

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

As recognized, adventure as with ease as experience virtually lesson, amusement, as capably as conformity can be gotten by just checking out a ebook **Algorithms For Reinforcement Learning Synthesis Lectures On Artificial Intelligence And Machine Learning** in addition to it is not directly done, you could resign yourself to even more on the order of this life, around the world.

We find the money for you this proper as with ease as simple showing off to get those all. We have enough money Algorithms For Reinforcement Learning Synthesis Lectures On Artificial Intelligence And Machine Learning and numerous ebook collections from fictions to scientific research in any way. in the midst of them is this Algorithms For Reinforcement Learning Synthesis Lectures On Artificial Intelligence And Machine Learning that can be your partner.

Metric Learning - Aurélien Muise 2022-05-31
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

Tensor Voting - Philippos Mordohai 2007

This lecture presents research on a general framework for perceptual organization that was developed mainly at the Institute for Robotics and Intelligent Systems of the University of Southern California. It is not written as a historical recount of the work, since the sequence of the presentation is not in chronological order. It aims at presenting an approach to a wide range of problems in computer vision and machine learning that is data-driven, local and requires a minimal number of assumptions. The tensor voting framework combines these properties and provides a unified perceptual organization methodology applicable in situations that may seem heterogeneous initially. Authors Philippos Mordohai and Gerard Medioni show how several problems can be posed as the organization of the inputs into salient perceptual structures, which are inferred via tensor voting. The book extends

the original tensor voting framework with the addition of boundary inference capabilities, a novel re-formulation of the framework applicable to high-dimensional spaces and the development of algorithms for computer vision and machine learning problems. The authors provide a complete analysis for some problems and briefly outline the approach for other applications and provide references to relevant sources.

Essentials of Game Theory - Kevin Leyton-Brown
2008-07-08

Game theory is the mathematical study of interaction among independent, self-interested agents. The audience for game theory has grown dramatically in recent years, and now spans disciplines as diverse as political science, biology, psychology, economics, linguistics, sociology, and computer science, among others. What has been missing is a relatively short introduction to the field covering the common basis that anyone with a professional interest in game theory is likely to require. Such a text

would minimize notation, ruthlessly focus on essentials, and yet not sacrifice rigor. This Synthesis Lecture aims to fill this gap by providing a concise and accessible introduction to the field. It covers the main classes of games, their representations, and the main concepts used to analyze them.

Active Learning - Burr Settles 2012-07-01

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

Machine Learning Algorithms for Signal and Image Processing - Suman Lata Tripathi
2022-12-01

Enables readers to understand the fundamental concepts of machine and deep learning techniques with interactive, real-life applications within signal and image processing Machine

Learning Algorithms for Signal and Image Processing aids the reader in designing and developing real-world applications using advances in machine learning to aid and enhance speech signal processing, image processing, computer vision, biomedical signal processing, adaptive filtering, and text processing. It includes signal processing techniques applied for pre-processing, feature extraction, source separation, or data decompositions to achieve machine learning tasks. Written by well-qualified authors and contributed to by a team of experts within the field, the work covers a wide range of important topics, such as: Speech recognition, image reconstruction, object classification and detection, and text processing Healthcare monitoring, biomedical systems, and green energy How various machine and deep learning techniques can improve accuracy, precision rate recall rate, and processing time Real applications and examples, including smart sign

language recognition, fake news detection in social media, structural damage prediction, and epileptic seizure detection Professionals within the field of signal and image processing seeking to adapt their work further will find immense value in this easy-to-understand yet extremely comprehensive reference work. It is also a worthy resource for students and researchers in related fields who are looking to thoroughly understand the historical and recent developments that have been made in the field.

Reinforcement Learning, second edition -
Richard S. Sutton 2018-11-13

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.

Markov Logic - Pedro Domingos 2009-05-08

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 viable 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.

Foundations of Deep Reinforcement Learning - Laura Graesser 2019-11-20

The Contemporary Introduction to Deep Reinforcement Learning that Combines Theory and Practice Deep reinforcement learning (deep RL) combines deep learning and reinforcement learning, in which artificial agents learn to solve sequential decision-making problems. In the past decade deep RL has achieved remarkable results on a range of problems, from single and multiplayer games—such as Go, Atari games,

and DotA 2—to robotics. Foundations of Deep Reinforcement Learning is an introduction to deep RL that uniquely combines both theory and implementation. It starts with intuition, then carefully explains the theory of deep RL algorithms, discusses implementations in its companion software library SLM Lab, and finishes with the practical details of getting deep RL to work. This guide is ideal for both computer science students and software engineers who are familiar with basic machine learning concepts and have a working understanding of Python. Understand each key aspect of a deep RL problem Explore policy- and value-based algorithms, including REINFORCE, SARSA, DQN, Double DQN, and Prioritized Experience Replay (PER) Delve into combined algorithms, including Actor-Critic and Proximal Policy Optimization (PPO) Understand how algorithms can be parallelized synchronously and asynchronously Run algorithms in SLM Lab and learn the practical implementation details for

getting deep RL to work Explore algorithm benchmark results with tuned hyperparameters Understand how deep RL environments are designed Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Multi-Objective Decision Making - Diederik M. Roijers 2017-04-20

Many real-world decision problems have multiple objectives. For example, when choosing a medical treatment plan, we want to maximize the efficacy of the treatment, but also minimize the side effects. These objectives typically conflict, e.g., we can often increase the efficacy of the treatment, but at the cost of more severe side effects. In this book, we outline how to deal with multiple objectives in decision-theoretic planning and reinforcement learning algorithms. To illustrate this, we employ the popular problem classes of multi-objective Markov decision processes (MOMDPs) and multi-

objective coordination graphs (MO-CoGs). First, we discuss different use cases for multi-objective decision making, and why they often necessitate explicitly multi-objective algorithms. We advocate a utility-based approach to multi-objective decision making, i.e., that what constitutes an optimal solution to a multi-objective decision problem should be derived from the available information about user utility. We show how different assumptions about user utility and what types of policies are allowed lead to different solution concepts, which we outline in a taxonomy of multi-objective decision problems. Second, we show how to create new methods for multi-objective decision making using existing single-objective methods as a basis. Focusing on planning, we describe two ways to creating multi-objective algorithms: in the inner loop approach, the inner workings of a single-objective method are adapted to work with multi-objective solution concepts; in the outer loop approach, a wrapper is created

around a single-objective method that solves the multi-objective problem as a series of single-objective problems. After discussing the creation of such methods for the planning setting, we discuss how these approaches apply to the learning setting. Next, we discuss three promising application domains for multi-objective decision making algorithms: energy, health, and infrastructure and transportation. Finally, we conclude by outlining important open problems and promising future directions.

Planning with Markov Decision Processes - Mausam 2012

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

General Game Playing - Michael Genesereth
2014-03-01

General game players are computer systems able to play strategy games based solely on formal game descriptions supplied at "runtime" (in other words, they don't know the rules until the game starts). Unlike specialized game players, such as Deep Blue, general game players cannot rely on algorithms designed in advance for specific games; they must discover such algorithms themselves. General game playing expertise depends on intelligence on the

part of the game player and not just intelligence of the programmer of the game player. GGP is an interesting application in its own right. It is intellectually engaging and more than a little fun. But it is much more than that. It provides a theoretical framework for modeling discrete dynamic systems and defining rationality in a way that takes into account problem representation and complexities like incompleteness of information and resource bounds. It has practical applications in areas where these features are important, e.g., in business and law. More fundamentally, it raises questions about the nature of intelligence and serves as a laboratory in which to evaluate competing approaches to artificial intelligence. This book is an elementary introduction to General Game Playing (GGP). (1) It presents the theory of General Game Playing and leading GGP technologies. (2) It shows how to create GGP programs capable of competing against other programs and humans. (3) It offers a

glimpse of some of the real-world applications of General Game Playing. Table of Contents: Preface / Introduction / Game Description / Game Management / Game Playing / Small Single-Player Games / Small Multiple-Player Games / Heuristic Search / Probabilistic Search / Propositional Nets / General Game Playing With Propnets / Factoring / Discovery of Heuristics / Logic / Analyzing Games with Logic / Solving Single-Player Games with Logic / Discovering Heuristics with Logic / Games with Incomplete Information / Games with Historical Constraints / Incomplete Game Descriptions / Advanced General Game Playing / Authors' Biographies

Reasoning With Probabilistic and Deterministic Graphical Models - Rina Dechter 2019-02-14

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.

Algorithms for Reinforcement Learning - Csaba Grossi 2022-05-31

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
Multi-Objective Decision Making - Diederik M. Roijers 2017-04-20

Many real-world decision problems have multiple objectives. For example, when choosing a medical treatment plan, we want to maximize the efficacy of the treatment, but also minimize the side effects. These objectives typically conflict, e.g., we can often increase the efficacy of the treatment, but at the cost of more severe side effects. In this book, we outline how to deal with multiple objectives in decision-theoretic planning and reinforcement learning algorithms. To illustrate this, we employ the popular problem classes of multi-objective Markov decision processes (MOMDPs) and multi-objective coordination graphs (MO-CoGs). First, we discuss different use cases for multi-objective decision making, and why they often necessitate explicitly multi-objective algorithms. We advocate a utility-based approach to multi-objective decision making, i.e., that what constitutes an optimal solution to a multi-objective decision problem should be derived from the available information about user utility.

We show how different assumptions about user utility and what types of policies are allowed lead to different solution concepts, which we outline in a taxonomy of multi-objective decision problems. Second, we show how to create new methods for multi-objective decision making using existing single-objective methods as a basis. Focusing on planning, we describe two ways to creating multi-objective algorithms: in the inner loop approach, the inner workings of a single-objective method are adapted to work with multi-objective solution concepts; in the outer loop approach, a wrapper is created around a single-objective method that solves the multi-objective problem as a series of single-objective problems. After discussing the creation of such methods for the planning setting, we discuss how these approaches apply to the learning setting. Next, we discuss three promising application domains for multi-objective decision making algorithms: energy, health, and infrastructure and transportation.

Finally, we conclude by outlining important open problems and promising future directions.

[Handbook of Reinforcement Learning and Control](#) - Kyriakos G. Vamvoudakis 2021-06-23

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.

Representations and Techniques for 3D Object Recognition and Scene

Interpretation - Derek Santhanam 2022-05-31

One of the grand challenges of artificial intelligence is to enable computers to interpret 3D scenes and objects from imagery. This book organizes and introduces major concepts in 3D scene and object representation and inference from still images, with a focus on recent efforts to fuse models of geometry and perspective with statistical machine learning. The book is organized into three sections: (1) Interpretation of Physical Space; (2) Recognition of 3D Objects; and (3) Integrated 3D Scene Interpretation. The first discusses representations of spatial layout and techniques to interpret physical scenes from images. The second section introduces representations for 3D object categories that

account for the intrinsically 3D nature of objects and provide robustness to change in viewpoints. The third section discusses strategies to unite inference of scene geometry and object pose and identity into a coherent scene interpretation. Each section broadly surveys important ideas from cognitive science and artificial intelligence research, organizes and discusses key concepts and techniques from recent work in computer vision, and describes a few sample approaches in detail. Newcomers to computer vision will benefit from introductions to basic concepts, such as single-view geometry and image classification, while experts and novices alike may find inspiration from the book's organization and discussion of the most recent ideas in 3D scene understanding and 3D object recognition. Specific topics include: mathematics of perspective geometry; visual elements of the physical scene, structural 3D scene representations; techniques and features for image and region categorization; historical

perspective, computational models, and datasets and machine learning techniques for 3D object recognition; inferences of geometrical attributes of objects, such as size and pose; and probabilistic and feature-passing approaches for contextual reasoning about 3D objects and scenes. Table of Contents: Background on 3D Scene Models / Single-view Geometry / Modeling the Physical Scene / Categorizing Images and Regions / Examples of 3D Scene Interpretation / Background on 3D Recognition / Modeling 3D Objects / Recognizing and Understanding 3D Objects / Examples of 2D 1/2 Layout Models / Reasoning about Objects and Scenes / Cascades of Classifiers / Conclusion and Future Directions

Deep Reinforcement Learning Hands-On - Maxim Lapan 2020-01-31

New edition of the bestselling guide to deep reinforcement learning and how it's used to solve complex real-world problems. Revised and expanded to include multi-agent methods, discrete optimization, RL in robotics, advanced

exploration techniques, and more Key Features
Second edition of the bestselling introduction to deep reinforcement learning, expanded with six new chapters Learn advanced exploration techniques including noisy networks, pseudo-count, and network distillation methods Apply RL methods to cheap hardware robotics platforms Book Description Deep Reinforcement Learning Hands-On, Second Edition is an updated and expanded version of the bestselling guide to the very latest reinforcement learning (RL) tools and techniques. It provides you with an introduction to the fundamentals of RL, along with the hands-on ability to code intelligent learning agents to perform a range of practical tasks. With six new chapters devoted to a variety of up-to-the-minute developments in RL, including discrete optimization (solving the Rubik's Cube), multi-agent methods, Microsoft's TextWorld environment, advanced exploration techniques, and more, you will come away from this book with a deep understanding of the latest

innovations in this emerging field. In addition, you will gain actionable insights into such topic areas as deep Q-networks, policy gradient methods, continuous control problems, and highly scalable, non-gradient methods. You will also discover how to build a real hardware robot trained with RL for less than \$100 and solve the Pong environment in just 30 minutes of training using step-by-step code optimization. In short, Deep Reinforcement Learning Hands-On, Second Edition, is your companion to navigating the exciting complexities of RL as it helps you attain experience and knowledge through real-world examples. What you will learn Understand the deep learning context of RL and implement complex deep learning models Evaluate RL methods including cross-entropy, DQN, actor-critic, TRPO, PPO, DDPG, D4PG, and others Build a practical hardware robot trained with RL methods for less than \$100 Discover Microsoft's TextWorld environment, which is an interactive fiction games platform Use discrete optimization

in RL to solve a Rubik's Cube Teach your agent to play Connect 4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI chatbots Discover advanced exploration techniques, including noisy networks and network distillation techniques Who this book is for Some fluency in Python is assumed. Sound understanding of the fundamentals of deep learning will be helpful. This book is an introduction to deep RL and requires no background in RL

Reinforcement Learning - Marco Wiering
2012-03-05

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 subfields 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.

Machine and Deep Learning Algorithms and Applications - Uday Shankar 2022-05-31

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.

Human Computation - Edith Law 2011-08-08

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

Lifelong Machine Learning - Zhiyuan Chen

2018-08-14

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.

Deep Learning for Computer Architects -

Brandon Reagen 2017-08-22

This is a primer written for computer architects in the new and rapidly evolving field of deep learning. It reviews how machine learning has evolved since its inception in the 1960s and tracks the key developments leading up to the emergence of the powerful deep learning techniques that emerged in the last decade. Machine learning, and specifically deep learning, has been hugely disruptive in many fields of computer science. The success of deep learning techniques in solving notoriously difficult classification and regression problems has resulted in their rapid adoption in solving real-world problems. The emergence of deep learning is widely attributed to a virtuous cycle whereby fundamental advancements in training deeper models were enabled by the availability of massive datasets and high-performance computer hardware. It also reviews representative workloads, including the most commonly used datasets and seminal networks across a variety of domains. In addition to

discussing the workloads themselves, it also details the most popular deep learning tools and show how aspiring practitioners can use the tools with the workloads to characterize and optimize DNNs. The remainder of the book is dedicated to the design and optimization of hardware and architectures for machine learning. As high-performance hardware was so instrumental in the success of machine learning becoming a practical solution, this chapter recounts a variety of optimizations proposed recently to further improve future designs. Finally, it presents a review of recent research published in the area as well as a taxonomy to help readers understand how various contributions fall in context.

Introduction to Semi-supervised Learning -
Xiaojin Zhu 2009

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
Abstract Dynamic Programming - Dimitri Bertsekas 2022-01-01

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).

Learning and Decision-Making from Rank Data - Lirong Xia 2019-02-06

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.

Federated Learning - Qiang Yang 2020-11-25

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.”

Algorithms for Reinforcement Learning -

Csaba Szepesvari 2010

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.

Foundations of Machine Learning, second edition - Mehryar Mohri 2018-12-25

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.

Visual Object Recognition - Kristen Grauman 2011

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

reliably today, and overviews key concepts that could be employed in systems requiring visual categorization.

Bandit Algorithms - Tor Lattimore 2020-07-16

A comprehensive and rigorous introduction for graduate students and researchers, with applications in sequential decision-making problems.

Neural Network Methods in Natural Language Processing - Yoav Goldberg

2017-04-17

Neural networks are a family of powerful machine learning models and this book focuses on their application to natural language data. The first half of the book (Parts I and II) covers the basics of supervised machine learning and feed-forward neural networks, the basics of working with machine learning over language data, and the use of vector-based rather than symbolic representations for words. It also covers the computation-graph abstraction, which allows to easily define and train arbitrary neural

networks, and is the basis behind the design of contemporary neural network software libraries. The second part of the book (Parts III and IV) introduces more specialized neural network architectures, including 1D convolutional neural networks, recurrent neural networks, conditioned-generation models, and attention-based models. These architectures and techniques are the driving force behind state-of-the-art algorithms for machine translation, syntactic parsing, and many other applications. Finally, we also discuss tree-shaped networks, structured prediction, and the prospects of multi-task learning.

Graph Representation Learning - William L. Hamilton 2020-09-16

This book is a foundational guide to graph representation learning, including state-of-the-art advances, and introduces the highly successful graph neural network (GNN) formalism. 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. 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.

Reinforcement Learning and Optimal Control -
Dimitri Bertsekas 2019-07-01

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.

Graph Machine Learning - Claudio Stamile
2021-06-25

Build machine learning algorithms using graph data and efficiently exploit topological information within your models
Key Features
Implement machine learning techniques and algorithms in graph data
Identify the relationship between nodes in order to make better business decisions
Apply graph-based machine learning methods to solve real-life problems
Book Description
Graph Machine Learning will introduce you to a set of tools used

for processing network data and leveraging the power of the relation between entities that can be used for predictive, modeling, and analytics tasks. The first chapters will introduce you to graph theory and graph machine learning, as well as the scope of their potential use. You'll then learn all you need to know about the main machine learning models for graph representation learning: their purpose, how they work, and how they can be implemented in a wide range of supervised and unsupervised learning applications. You'll build a complete machine learning pipeline, including data processing, model training, and prediction in order to exploit the full potential of graph data. After covering the basics, you'll be taken through real-world scenarios such as extracting data from social networks, text analytics, and natural language processing (NLP) using graphs and financial transaction systems on graphs. You'll also learn how to build and scale out data-driven applications for graph analytics to store,

query, and process network information, and explore the latest trends on graphs. By the end of this machine learning book, you will have learned essential concepts of graph theory and all the algorithms and techniques used to build successful machine learning applications. What you will learn

Write Python scripts to extract features from graphs

Distinguish between the main graph representation learning techniques

Learn how to extract data from social networks, financial transaction systems, for text analysis, and more

Implement the main unsupervised and supervised graph embedding techniques

Get to grips with shallow embedding methods, graph neural networks, graph regularization methods, and more

Deploy and scale out your application seamlessly

Who this book is for

This book is for data scientists, data analysts, graph analysts, and graph professionals who want to leverage the information embedded in the connections and relations between data points to boost their

analysis and model performance using machine learning. It will also be useful for machine learning developers or anyone who wants to build ML-driven graph databases. A beginner-level understanding of graph databases and graph data is required, alongside a solid understanding of ML basics. You'll also need intermediate-level Python programming knowledge to get started with this book.

Automated Machine Learning - Frank Hutter
2019-05-17

This open access book presents the first comprehensive overview of general methods in Automated Machine Learning (AutoML), collects descriptions of existing systems based on these methods, and discusses the first series of international challenges of AutoML systems. The recent success of commercial ML applications and the rapid growth of the field has created a high demand for off-the-shelf ML methods that can be used easily and without expert knowledge. However, many of the recent

machine learning successes crucially rely on human experts, who manually select appropriate ML architectures (deep learning architectures or more traditional ML workflows) and their hyperparameters. To overcome this problem, the field of AutoML targets a progressive automation of machine learning, based on principles from optimization and machine learning itself. This book serves as a point of entry into this quickly-developing field for researchers and advanced students alike, as well as providing a reference for practitioners aiming to use AutoML in their work.

Deep Learning Systems - Andres Rodriguez
2020-10-26

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 co-design 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.

Data Integration - Michael Genesereth 2010
Data integration is a critical problem in our increasingly interconnected but inevitably heterogeneous world. There are numerous data sources available in organizational databases and on public information systems like the World Wide Web. Not surprisingly, the sources often use different vocabularies and different data structures, being created, as they are, by different people, at different times, for different

purposes. The goal of data integration is to provide programmatic and human users with integrated access to multiple, heterogeneous data sources, giving each user the illusion of a single, homogeneous database designed for his or her specific need. The good news is that, in many cases, the data integration process can be automated. This book is an introduction to the problem of data integration and a rigorous account of one of the leading approaches to solving this problem, viz., the relational logic approach. Relational logic provides a theoretical framework for discussing data integration. Moreover, in many important cases, it provides algorithms for solving the problem in a computationally practical way. In many respects, relational logic does for data integration what relational algebra did for database theory several decades ago. A companion web site provides interactive demonstrations of the algorithms. Table of Contents: Preface / Interactive Edition / Introduction / Basic

Concepts / Query Folding / Query Planning /
Master Schema Management / Appendix /
References / Index / Author Biography

Insights in Reinforcement Learning - Hado
Philip van Hasselt 2011

Algorithms for Reinforcement Learning - Csaba
Szepesvari 2010-08-08

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

Introduction to Graph Neural Networks -
Zhiyuan Liu 2020-03-20

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. 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. 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.