of applications attracted the attention of at least 15 distinct fields of research, using eight distinct notational systems which produced a vast array of analytical tools. A byproduct is that powerful tools developed in one community may be unknown to other communities. Reinforcement Learning and Stochastic Optimization offers a single canonical framework that can model any sequential decision problem using five core components: state variables, decision variables, exogenous information variables, transition function, and objective function. This book highlights twelve types of uncertainty that might enter any model and pulls together the diverse set of methods for maki decisions, known as policies, into four fundamental classes that span every method suggested in the academic literature or used in practice. Reinforcement Learning and Stochastic Optimization is the first book to provide a balanced treatment of the different methods for modeling and solving sequential decision problems, following the style used by most books on machine lear optimization, and simulation. The presentation is designed for readers with a course in probability and statistics, and an interest in modeling and applications. Linear programming is occasionally used for specific problem classes. The book is designed for readers who are new to the field, as well as those with some background in optimization under uncertainty. Throughout this b readers will find references to over 100 different applications, spanning pure learning problems, dynamic resource allocation problems, general state-dependent problems, and hybrid learning/resource allocation problems such as those that arose in the COVID pandemic. There are 370 exercises, organized into seven groups, ranging from review questions, modeling, computati problem solving, theory, programming exercises and a “diary problem” that a reader chooses at the beginning of the book, and which is used as a basis for questions throughout the rest of the book.

Most subfields of computer science have an interface layer via which applications communicate with the infrastructure, and this is key to their success (e.g., the Internet in networking, the relational model in databases, etc.). So far this interface layer has been missing in AI. First-order logic and probabilistic graphical models each have some of the necessary features, but a viabl interface layer requires combining both. Markov logic is a powerful new language that accomplishes this by attaching weights to first-order formulas and treating them as templates for features of Markov random fields. Most statistical models in wide use are special cases of Markov logic, and first-order logic is its infinite-weight limit. Inference algorithms for Markov logic combine ideas from satisfiability, Markov chain Monte Carlo, belief propagation, and resolution. Learning algorithms make use of conditional likelihood, convex optimization, and inductive logic programming. Markov logic has been successfully applied to problems in information extraction and integration, natural language processing, robot mapping, social networks, computational biology, and others, and is the basis of the open-source Alchemy system.

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 requi visual categorization.

Control Systems and Reinforcement Learning
Abstract Dynamic Programming
Reinforcement Learning for Optimal Feedback Control
Reinforcement Learning Algorithms: Analysis and Applications
A Unified Framework for Sequential Decisions
An AI Perspective

MACHINE LEARNING AND DATA SCIENCE Written and edited by a team of experts in the field, this collection of papers reflects the most up-to-date and comprehensive current state of machine learning and data science for industry, government, and academia. Machine learning (ML) and data science (DS) are very active topics with an extensive scope, both in terms of theory and applications. They have been established as an important emergent scientific field and paradigm driving research evolution in such disciplines as statistics, computing science and intelligence science, and practical transformation in such domains as science, engineering, the public sector, business, social science, and lifestyle. Simultaneously, their applications provide important challenges that can often be addressed only with innovative machine learning and data science algorithms. These algorithms encompass the larger areas of artificial intelligence, data analytics, machine learning, pattern recognition, natural language understanding, and big data manipulation. They also tackle related new scientific challenges, ranging from data capture, creation, storage, retrieval, sharing, analysis, optimization, and visualization, to integrative analysis across heterogeneous and interdependent complex resources for better decision-making, collaboration, and, ultimately, value creation.

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.

Algorithms for Reinforcement LearningMorgan & Claypool Publishers

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.

Reinforcement Learning
Algorithms for Reinforcement Learning
Machine Learning and Data Science
Metric Learning
Proceedings of MISS 2021
An Interface Layer for Artificial Intelligence

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

Reinforcement learning (RL) and adaptive dynamic programming (ADP) has been one of the most critical research fields in science and engineering for modern complex systems. This book describes the latest RL and ADP techniques for decision and control in human engineered systems, covering both single player decision and control and multi-player games. Edited by the pioneers of RL and ADP research, the book brings together ideas and methods from many fields and provides an important and timely guidance on controlling a wide variety of systems, such as robots, industrial processes, and economic decision-making.

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).

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. Table of Contents: Introduction / Metrics / Properties of Metric Learning Algorithms / Linear Metric Learning / Nonlinear and Local Metric Learning / Metric Learning for Special Settings / Metric Learning for Structured Data / Generalization Guarantees for Metric Learning / Applications / Conclusion / Bibliography / Authors' Biographies

Markov Logic
3rd Edition
An Introduction
Learning and Decision-Making from Rank Data
VLSI and Hardware Implementations using Modern Machine Learning Methods
State-of-the-Art

Quantum Machine Learning bridges the gap between abstract developments in quantum computing and the applied research on machine learning. Paring down the complexity of the disciplines involved, it focuses on providing a synthesis that explains the most important machine learning algorithms in a quantum framework. Theoretical advances in quantum computing are hard to follow for computer scientists, and sometimes even for researchers involved in the field. The lack of a step-by-step guide hampers the broader understanding of this emergent interdisciplinary body of research. Quantum Machine Learning sets the scene for a deeper understanding of the subject for readers of different backgrounds. The author has carefully constructed a clear comparison of classical learning algorithms and their quantum counterparts, thus making differences in computational complexity and learning performance apparent. This book synthesizes of a broad array of research into a manageable and concise presentation, with practical examples and applications. Bridges the gap between abstract developments in quantum computing with the applied research on machine learning Provides the theoretical minimum of machine learning, quantum mechanics, and quantum computing Gives step-by-step guidance to a broader understanding of this emergent interdisciplinary body of research

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.

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. 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 while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

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.

Reinforcement Learning and Optimal Control
Fundamentals and Applications
Lifelong Machine Learning, Second Edition
Machine and Deep Learning Algorithms and Applications
Efficient Reinforcement Learning Using Gaussian Processes
Machine Intelligence and Smart Systems

This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to codesign algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers that optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment platform constraints when designing models. Performance engineers should support optimizations across diverse models, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system stack in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

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.

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.

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.

Reinforcement Learning and Stochastic Optimization
European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II
A Lyapunov-Based Approach
Graph Representation Learning
Convex Optimization Meets Formal Methods
Second Edition

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.”

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.

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

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 Universit ä t 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.

Foundations of Machine Learning, second edition
Planning with Markov Decision Processes
Introduction to Semi-supervised Learning

Compiling Algorithms for Heterogeneous Systems
Synthetic Data for Deep Learning

This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (4th edition, Athena Scientific, 2017), Abstract Dynamic Programming (2nd edition, Athena Scientific, 2018), and Nonlinear Programming (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations). The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator. The book is related and supplemented by the companion research monograph Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and slides from a 2021 course at ASU, which address a selection of topics from both books.

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.

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.

A high school student can create deep Q-learning code to control her robot, without any understanding of the meaning of 'deep' or 'Q', or why the code sometimes fails. This book is designed to explain the science behind reinforcement learning and optimal control in a way that is accessible to students with a background in calculus and matrix algebra. A unique focus is algorithm design to obtain the fastest possible speed of convergence for learning algorithms, along with insight into why reinforcement learning sometimes fails. Advanced stochastic process theory is avoided at the start by substituting random exploration with more intuitive deterministic probing for learning. Once these ideas are understood, it is not difficult to master techniques rooted in stochastic control. These topics are covered in the second part of the book, starting with Markov chain theory and ending with a fresh look at actor-critic methods for reinforcement learning.

Handbook of Reinforcement Learning and Control
Human Computation
Privacy and Incentive
Visual Object Recognition
Reinforcement Learning and Approximate Dynamic Programming for Feedback Control
Algorithms, Compilers, and Processors for Large-Scale Production

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

This book introduces basic machine learning concepts and applications for a broad audience that includes students, faculty, and industry practitioners. We begin by describing how machine learning provides capabilities to computers and embedded systems to learn from data. A typical machine learning algorithm involves training, and generally the performance of a machine learning model improves with more training data. Deep learning is a sub-area of machine learning that involves extensive use of layers of artificial neural networks typically trained on massive amounts of data. Machine and deep learning methods are often used in contemporary data science tasks to address the growing data sets and detect, cluster, and classify data patterns. Although machine learning commercial interest has grown relatively recently, the roots of machine learning go back to decades ago. We note that nearly all organizations, including industry, government, defense, and health, are using machine learning to address a variety of needs and applications. The machine learning paradigms presented can be broadly divided into the following three categories: supervised learning, unsupervised learning, and semi-supervised learning. Supervised learning algorithms focus on learning a mapping function, and they are trained with supervision on labeled data. Supervised learning is further sub-divided into classification and regression algorithms. Unsupervised learning typically does not have access to ground truth, and often the goal is to learn or uncover the hidden pattern in the data. Through semi-supervised learning, one can effectively utilize a large volume of unlabeled data and a limited amount of labeled data to improve machine learning model performances. Deep learning and neural networks are also covered in this book. Deep neural networks have attracted a lot of interest during the last ten years due to the availability of graphics processing units (GPU) computational power, big data, and new software platforms. They have strong capabilities in terms of learning complex mapping functions for different types of data. We organize the book as follows. The book starts by introducing concepts in supervised, unsupervised, and semi-supervised learning. Several algorithms and their inner workings are presented within these three categories. We then continue with a brief introduction to artificial neural network algorithms and their properties. In addition, we cover an array of applications and provide extensive bibliography. The book ends with a summary of the key machine learning concepts.

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

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

Adaptive Representations for Reinforcement Learning

Active Learning

Lifelong Machine Learning

Deep Learning Systems

Modeling and Optimization in Software-Defined Networks

Federated Learning

This book presents new algorithms for reinforcement learning, a form of machine learning in which an autonomous agent seeks a control policy for a sequential decision task. Since current methods typically rely on manually designed solution representations, agents that automatically adapt their own representations have the potential to dramatically improve performance. This book introduces two novel approaches for automatically discovering high-performing representations. The first approach synthesizes temporal difference methods, the traditional approach to reinforcement learning, with evolutionary methods, which can learn representations for a broad class of optimization problems. This synthesis is accomplished by customizing evolutionary methods to the on-line nature of reinforcement learning and using them to evolve representations for value function approximators. The second approach automatically learns representations based on piecewise-constant approximations of value functions. It begins with coarse representations and gradually refines them during learning, analyzing the current policy and value function to deduce the best refinements. This book also introduces a novel method for devising input representations. This method addresses the feature selection problem by extending an algorithm that evolves the topology and weights of neural networks such that it evolves their inputs too. In addition to introducing these new methods, this book presents extensive empirical results in multiple domains demonstrating that these techniques can substantially improve performance over methods with manual representations.

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.

This book is a collection of peer-reviewed best selected research papers presented at the Second International Conference on Machine Intelligence and Smart Systems (MISS 2021), organized during September 24–25, 2021, in Gwalior, India. The book presents new advances and research results in the fields of machine intelligence, artificial intelligence and smart systems. It includes main paradigms of machine intelligence algorithms, namely (1) neural networks, (2) evolutionary computation, (3) swarm intelligence, (4) fuzzy systems and (5) immunological computation. Scientists, engineers, academicians, technology developers, researchers, students and government officials will find this book useful in handling their complicated real-world issues by using machine intelligence methodologies.

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. Table of Contents: Markov Decision Processes / Value Prediction Problems / Control / For Further Exploration

Machine Learning and Knowledge Discovery in Databases

Verification, Synthesis, and Learning in Markov Decision Processes

Insights in Reinforcement Learning

Reinforcement Learning, second edition

Quantum Machine Learning

What Quantum Computing Means to Data Mining

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.