

Reinforcement Learning An Introduction

Yeah, reviewing a book **Reinforcement Learning An Introduction** could add your close links listings. This is just one of the solutions for you to be successful. As understood, achievement does not recommend that you have fabulous points.

Comprehending as without difficulty as harmony even more than new will pay for each success. next to, the broadcast as with ease as sharpness of this Reinforcement Learning An Introduction can be taken as competently as picked to act.

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.

[Applied Reinforcement Learning with Python](#) - Taweh Beysolow II
2019-08-23

Delve into the world of reinforcement learning algorithms and apply them to different use-cases via Python. This book covers important topics such as policy gradients and Q learning, and utilizes frameworks such as

Tensorflow, Keras, and OpenAI Gym. Applied Reinforcement Learning with Python introduces you to the theory behind reinforcement learning (RL) algorithms and the code that will be used to implement them. You will take a guided tour through features of OpenAI Gym, from utilizing standard libraries to creating your own environments, then discover how to frame reinforcement learning problems so you can research, develop, and deploy RL-based solutions. What You'll Learn Implement reinforcement learning with Python Work with AI frameworks such as OpenAI Gym, Tensorflow, and KerasDeploy and train reinforcement learning-based solutions via cloud resourcesApply practical applications of reinforcement learning Who This Book Is For Data scientists, machine learning engineers and software engineers familiar with machine learning and deep learning concepts.

An Introduction to Machine Learning - Gopinath Rebala 2019-05-07

Just like electricity, Machine Learning will revolutionize our life in many ways - some of which are not even conceivable today. This book provides a thorough conceptual understanding of Machine Learning techniques and algorithms. Many of the mathematical concepts are explained in an intuitive manner. The book starts with an overview of machine learning and the underlying Mathematical and Statistical concepts before moving onto machine learning topics. It gradually builds up the depth, covering many of the present day machine learning algorithms, ending in Deep Learning and Reinforcement Learning algorithms. The book also covers

some of the popular Machine Learning applications. The material in this book is agnostic to any specific programming language or hardware so that readers can try these concepts on whichever platforms they are already familiar with. Offers a comprehensive introduction to Machine Learning, while not assuming any prior knowledge of the topic; Provides a complete overview of available techniques and algorithms in conceptual terms, covering various application domains of machine learning; Not tied to any specific software language or hardware implementation.

Mastering Reinforcement Learning with Python - Enes Bilgin
2020-12-18

Get hands-on experience in creating state-of-the-art reinforcement learning agents using TensorFlow and RLlib to solve complex real-world business and industry problems with the help of expert tips and best practices

Key Features Understand how large-scale state-of-the-art RL algorithms and approaches work Apply RL to solve complex problems in marketing, robotics, supply chain, finance, cybersecurity, and more Explore tips and best practices from experts that will enable you to overcome real-world RL challenges

Book Description Reinforcement learning (RL) is a field of artificial intelligence (AI) used for creating self-learning autonomous agents. Building on a strong theoretical foundation, this book takes a practical approach and uses examples inspired by real-world industry problems to teach you about state-of-the-art RL. Starting with bandit problems, Markov decision processes, and dynamic programming, the book provides an in-depth review of the classical RL techniques, such as Monte Carlo methods and temporal-difference learning. After that, you will learn about deep Q-learning, policy gradient algorithms, actor-critic methods, model-based methods, and multi-agent reinforcement learning. Then, you'll be introduced to some of the key approaches behind the most successful RL implementations, such as domain randomization and curiosity-driven learning. As you advance, you'll explore many novel algorithms with advanced implementations using modern Python libraries such as TensorFlow and Ray's RLlib package. You'll also find out how to implement RL in areas such as

robotics, supply chain management, marketing, finance, smart cities, and cybersecurity while assessing the trade-offs between different approaches and avoiding common pitfalls. By the end of this book, you'll have mastered how to train and deploy your own RL agents for solving RL problems. What you will learn

Model and solve complex sequential decision-making problems using RL Develop a solid understanding of how state-of-the-art RL methods work Use Python and TensorFlow to code RL algorithms from scratch Parallelize and scale up your RL implementations using Ray's RLlib package Get in-depth knowledge of a wide variety of RL topics Understand the trade-offs between different RL approaches Discover and address the challenges of implementing RL in the real world

Who this book is for This book is for expert machine learning practitioners and researchers looking to focus on hands-on reinforcement learning with Python by implementing advanced deep reinforcement learning concepts in real-world projects. Reinforcement learning experts who want to advance their knowledge to tackle large-scale and complex sequential decision-making problems will also find this book useful. Working knowledge of Python programming and deep learning along with prior experience in reinforcement learning is required.

[Learning to Play](#) - Aske Plaat 2020-12-23

In this textbook the author takes as inspiration recent breakthroughs in game playing to explain how and why deep reinforcement learning works. In particular he shows why two-person games of tactics and strategy fascinate scientists, programmers, and game enthusiasts and unite them in a common goal: to create artificial intelligence (AI). After an introduction to the core concepts, environment, and communities of intelligence and games, the book is organized into chapters on reinforcement learning, heuristic planning, adaptive sampling, function approximation, and self-play. The author takes a hands-on approach throughout, with Python code examples and exercises that help the reader understand how AI learns to play. He also supports the main text with detailed pointers to online machine learning frameworks, technical details for AlphaGo, notes on how to play and program Go and chess, and

a comprehensive bibliography. The content is class-tested and suitable for advanced undergraduate and graduate courses on artificial intelligence and games. It's also appropriate for self-study by professionals engaged with applications of machine learning and with games development. Finally it's valuable for any reader engaged with the philosophical implications of artificial and general intelligence, games represent a modern Turing test of the power and limitations of AI.

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.

An Introduction to Deep Reinforcement Learning - Vincent Francois-Lavet 2018-12-20

Deep reinforcement learning is the combination of reinforcement learning (RL) and deep learning. This field of research has recently been able to solve a wide range of complex decision-making tasks that were

previously out of reach for a machine. Deep RL opens up many new applications in domains such as healthcare, robotics, smart grids, finance, and many more. This book provides the reader with a starting point for understanding the topic. Although written at a research level it provides a comprehensive and accessible introduction to deep reinforcement learning models, algorithms and techniques. Particular focus is on the aspects related to generalization and how deep RL can be used for practical applications. Written by recognized experts, this book is an important introduction to Deep Reinforcement Learning for practitioners, researchers and students alike.

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.

Mathematics for Machine Learning - Marc Peter Deisenroth

2020-04-23

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Transfer in Reinforcement Learning Domains - Matthew Taylor

2009-06-05

In reinforcement learning (RL) problems, learning agents sequentially execute actions with the goal of maximizing a reward signal. The RL framework has gained popularity with the development of algorithms capable of mastering increasingly complex problems, but learning difficult tasks is often slow or infeasible when RL agents begin with no prior knowledge. The key insight behind "transfer learning" is that generalization may occur not only within tasks, but also across tasks. While transfer has been studied in the psychological literature for many years, the RL community has only recently begun to investigate the benefits of transferring knowledge. This book provides an introduction to the RL transfer problem and discusses methods which demonstrate the promise of this exciting area of research. The key contributions of this

book are: Definition of the transfer problem in RL domains Background on RL, sufficient to allow a wide audience to understand discussed transfer concepts Taxonomy for transfer methods in RL Survey of existing approaches In-depth presentation of selected transfer methods Discussion of key open questions By way of the research presented in this book, the author has established himself as the pre-eminent worldwide expert on transfer learning in sequential decision making tasks. A particular strength of the research is its very thorough and methodical empirical evaluation, which Matthew presents, motivates, and analyzes clearly in prose throughout the book. Whether this is your initial introduction to the concept of transfer learning, or whether you are a practitioner in the field looking for nuanced details, I trust that you will find this book to be an enjoyable and enlightening read. Peter Stone, Associate Professor of Computer Science

Grokking Deep Reinforcement Learning - Miguel Morales 2020-11-10

Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. Summary We all learn through trial and error. We avoid the things that cause us to experience pain and failure. We embrace and build on the things that give us reward and success. This common pattern is the foundation of deep reinforcement learning: building machine learning systems that explore and learn based on the responses of the environment. Grokking Deep Reinforcement Learning introduces this powerful machine learning approach, using examples, illustrations, exercises, and crystal-clear teaching. You'll love the perfectly paced teaching and the clever, engaging writing style as you dig into this awesome exploration of reinforcement learning fundamentals, effective deep learning techniques, and practical applications in this emerging field. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology We learn by interacting with our environment, and the rewards or punishments we experience guide our future behavior.

Deep reinforcement learning brings that same natural process to artificial intelligence, analyzing results to uncover the most efficient ways forward. DRL agents can improve marketing campaigns, predict stock performance, and beat grand masters in Go and chess. About the book Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. What's inside An introduction to reinforcement learning DRL agents with human-like behaviors Applying DRL to complex situations About the reader For developers with basic deep learning experience. About the author Miguel Morales works on reinforcement learning at Lockheed Martin and is an instructor for the Georgia Institute of Technology's Reinforcement Learning and Decision Making course. Table of Contents 1 Introduction to deep reinforcement learning 2 Mathematical foundations of reinforcement learning 3 Balancing immediate and long-term goals 4 Balancing the gathering and use of information 5 Evaluating agents' behaviors 6 Improving agents' behaviors 7 Achieving goals more effectively and efficiently 8 Introduction to value-based deep reinforcement learning 9 More stable value-based methods 10 Sample-efficient value-based methods 11 Policy-gradient and actor-critic methods 12 Advanced actor-critic methods 13 Toward artificial general intelligence

Reinforcement Learning - Phil Winder Ph.D. 2020-11-06

Reinforcement learning (RL) will deliver one of the biggest breakthroughs in AI over the next decade, enabling algorithms to learn from their environment to achieve arbitrary goals. This exciting development avoids constraints found in traditional machine learning (ML) algorithms. This practical book shows data science and AI professionals how to learn by reinforcement and enable a machine to learn by itself. Author Phil Winder of Winder Research covers everything from basic building blocks to state-of-the-art practices. You'll explore the current state of RL, focus on industrial applications, learn numerous

algorithms, and benefit from dedicated chapters on deploying RL solutions to production. This is no cookbook; doesn't shy away from math and expects familiarity with ML. Learn what RL is and how the algorithms help solve problems Become grounded in RL fundamentals including Markov decision processes, dynamic programming, and temporal difference learning Dive deep into a range of value and policy gradient methods Apply advanced RL solutions such as meta learning, hierarchical learning, multi-agent, and imitation learning Understand cutting-edge deep RL algorithms including Rainbow, PPO, TD3, SAC, and more Get practical examples through the accompanying website

Statistical Reinforcement Learning - Masashi Sugiyama 2015-03-16
Reinforcement learning is a mathematical framework for developing computer agents that can learn an optimal behavior by relating generic reward signals with its past actions. With numerous successful applications in business intelligence, plant control, and gaming, the RL framework is ideal for decision making in unknown environments with large amo

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.

Reinforcement Learning and Stochastic Optimization - Warren B. Powell 2022-03-15

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 making 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 learning, 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 book, 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, computation, 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.

Artificial Intelligent Techniques for Wireless Communication and Networking - R. Kanthavel 2022-02-24

ARTIFICIAL INTELLIGENT TECHNIQUES FOR WIRELESS COMMUNICATION AND NETWORKING The 20 chapters address AI principles and techniques used in wireless communication and networking and outline their benefit, function, and future role in the field. Wireless communication and networking based on AI concepts and techniques are explored in this book, specifically focusing on the current research in the field by highlighting empirical results along with theoretical concepts. The possibility of applying AI mechanisms towards security aspects in the communication domain is elaborated; also explored is the application side of integrated technologies that enhance AI-based innovations, insights, intelligent predictions, cost optimization, inventory management, identification processes, classification mechanisms, cooperative spectrum sensing techniques, ad-hoc network architecture, and protocol and simulation-based environments. Audience Researchers, industry IT engineers, and graduate students working on and implementing AI-based wireless sensor networks, 5G, IoT, deep learning, reinforcement learning, and robotics in WSN, and related technologies.

Reinforcement Learning with TensorFlow - Sayon Dutta 2018-04-24
Leverage the power of the Reinforcement Learning techniques to develop self-learning systems using Tensorflow Key Features Learn reinforcement learning concepts and their implementation using TensorFlow Discover different problem-solving methods for Reinforcement Learning Apply reinforcement learning for autonomous driving cars, robo-brokers, and more Book Description Reinforcement Learning (RL), allows you to develop smart, quick and self-learning systems in your business surroundings. It is an effective method to train your learning agents and solve a variety of problems in Artificial Intelligence—from games, self-driving cars and robots to enterprise applications that range from datacenter energy saving (cooling data centers) to smart warehousing solutions. The book covers the major advancements and successes achieved in deep reinforcement learning by synergizing deep neural network architectures with reinforcement learning. The book also introduces readers to the concept of Reinforcement Learning, its advantages and why it’s gaining so much popularity. The book also discusses on MDPs, Monte Carlo tree searches, dynamic programming such as policy and value iteration, temporal difference learning such as Q-learning and SARSA. You will use TensorFlow and OpenAI Gym to build simple neural network models that learn from their own actions. You will also see how reinforcement learning algorithms play a role in games, image processing and NLP. By the end of this book, you will have a firm understanding of what reinforcement learning is and how to put your knowledge to practical use by leveraging the power of TensorFlow and OpenAI Gym. What you will learn Implement state-of-the-art Reinforcement Learning algorithms from the basics Discover various techniques of Reinforcement Learning such as MDP, Q Learning and more Learn the applications of Reinforcement Learning in advertisement, image processing, and NLP Teach a Reinforcement Learning model to play a game using TensorFlow and the OpenAI gym Understand how Reinforcement Learning Applications are used in robotics Who this book is for If you want to get started with reinforcement learning using TensorFlow in the most

practical way, this book will be a useful resource. The book assumes prior knowledge of machine learning and neural network programming concepts, as well as some understanding of the TensorFlow framework. No previous experience with Reinforcement Learning is required.

Reinforcement Learning and Dynamic Programming Using Function Approximators - Lucian Busoniu 2017-07-28

From household appliances to applications in robotics, engineered systems involving complex dynamics can only be as effective as the algorithms that control them. While Dynamic Programming (DP) has provided researchers with a way to optimally solve decision and control problems involving complex dynamic systems, its practical value was limited by algorithms that lacked the capacity to scale up to realistic problems. However, in recent years, dramatic developments in Reinforcement Learning (RL), the model-free counterpart of DP, changed our understanding of what is possible. Those developments led to the creation of reliable methods that can be applied even when a mathematical model of the system is unavailable, allowing researchers to solve challenging control problems in engineering, as well as in a variety of other disciplines, including economics, medicine, and artificial intelligence. *Reinforcement Learning and Dynamic Programming Using Function Approximators* provides a comprehensive and unparalleled exploration of the field of RL and DP. With a focus on continuous-variable problems, this seminal text details essential developments that have substantially altered the field over the past decade. In its pages, pioneering experts provide a concise introduction to classical RL and DP, followed by an extensive presentation of the state-of-the-art and novel methods in RL and DP with approximation. Combining algorithm development with theoretical guarantees, they elaborate on their work with illustrative examples and insightful comparisons. Three individual chapters are dedicated to representative algorithms from each of the major classes of techniques: value iteration, policy iteration, and policy search. The features and performance of these algorithms are highlighted in extensive experimental studies on a range of control applications. The recent development of applications involving complex

systems has led to a surge of interest in RL and DP methods and the subsequent need for a quality resource on the subject. For graduate students and others new to the field, this book offers a thorough introduction to both the basics and emerging methods. And for those researchers and practitioners working in the fields of optimal and adaptive control, machine learning, artificial intelligence, and operations research, this resource offers a combination of practical algorithms, theoretical analysis, and comprehensive examples that they will be able to adapt and apply to their own work. Access the authors' website at www.dcsc.tudelft.nl/rlbook/ for additional material, including computer code used in the studies and information concerning new developments.

Deep Learning - Ian Goodfellow 2016-11-10

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, *Deep Learning* is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering

such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

[Introduction to Deep Learning](#) - Eugene Charniak 2019-01-29

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Reinforcement Learning, second edition - Richard S. Sutton
1998-02-26

Richard Sutton and Andrew Barto provide a clear and simple account of the key ideas and algorithms of reinforcement learning. Their discussion

ranges from the history of the field's intellectual foundations to the most recent developments and applications. 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 when interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the key ideas and algorithms of reinforcement learning. Their discussion ranges from the history of the field's intellectual foundations to the most recent developments and applications. The only necessary mathematical background is familiarity with elementary concepts of probability. The book is divided into three parts. Part I defines the reinforcement learning problem in terms of Markov decision processes. Part II provides basic solution methods: dynamic programming, Monte Carlo methods, and temporal-difference learning. Part III presents a unified view of the solution methods and incorporates artificial neural networks, eligibility traces, and planning; the two final chapters present case studies and consider the future of reinforcement learning.

[Deep Learning for Coders with fastai and PyTorch](#) - Jeremy Howard
2020-06-29

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn

your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Introduction to Machine Learning - Ethem Alpaydin 2014-08-22

Introduction -- Supervised learning -- Bayesian decision theory --

Parametric methods -- Multivariate methods -- Dimensionality reduction -

- Clustering -- Nonparametric methods -- Decision trees -- Linear

discrimination -- Multilayer perceptrons -- Local models -- Kernel

machines -- Graphical models -- Brief contents -- Hidden markov models -

- Bayesian estimation -- Combining multiple learners -- Reinforcement

learning -- Design and analysis of machine learning experiments.

The International Conference on Deep Learning, Big Data and Blockchain (Deep-BDB 2021) - Irfan Awan 2021-08-07

The role of deep learning for the analysis and learning of massive amounts of data from all aspects of daily-life has dramatically changed over the last few years. It is increasingly helping uncover trends leading to great successes. This book includes a collection of research manuscripts presenting state-of-the-art work in the areas of deep learning, blockchain and big data. All the manuscripts included in this book have been peer-reviewed based on aspects of novelty, originality and rigour. The main topics covered in the book include machine learning and time series, blockchain technologies and applications, data security, deep learning, and Internet of Things.

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 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. *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.

AI Crash Course - Hadelin de Ponteves 2019-11-29

Unlock the power of artificial intelligence with top Udemy AI instructor Hadelin de Ponteves. Key Features

- Learn from friendly, plain English explanations and practical activities
- Put ideas into action with 5 hands-on projects that show step-by-step how to build intelligent software
- Use AI to win classic video games and construct a virtual self-driving car

Book Description: Welcome to the Robot World ... and start building intelligent software now! Through his best-selling video courses, Hadelin de Ponteves has taught hundreds of thousands of people to write AI software. Now, for the first time, his hands-on, energetic approach is available as a book. Starting with the basics before easing you into more complicated formulas and notation, *AI Crash Course* gives you everything you need to build AI systems with reinforcement learning and deep learning. Five full working projects put the ideas into action, showing step-by-step how to build intelligent software using the best and easiest

tools for AI programming, including Python, TensorFlow, Keras, and PyTorch. *AI Crash Course* teaches everyone to build an AI to work in their applications. Once you've read this book, you're only limited by your imagination. What you will learn

- Master the basics of AI without any previous experience
- Build fun projects, including a virtual-self-driving car and a robot warehouse worker
- Use AI to solve real-world business problems
- Learn how to code in Python
- Discover the 5 principles of reinforcement learning
- Create your own AI toolkit

Who this book is for: If you want to add AI to your skillset, this book is for you. It doesn't require data science or machine learning knowledge. Just maths basics (high school level).

Deep Reinforcement Learning with Python - Nimish Sanghi 2021-06-12

Deep reinforcement learning is a fast-growing discipline that is making a significant impact in fields of autonomous vehicles, robotics, healthcare, finance, and many more. This book covers deep reinforcement learning using deep-q learning and policy gradient models with coding exercise. You'll begin by reviewing the Markov decision processes, Bellman equations, and dynamic programming that form the core concepts and foundation of deep reinforcement learning. Next, you'll study model-free learning followed by function approximation using neural networks and deep learning. This is followed by various deep reinforcement learning algorithms such as deep q-networks, various flavors of actor-critic methods, and other policy-based methods. You'll also look at exploration vs exploitation dilemma, a key consideration in reinforcement learning algorithms, along with Monte Carlo tree search (MCTS), which played a key role in the success of AlphaGo. The final chapters conclude with deep reinforcement learning implementation using popular deep learning frameworks such as TensorFlow and PyTorch. In the end, you'll understand deep reinforcement learning along with deep q networks and policy gradient models implementation with TensorFlow, PyTorch, and Open AI Gym. What You'll Learn

- Examine deep reinforcement learning
- Implement deep learning algorithms using OpenAI's Gym environment
- Code your own game playing agents for Atari using actor-critic algorithms
- Apply best practices for model building and algorithm

training Who This Book Is For Machine learning developers and architects who want to stay ahead of the curve in the field of AI and deep learning.

Deep Learning Illustrated - Jon Krohn 2019-08-05

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move

forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Deep Learning and the Game of Go - Kevin Ferguson 2019-01-06

Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild

Learning by practice: reinforcement learning
Reinforcement learning with policy gradients
Reinforcement learning with value methods
Reinforcement learning with actor-critic methods
PART 3 - GREATER THAN THE SUM OF ITS PARTS
AlphaGo: Bringing it all together
AlphaGo Zero: Integrating tree search with reinforcement learning

An Introduction to Machine Learning - Miroslav Kubat 2021-09-25
This textbook offers a comprehensive introduction to Machine Learning techniques and algorithms. This Third Edition covers newer approaches that have become highly topical, including deep learning, and auto-encoding, introductory information about temporal learning and hidden Markov models, and a much more detailed treatment of reinforcement learning. The book is written in an easy-to-understand manner with many examples and pictures, and with a lot of practical advice and discussions of simple applications. The main topics include Bayesian classifiers, nearest-neighbor classifiers, linear and polynomial classifiers, decision trees, rule-induction programs, artificial neural networks, support vector machines, boosting algorithms, unsupervised learning (including Kohonen networks and auto-encoding), deep learning, reinforcement learning, temporal learning (including long short-term memory), hidden Markov models, and the genetic algorithm. Special attention is devoted to performance evaluation, statistical assessment, and to many practical issues ranging from feature selection and feature construction to bias, context, multi-label domains, and the problem of imbalanced classes.

Introduction to Artificial Intelligence - Wolfgang Ertel 2018-01-18
This accessible and engaging textbook presents a concise introduction to the exciting field of artificial intelligence (AI). The broad-ranging discussion covers the key subdisciplines within the field, describing practical algorithms and concrete applications in the areas of agents, logic, search, reasoning under uncertainty, machine learning, neural networks, and reinforcement learning. Fully revised and updated, this much-anticipated second edition also includes new material on deep learning. Topics and features: presents an application-focused and hands-on approach to learning, with supplementary teaching resources provided at an associated website; contains numerous study exercises

and solutions, highlighted examples, definitions, theorems, and illustrative cartoons; includes chapters on predicate logic, PROLOG, heuristic search, probabilistic reasoning, machine learning and data mining, neural networks and reinforcement learning; reports on developments in deep learning, including applications of neural networks to generate creative content such as text, music and art (NEW); examines performance evaluation of clustering algorithms, and presents two practical examples explaining Bayes' theorem and its relevance in everyday life (NEW); discusses search algorithms, analyzing the cycle check, explaining route planning for car navigation systems, and introducing Monte Carlo Tree Search (NEW); includes a section in the introduction on AI and society, discussing the implications of AI on topics such as employment and transportation (NEW). Ideal for foundation courses or modules on AI, this easy-to-read textbook offers an excellent overview of the field for students of computer science and other technical disciplines, requiring no more than a high-school level of knowledge of mathematics to understand the material.

Deep Reinforcement Learning - Mohit Sewak 2019-06-27
This book starts by presenting the basics of reinforcement learning using highly intuitive and easy-to-understand examples and applications, and then introduces the cutting-edge research advances that make reinforcement learning capable of out-performing most state-of-art systems, and even humans in a number of applications. The book not only equips readers with an understanding of multiple advanced and innovative algorithms, but also prepares them to implement systems such as those created by Google Deep Mind in actual code. This book is intended for readers who want to both understand and apply advanced concepts in a field that combines the best of two worlds - deep learning and reinforcement learning - to tap the potential of 'advanced artificial intelligence' for creating real-world applications and game-winning algorithms.

Deep Reinforcement Learning - Hao Dong 2020-06-29
Deep reinforcement learning (DRL) is the combination of reinforcement learning (RL) and deep learning. It has been able to solve a wide range of

complex decision-making tasks that were previously out of reach for a machine, and famously contributed to the success of AlphaGo. Furthermore, it opens up numerous new applications in domains such as healthcare, robotics, smart grids and finance. Divided into three main parts, this book provides a comprehensive and self-contained introduction to DRL. The first part introduces the foundations of deep learning, reinforcement learning (RL) and widely used deep RL methods and discusses their implementation. The second part covers selected DRL research topics, which are useful for those wanting to specialize in DRL research. To help readers gain a deep understanding of DRL and quickly apply the techniques in practice, the third part presents mass applications, such as the intelligent transportation system and learning to run, with detailed explanations. The book is intended for computer science students, both undergraduate and postgraduate, who would like to learn DRL from scratch, practice its implementation, and explore the research topics. It also appeals to engineers and practitioners who do not have strong machine learning background, but want to quickly understand how DRL works and use the techniques in their applications.

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.

Deep Reinforcement Learning Hands-On - Maxim Lapan 2018-06-21

This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. Key Features Explore deep reinforcement learning (RL), from the first principles to the latest algorithms Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms Keep up with the very latest industry developments, including AI-driven chatbots Book Description Recent developments in reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning Hands-On is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. What you will learn Understand the DL context of RL and implement complex DL

models Learn the foundation of RL: Markov decision processes Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others Discover how to deal with discrete and continuous action spaces in various environments Defeat Atari arcade games using the value iteration method Create your own OpenAI Gym environment to train a stock trading agent Teach your agent to play Connect4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI-driven chatbots Who this book is for Some fluency in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL.

[An Introduction to Statistical Learning](#) - Gareth James 2013-06-24

An Introduction to Statistical Learning provides an accessible overview of the field of statistical learning, an essential toolset for making sense of the vast and complex data sets that have emerged in fields ranging from biology to finance to marketing to astrophysics in the past twenty years. This book presents some of the most important modeling and prediction techniques, along with relevant applications. Topics include linear regression, classification, resampling methods, shrinkage approaches, tree-based methods, support vector machines, clustering, and more. Color graphics and real-world examples are used to illustrate the methods presented. Since the goal of this textbook is to facilitate the use of these statistical learning techniques by practitioners in science, industry, and other fields, each chapter contains a tutorial on implementing the analyses and methods presented in R, an extremely popular open source statistical software platform. Two of the authors co-wrote The Elements of Statistical Learning (Hastie, Tibshirani and Friedman, 2nd edition 2009), a popular reference book for statistics and machine learning researchers. An Introduction to Statistical Learning covers many of the same topics, but at a level accessible to a much broader audience. This book is targeted at statisticians and non-statisticians alike who wish to use cutting-edge statistical learning techniques to analyze their data. The text assumes only a previous

course in linear regression and no knowledge of matrix algebra.

Decision Making Under Uncertainty - Mykel J. Kochenderfer
2015-07-24

An introduction to decision making under uncertainty from a computational perspective, covering both theory and applications ranging from speech recognition to airborne collision avoidance. Many important problems involve decision making under uncertainty—that is, choosing actions based on often imperfect observations, with unknown outcomes. Designers of automated decision support systems must take into account the various sources of uncertainty while balancing the multiple objectives of the system. This book provides an introduction to the challenges of decision making under uncertainty from a computational perspective. It presents both the theory behind decision making models and algorithms and a collection of example applications that range from speech recognition to aircraft collision avoidance. Focusing on two methods for designing decision agents, planning and reinforcement learning, the book covers probabilistic models, introducing Bayesian networks as a graphical model that captures probabilistic relationships between variables; utility theory as a framework for understanding optimal decision making under uncertainty; Markov decision processes as a method for modeling sequential problems; model uncertainty; state uncertainty; and cooperative decision making involving multiple interacting agents. A series of applications shows how the theoretical concepts can be applied to systems for attribute-based person search, speech applications, collision avoidance, and unmanned aircraft persistent surveillance. Decision Making Under Uncertainty unifies research from different communities using consistent notation, and is accessible to students and researchers across engineering disciplines who have some prior exposure to probability theory and calculus. It can be used as a text for advanced undergraduate and graduate students in fields including computer science, aerospace and electrical engineering, and management science. It will also be a valuable professional reference for researchers in a variety of disciplines.

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 Reinforcement Learning in Action](#) - Alexander Zai 2020-04-28

Summary Humans learn best from feedback—we are encouraged to take actions that lead to positive results while deterred by decisions with negative consequences. This reinforcement process can be applied to computer programs allowing them to solve more complex problems that classical programming cannot. Deep Reinforcement Learning in Action teaches you the fundamental concepts and terminology of deep reinforcement learning, along with the practical skills and techniques you'll need to implement it into your own projects. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Deep reinforcement learning AI systems rapidly adapt to new environments, a vast improvement over standard neural networks. A DRL agent learns like people do, taking in raw data such as sensor input and refining its responses and predictions through trial and error. About the book Deep Reinforcement Learning in Action teaches you how to program AI agents that adapt and improve based on direct feedback from their environment. In this example-rich tutorial, you'll master foundational and advanced DRL techniques by taking on interesting challenges like navigating a maze and playing video games. Along the way, you'll work with core algorithms, including deep Q-networks and policy gradients, along with industry-standard tools like PyTorch and OpenAI Gym. What's inside Building and training DRL networks The most popular DRL algorithms for learning and problem solving Evolutionary algorithms for curiosity and multi-agent learning All examples available as Jupyter Notebooks About the reader For readers with intermediate skills in Python and deep learning. About the author Alexander Zai is a machine learning engineer at Amazon AI. Brandon Brown is a machine learning and data analysis blogger. Table of Contents PART 1 - FOUNDATIONS 1. What is reinforcement learning? 2. Modeling reinforcement learning problems: Markov decision processes 3. Predicting the best states and actions: Deep Q-networks 4. Learning to pick the best policy: Policy gradient methods 5. Tackling more complex problems with actor-critic methods PART 2 - ABOVE AND BEYOND 6. Alternative optimization methods: Evolutionary algorithms 7. Distributional DQN: Getting the full story

8. Curiosity-driven exploration 9. Multi-agent reinforcement learning 10. Interpretable reinforcement learning: Attention and relational models 11. In conclusion: A review and roadmap