

Read Book Tutorial Deep Reinforcement Learning Pdf For Free

Deep Reinforcement Learning Foundations of Deep Reinforcement Learning An Introduction to Deep Reinforcement Learning Deep Reinforcement Learning in Action Grokking Deep Reinforcement Learning Deep Reinforcement Learning Hands-On Deep Reinforcement Learning Deep Reinforcement Learning Hands-On Reinforcement Learning, second edition Grokking Deep Reinforcement Learning An Introduction to Deep Reinforcement Learning Deep Reinforcement Learning with Python Hands-On Reinforcement Learning with Python Deep Reinforcement Learning Hands-On - Second Edition Artificial Intelligence Deep Reinforcement Learning for Wireless Networks The Reinforcement Learning Workshop Reinforcement Learning Hands-On Deep Learning for Games Deep Reinforcement Learning with Python Mastering Reinforcement Learning with Python Python Reinforcement Learning Projects Deep Reinforcement Learning Algorithms for Reinforcement Learning Deep Reinforcement Learning with Guaranteed Performance TensorFlow 2 Reinforcement Learning Cookbook Learning to Play Advanced Deep Learning with Keras Practical Deep Reinforcement Learning with Python Grokking Deep Learning A Review of Recent Advancements in Deep Reinforcement Learning Reinforcement Learning Algorithms: Analysis and Applications Reinforcement Learning with TensorFlow Hands-On Q-Learning with Python Deep Reinforcement Learning: Emerging Trends in Macroeconomics and Future Prospects Deep Reinforcement Learning Hands-On TensorFlow Reinforcement Learning Quick Start Guide Deep Learning and the Game of Go Hands-On Deep Learning with R Reinforcement Learning Algorithms with Python

Deep Reinforcement Learning Apr 29 2023 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.

Hands-On Q-Learning with Python Jun 26 2020 Leverage the power of reward-based training for your deep learning models with Python Key Features Understand Q-learning algorithms to train neural networks using Markov Decision Process (MDP) Study practical deep reinforcement learning using Q-Networks Explore state-based unsupervised learning for machine learning models Book Description Q-learning is a machine learning algorithm used to solve optimization problems in artificial intelligence (AI). It is one of the most popular fields of study among AI researchers. This book starts off by introducing you to reinforcement learning and Q-learning, in addition to helping you get familiar with OpenAI Gym as well as libraries such as Keras and TensorFlow. A few chapters into the book, you will gain insights into model-free Q-learning and use deep Q-networks and double deep Q-networks to solve complex problems. This book will guide you in exploring use cases such as self-driving vehicles and OpenAI Gym's CartPole problem. You will also learn how to tune and optimize Q-networks and their hyperparameters. As you progress, you will understand the reinforcement learning approach to solving real-world problems. You will also explore how to use Q-learning and related algorithms in real-world applications such as scientific research. Toward the end, you'll gain a sense of what's in store for reinforcement learning. By the end of this book, you will be equipped with the skills you need to solve reinforcement learning problems using Q-learning algorithms with OpenAI Gym, Keras, and TensorFlow. What you will learn Explore the fundamentals of reinforcement learning and the state-action-reward process Understand Markov decision processes Get well versed with libraries such as Keras, and TensorFlow Create and deploy model-free learning and deep Q-learning agents with TensorFlow, Keras, and OpenAI Gym Choose and optimize a Q-Network's learning parameters and fine-tune its performance Discover real-world applications and use cases of Q-learning Who this book is for If you are a machine learning developer, engineer, or professional who wants to delve into the deep learning approach for a complex environment, then this is the book for you. Proficiency in Python programming and basic understanding of decision-making in reinforcement learning is assumed.

Deep Reinforcement Learning with Python May 18 2022 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.

An Introduction to Deep Reinforcement Learning Jun 19 2022 Deep reinforcement learning is the combination of reinforcement learning (RL) and deep learning. This field of research has been able to solve a wide range of complex decision-making tasks that were previously out of reach for a machine. Thus, deep RL opens up many new applications in domains such as healthcare, robotics, smart grids, finance, and many more. This manuscript provides an 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. We assume the reader is familiar with basic machine learning concepts.

TensorFlow Reinforcement Learning Quick Start Guide Mar 24 2020 Leverage the power of TensorFlow to Create powerful software agents that can self-learn to perform real-world tasks Key Features Explore efficient Reinforcement Learning algorithms and code them using TensorFlow and Python Train Reinforcement Learning agents for problems, ranging from computer games to autonomous driving. Formulate and devise selective algorithms and techniques in your applications in no time. Book Description Advances in reinforcement learning algorithms have made it possible to use them for optimal control in several different industrial applications. With this book, you will apply Reinforcement Learning to a range of problems, from computer games to autonomous driving. The book starts by introducing you to essential Reinforcement Learning concepts such as agents, environments, rewards, and advantage functions. You will also master the distinctions between on-policy and off-policy algorithms, as well as model-free and model-based algorithms. You will also learn about several Reinforcement Learning algorithms, such as SARSA, Deep Q-Networks (DQN), Deep Deterministic Policy Gradients (DDPG), Asynchronous Advantage Actor-Critic (A3C), Trust Region Policy Optimization (TRPO), and Proximal Policy Optimization (PPO). The book will also show you how to code these algorithms in TensorFlow and Python and apply them to solve computer games from OpenAI Gym. Finally, you will also learn how to train a car to drive autonomously in the Torcs racing car simulator. By the end of the book, you will be able to design, build, train, and evaluate feed-forward neural networks and convolutional neural networks. You will also have mastered coding state-of-the-art algorithms and also training agents for various control problems. What you will learn Understand the theory and concepts behind modern Reinforcement Learning algorithms Code state-of-the-art Reinforcement Learning algorithms with discrete or continuous actions Develop Reinforcement Learning algorithms and apply them to training agents to play computer games Explore DQN, DDQN, and Dueling architectures to play Atari's Breakout using TensorFlow Use A3C to play CartPole and LunarLander Train an agent to drive a car autonomously in a simulator Who this book is for Data scientists and AI developers who wish to quickly get started with training effective reinforcement learning models in TensorFlow will find this book very useful. Prior knowledge of machine learning and deep learning concepts (as well as exposure to Python programming) will be useful.

Advanced Deep Learning with Keras Jan 02 2021 A comprehensive guide to advanced deep learning techniques, including Autoencoders, GANs, VAEs, and Deep Reinforcement Learning, that drive today's most impressive AI results Key Features Explore the most advanced deep learning techniques that drive modern AI results Implement Deep Neural Networks, Autoencoders, GANs, VAEs, and Deep Reinforcement Learning A wide study of GANs, including Improved GANs, Cross-Domain GANs and Disentangled Representation GANs Book Description Recent developments in deep learning, including GANs, Variational Autoencoders, and Deep Reinforcement Learning, are creating impressive AI results in our news headlines - such as AlphaGo Zero beating world chess champions, and generative AI that can create art paintings that sell for over \$400k because they are so human-like. Advanced Deep Learning with Keras is a comprehensive guide to the advanced deep learning techniques available today, so you can create your own cutting-edge AI. Using Keras as an open-source deep learning library, you'll find hands-on projects throughout that show you how to create more effective AI with the latest techniques. The journey begins with an overview of MLPs, CNNs, and RNNs, which are the building blocks for the more advanced techniques in the book. You'll learn how to implement deep learning models with Keras and TensorFlow, and move forwards to advanced techniques, as you explore deep neural network architectures, including ResNet and DenseNet, and how to create Autoencoders. You then learn all about Generative Adversarial Networks (GANs), and how they can open new levels of AI performance. Variational AutoEncoders (VAEs) are implemented, and you'll see how GANs and VAEs have the generative power to synthesize data that can be extremely convincing to humans - a major stride forward for modern AI. To complete this set of advanced techniques, you'll learn how to implement Deep Reinforcement

Learning (DRL) such as Deep Q-Learning and Policy Gradient Methods, which are critical to many modern results in AI. What you will learn Cutting-edge techniques in human-like AI performance Implement advanced deep learning models using Keras The building blocks for advanced techniques - MLPs, CNNs, and RNNs Deep neural networks - ResNet and DenseNet Autoencoders and Variational AutoEncoders (VAEs) Generative Adversarial Networks (GANs) and creative AI techniques Disentangled Representation GANs, and Cross-Domain GANs Deep Reinforcement Learning (DRL) methods and implementation Produce industry-standard applications using OpenAI gym Deep Q-Learning and Policy Gradient Methods Who this book is for Some fluency with Python is assumed. As an advanced book, you'll be familiar with some machine learning approaches, and some practical experience with DL will be helpful. Knowledge of Keras or TensorFlow is not required but would be helpful.

Foundations of Deep Reinforcement Learning Mar 28 2023 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.

TensorFlow 2 Reinforcement Learning Cookbook Mar 04 2021 This cookbook will help you to gain a solid understanding of deep reinforcement learning (RL) algorithms with the help of concise, easy-to-follow implementations from scratch. You'll learn how to implement these algorithms with minimal code and develop AI applications to solve real-world and business problems using RL.

Deep Learning and the Game of Go Feb 21 2020 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

Deep Reinforcement Learning Jun 07 2021 Deep reinforcement learning has attracted considerable attention recently. Impressive results have been achieved in such diverse fields as autonomous driving, game playing, molecular recombination, and robotics. In all these fields, computer programs have taught themselves to understand problems that were previously considered to be very difficult. In the game of Go, the program AlphaGo has even learned to outmatch three of the world's leading players. Deep reinforcement learning takes its inspiration from the fields of biology and psychology. Biology has inspired the creation of artificial neural networks and deep learning, while psychology studies how animals and humans learn, and how subjects' desired behavior can be reinforced with positive and negative stimuli. When we see how reinforcement learning teaches a simulated robot to walk, we are reminded of how children learn, through playful exploration. Techniques that are inspired by biology and psychology work amazingly well in computers: animal behavior and the structure of the brain as new blueprints for science and engineering. In fact, computers truly seem to possess aspects of human behavior; as such, this field goes to the heart of the dream of artificial intelligence. These research advances have not gone unnoticed by educators. Many universities have begun offering courses on the subject of deep reinforcement learning. The aim of this book is to provide an overview of the field, at the proper level of detail for a graduate course in artificial intelligence. It covers the complete field, from the basic algorithms of Deep Q-learning, to advanced topics such as multi-agent reinforcement learning and meta learning.

Grokking Deep Reinforcement Learning Dec 25 2022 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 Algorithms with Python Dec 21 2019 Develop self-learning algorithms and agents using TensorFlow and other Python tools, frameworks, and libraries Key Features Learn, develop, and deploy advanced reinforcement learning algorithms to solve a variety of tasks Understand and develop model-free and model-based algorithms for building self-learning agents Work with advanced Reinforcement Learning concepts and algorithms such as imitation learning and evolution strategies Book Description Reinforcement Learning (RL) is a popular and promising branch of AI that involves making smarter models and agents that can automatically determine ideal behavior based on changing requirements. This book will help you master RL algorithms and understand their implementation as you build self-learning agents. Starting with an introduction to the tools, libraries, and setup needed to work in the RL environment, this book covers the building blocks of RL and delves into value-based methods, such as the application of Q-learning and SARSA algorithms. You'll learn how to use a combination of Q-learning and neural networks to solve complex problems. Furthermore, you'll study the policy gradient methods, TRPO, and PPO, to improve performance and stability, before moving on to the DDPG and TD3 deterministic algorithms. This book also covers how imitation learning techniques work and how Dagger can teach an agent to drive. You'll discover evolutionary strategies and black-box optimization techniques, and see how they can improve RL algorithms. Finally, you'll get to grips with exploration approaches, such as UCB and UCB1, and develop a meta-algorithm called ESBAS. By the end of the book, you'll have worked with key RL algorithms to overcome challenges in real-world applications, and be part of the RL research community. What you will learn Develop an agent to play CartPole using the OpenAI Gym interface Discover the model-based reinforcement learning paradigm Solve the Frozen Lake problem with dynamic programming Explore Q-learning and SARSA with a view to playing a taxi game Apply Deep Q-Networks (DQNs) to Atari games using Gym Study policy gradient algorithms, including Actor-Critic and REINFORCE Understand and apply PPO and TRPO in continuous locomotion environments Get to grips with evolution strategies for solving the lunar lander problem Who this book is for If you are an AI researcher, deep learning user, or anyone who wants to learn reinforcement learning from scratch, this book is for you. You'll also find this reinforcement learning book useful if you want to learn about the advancements in the field. Working knowledge of Python is necessary.

Deep Reinforcement Learning for Wireless Networks Jan 14 2022 This Springerbrief presents a deep reinforcement learning approach to wireless systems to improve system performance. Particularly, deep reinforcement learning approach is used in cache-enabled opportunistic interference alignment wireless networks and mobile social networks. Simulation results with different network parameters are presented to show the effectiveness of the proposed scheme. There is a phenomenal burst of research activities in artificial intelligence, deep reinforcement learning and wireless systems. Deep reinforcement learning has been successfully used to solve many practical problems. For example, Google DeepMind adopts this method on several artificial intelligent projects with big data (e.g., AlphaGo), and gets quite good results.. Graduate students in

electrical and computer engineering, as well as computer science will find this brief useful as a study guide. Researchers, engineers, computer scientists, programmers, and policy makers will also find this brief to be a useful tool.

Deep Reinforcement Learning in Action Jan 26 2023 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

Hands-On Deep Learning for Games Oct 11 2021 Understand the core concepts of deep learning and deep reinforcement learning by applying them to develop games Key Features Apply the power of deep learning to complex reasoning tasks by building a Game AI Exploit the most recent developments in machine learning and AI for building smart games Implement deep learning models and neural networks with Python Book Description The number of applications of deep learning and neural networks has multiplied in the last couple of years. Neural nets has enabled significant breakthroughs in everything from computer vision, voice generation, voice recognition and self-driving cars. Game development is also a key area where these techniques are being applied. This book will give an in depth view of the potential of deep learning and neural networks in game development. We will take a look at the foundations of multi-layer perceptron's to using convolutional and recurrent networks. In applications from GANs that create music or textures to self-driving cars and chatbots. Then we introduce deep reinforcement learning through the multi-armed bandit problem and other OpenAI Gym environments. As we progress through the book we will gain insights about DRL techniques such as Motivated Reinforcement Learning with Curiosity and Curriculum Learning. We also take a closer look at deep reinforcement learning and in particular the Unity ML-Agents toolkit. By the end of the book, we will look at how to apply DRL and the ML-Agents toolkit to enhance, test and automate your games or simulations. Finally, we will cover your possible next steps and possible areas for future learning. What you will learn Learn the foundations of neural networks and deep learning. Use advanced neural network architectures in applications to create music, textures, self-driving cars and chatbots. Understand the basics of reinforcement and DRL and how to apply it to solve a variety of problems. Working with Unity ML-Agents toolkit and how to install, setup and run the kit. Understand core concepts of DRL and the differences between discrete and continuous action environments. Use several advanced forms of learning in various scenarios from developing agents to testing games. Who this book is for This book is for game developers who wish to create highly interactive games by leveraging the power of machine and deep learning. No prior knowledge of machine learning, deep learning or neural networks is required this book will teach those concepts from scratch. A good understanding of Python is required.

Reinforcement Learning Algorithms: Analysis and Applications Aug 29 2020 This book reviews research developments in diverse areas of reinforcement learning such as model-free actor-critic methods, model-based learning and control, information geometry of policy searches, reward design, and exploration in biology and the behavioral sciences. Special emphasis is placed on advanced ideas, algorithms, methods, and applications. The contributed papers gathered here grew out of a lecture course on reinforcement learning held by Prof. Jan Peters in the winter semester 2018/2019 at Technische Universität Darmstadt. The book is intended for reinforcement learning students and researchers with a firm grasp of linear algebra, statistics, and optimization. Nevertheless, all key concepts are introduced in each chapter, making the content self-contained and accessible to a broader audience.

Grokking Deep Learning Oct 31 2020 Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

Deep Reinforcement Learning Hands-On Sep 22 2022 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

Deep Reinforcement Learning Hands-On Nov 24 2022 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.

Artificial Intelligence Feb 15 2022 Artificial Intelligence: A Modern Approach offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. Number one in its field, this textbook is ideal for one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence.

Deep Reinforcement Learning with Python Sep 10 2021 An example-rich guide for beginners to start their reinforcement and deep reinforcement learning journey with state-of-the-art distinct algorithms Key FeaturesCovers a vast spectrum of basic-to-advanced RL algorithms with mathematical explanations of each algorithmLearn how to implement algorithms with code by following examples with line-by-line explanationsExplore the latest RL methodologies such as DDPG, PPO, and the use of expert demonstrationsBook Description With significant enhancements in the quality and quantity of algorithms in recent years, this second edition of Hands-On Reinforcement Learning with Python has been revamped into an example-rich guide to learning state-of-the-art reinforcement learning (RL) and deep RL algorithms with TensorFlow 2 and the OpenAI Gym toolkit. In addition to exploring RL basics and foundational concepts such as Bellman equation, Markov decision processes, and dynamic programming algorithms, this second edition dives deep into the full spectrum of value-based, policy-based, and actor-critic RL methods. It explores state-of-the-art algorithms such as DQN, TRPO, PPO and ACKTR, DDPG, TD3, and SAC in depth, demystifying the underlying math and demonstrating implementations through simple code examples. The book has several new chapters dedicated to new RL techniques, including distributional RL, imitation learning, inverse RL, and meta RL. You will learn to leverage stable baselines, an improvement of OpenAI's baseline library, to effortlessly implement popular RL algorithms. The book concludes with an overview of promising approaches such as meta-learning and imagination augmented agents in research. By the end, you will become skilled in effectively employing RL and deep RL in your real-world projects. What you will learnUnderstand core RL concepts including the methodologies, math, and codeTrain an agent to solve Blackjack, FrozenLake, and many other problems using OpenAI GymTrain an agent to play Ms Pac-Man using a Deep Q NetworkLearn policy-based, value-based, and actor-critic methodsMaster the math behind DDPG, TD3, TRPO, PPO, and many othersExplore new avenues such as the distributional RL, meta RL, and inverse RLUse Stable Baselines to train an agent to walk and play Atari gamesWho this book is for If you're a machine learning developer with little or no experience with neural networks interested in artificial intelligence and want to learn about reinforcement learning from scratch, this book is for you. Basic familiarity with linear algebra, calculus, and the Python programming language is required. Some experience with TensorFlow would be a plus.

Mastering Reinforcement Learning with Python Aug 09 2021 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 FeaturesUnderstand how large-scale state-of-the-art RL algorithms and approaches workApply RL to solve complex problems in marketing, robotics, supply chain, finance, cybersecurity, and moreExplore tips and best practices from experts that will enable you to overcome real-world RL challengesBook 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 learnModel and solve complex sequential decision-making problems using RLDevelop a solid understanding of how state-of-the-art RL methods workUse Python and TensorFlow to code RL algorithms from scratchParallelize and scale up your RL implementations using Ray's RLlib packageGet in-depth knowledge of a wide variety of RL topicsUnderstand the trade-offs between different RL approachesDiscover and address the challenges of implementing RL in the real worldWho 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.

Reinforcement Learning, second edition Aug 21 2022 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 with Guaranteed Performance Apr 05 2021 This book discusses methods and algorithms for the near-optimal adaptive control of nonlinear systems, including the corresponding theoretical analysis and simulative examples, and presents two innovative methods for the redundancy resolution of redundant manipulators with consideration of parameter uncertainty and periodic disturbances. It also reports on a series of systematic investigations on a near-optimal adaptive control method based on the Taylor expansion, neural networks, estimator design approaches, and the idea of sliding mode control, focusing on the tracking control problem of nonlinear systems under different scenarios. The book culminates with a presentation of two new redundancy resolution methods; one addresses adaptive kinematic control of redundant manipulators, and the other centers on the effect of periodic input disturbance on redundancy resolution. Each self-contained chapter is clearly written, making the book accessible to graduate students as well as academic and industrial researchers in the fields of adaptive and optimal control, robotics, and dynamic neural networks.

A Review of Recent Advancements in Deep Reinforcement Learning Sep 29 2020 Bachelor Thesis from the year 2018 in the subject Computer Science - Commercial Information Technology, grade: 1.0, University of Duisburg-Essen, language: English, abstract: Reinforcement learning is a learning problem in which an actor has to behave optimally in its environment. Deep learning methods, on the other hand, are a subclass of representation learning, which in turn focuses on extracting the necessary features for the task (e.g. classification or detection). As such, they serve as powerful function approximators. The combination of those two paradigm results in deep reinforcement learning. This thesis gives an overview of the recent advancement in the field. The results are divided into two broad research directions: value-based and policy-based approaches. This research shows several algorithms from those directions and how they perform. Finally, multiple open research questions are addressed and new research directions are proposed.

Algorithms for Reinforcement Learning May 06 2021 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.

Deep Reinforcement Learning Oct 23 2022 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.

Hands-On Deep Learning with R Jan 22 2020 Explore and implement deep learning to solve various real-world problems using modern R libraries such as TensorFlow, MXNet, H2O, and Deepnet Key FeaturesUnderstand deep learning algorithms and architectures using R and determine which algorithm is best suited for a specific problemImprove models using parameter tuning, feature engineering, and ensemblingApply advanced neural network models such as deep autoencoders and generative adversarial networks (GANs) across different domainsBook Description Deep learning enables efficient and accurate learning from a massive amount of data. This book will help you overcome a number of challenges using various deep learning algorithms and architectures with R programming. This book starts with a brief overview of machine learning and deep learning and how to build your first neural network. You'll understand the architecture of various deep learning algorithms and their applicable fields, learn how to build deep learning models, optimize hyperparameters, and evaluate model performance. Various deep learning applications in image processing, natural language

processing (NLP), recommendation systems, and predictive analytics will also be covered. Later chapters will show you how to tackle recognition problems such as image recognition and signal detection, programmatically summarize documents, conduct topic modeling, and forecast stock market prices. Toward the end of the book, you will learn the common applications of GANs and how to build a face generation model using them. Finally, you'll get to grips with using reinforcement learning and deep reinforcement learning to solve various real-world problems. By the end of this deep learning book, you will be able to build and deploy your own deep learning applications using appropriate frameworks and algorithms. What you will learn: Design a feedforward neural network to see how the activation function computes an output. Create an image recognition model using convolutional neural networks (CNNs). Prepare data, decide hidden layers and neurons and train your model with the backpropagation algorithm. Apply text cleaning techniques to remove uninformative text using NLP. Build, train, and evaluate a GAN model for face generation. Understand the concept and implementation of reinforcement learning in R. Who this book is for: This book is for data scientists, machine learning engineers, and deep learning developers who are familiar with machine learning and are looking to enhance their knowledge of deep learning using practical examples. Anyone interested in increasing the efficiency of their machine learning applications and exploring various options in R will also find this book useful. Basic knowledge of machine learning techniques and working knowledge of the R programming language is expected.

Reinforcement Learning with TensorFlow Jul 28 2020 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, robobrokers, 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.

Deep Reinforcement Learning Hands-On - Second Edition Mar 16 2022

Deep Reinforcement Learning Hands-On Apr 24 2020 This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. About This Book: 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. 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. 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. In Detail: 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. Style and approach: *Deep Reinforcement Learning Hands-On* explains the art of building self-learning agents using algorithms ...

The Reinforcement Learning Workshop Dec 13 2021 Start with the basics of reinforcement learning and explore deep learning concepts such as deep Q-learning, deep recurrent Q-networks, and policy-based methods with this practical guide. Key Features: Use TensorFlow to write reinforcement learning agents for performing challenging tasks. Learn how to solve finite Markov decision problems. Train models to understand popular video games like Breakout. Book Description: Various intelligent applications such as video games, inventory management software, warehouse robots, and translation tools use reinforcement learning (RL) to make decisions and perform actions that maximize the probability of the desired outcome. This book will help you to get to grips with the techniques and the algorithms for implementing RL in your machine learning models. Starting with an introduction to RL, you'll be guided through different RL environments and frameworks. You'll learn how to implement your own custom environments and use OpenAI baselines to run RL algorithms. Once you've explored classic RL techniques such as Dynamic Programming, Monte Carlo, and TD Learning, you'll understand when to apply the different deep learning methods in RL and advance to deep Q-learning. The book will even help you understand the different stages of machine-based problem-solving by using DARN on a popular video game Breakout. Finally, you'll find out when to use a policy-based method to tackle an RL problem. By the end of *The Reinforcement Learning Workshop*, you'll be equipped with the knowledge and skills needed to solve challenging problems using reinforcement learning. What you will learn: Use OpenAI Gym as a framework to implement RL environments. Find out how to define and implement reward function. Explore Markov chain, Markov decision process, and the Bellman equation. Distinguish between Dynamic Programming, Monte Carlo, and Temporal Difference Learning. Understand the multi-armed bandit problem and explore various strategies to solve it. Build a deep Q model network for playing the video game Breakout. Who this book is for: If you are a data scientist, machine learning enthusiast, or a Python developer who wants to learn basic to advanced deep reinforcement learning algorithms, this workshop is for you. A basic understanding of the Python language is necessary.

Deep Reinforcement Learning: Emerging Trends in Macroeconomics and Future Prospects May 26 2020 The application of Deep Reinforcement Learning (DRL) in economics has been an area of active research in recent years. A number of recent works have shown how deep reinforcement learning can be used to study a variety of economic problems, including optimal policy-making, game theory, and bounded rationality. In this paper, after a theoretical introduction to deep reinforcement learning and various DRL algorithms, we provide an overview of the literature on deep reinforcement learning in economics, with a focus on the main applications of deep reinforcement learning in macromodeling. Then, we analyze the potentials and limitations of deep reinforcement learning in macroeconomics and identify a number of issues that need to be addressed in order for deep reinforcement learning to be more widely used in macro modeling.

Grokking Deep Reinforcement Learning Jul 20 2022 *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

Python Reinforcement Learning Projects Jul 08 2021 Implement state-of-the-art deep reinforcement learning algorithms using Python and its powerful libraries. Key Features: Implement Q-learning and Markov models with Python and OpenAI. Explore the power of TensorFlow to build self-learning models. Eight AI projects to gain confidence in building self-trained applications. Book Description: Reinforcement learning is one of the most exciting and rapidly growing fields in machine learning. This is

due to the many novel algorithms developed and incredible results published in recent years. In this book, you will learn about the core concepts of RL including Q-learning, policy gradients, Monte Carlo processes, and several deep reinforcement learning algorithms. As you make your way through the book, you'll work on projects with datasets of various modalities including image, text, and video. You will gain experience in several domains, including gaming, image processing, and physical simulations. You'll explore technologies such as TensorFlow and OpenAI Gym to implement deep learning reinforcement learning algorithms that also predict stock prices, generate natural language, and even build other neural networks. By the end of this book, you will have hands-on experience with eight reinforcement learning projects, each addressing different topics and/or algorithms. We hope these practical exercises will provide you with better intuition and insight about the field of reinforcement learning and how to apply its algorithms to various problems in real life. What you will learn Train and evaluate neural networks built using TensorFlow for RL Use RL algorithms in Python and TensorFlow to solve CartPole balancing Create deep reinforcement learning algorithms to play Atari games Deploy RL algorithms using OpenAI Universe Develop an agent to chat with humans Implement basic actor-critic algorithms for continuous control Apply advanced deep RL algorithms to games such as Minecraft Auto-generate an image classifier using RL Who this book is for Python Reinforcement Learning Projects is for data analysts, data scientists, and machine learning professionals, who have working knowledge of machine learning techniques and are looking to build better performing, automated, and optimized deep learning models. Individuals who want to work on self-learning model projects will also find this book useful.

Reinforcement Learning Nov 12 2021 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

Learning to Play Feb 03 2021 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.

Hands-On Reinforcement Learning with Python Apr 17 2022 Reinforcement learning is a self-evolving type of machine learning that takes us closer to achieving true artificial intelligence. This easy-to-follow guide explains everything from scratch using rich examples written in Python.

An Introduction to Deep Reinforcement Learning Feb 27 2023 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.

Practical Deep Reinforcement Learning with Python Dec 01 2020 Introducing Practical Smart Agents Development using Python, PyTorch, and TensorFlow KEY FEATURES ? Exposure to well-known RL techniques, including Monte-Carlo, Deep Q-Learning, Policy Gradient, and Actor-Critical. ? Hands-on experience with TensorFlow and PyTorch on Reinforcement Learning projects. ? Everything is concise, up-to-date, and visually explained with simplified mathematics. DESCRIPTION Reinforcement learning is a fascinating branch of AI that differs from standard machine learning in several ways. Adaptation and learning in an unpredictable environment is the part of this project. There are numerous real-world applications for reinforcement learning these days, including medical, gambling, human imitation activity, and robotics. This book introduces readers to reinforcement learning from a pragmatic point of view. The book does involve mathematics, but it does not attempt to overburden the reader, who is a beginner in the field of reinforcement learning. The book brings a lot of innovative methods to the reader's attention in much practical learning, including Monte-Carlo, Deep Q-Learning, Policy Gradient, and Actor-Critical methods. While you understand these techniques in detail, the book also provides a real implementation of these methods and techniques using the power of TensorFlow and PyTorch. The book covers some enticing projects that show the power of reinforcement learning, and not to mention that everything is concise, up-to-date, and visually explained. After finishing this book, the reader will have a thorough, intuitive understanding of modern reinforcement learning and its applications, which will tremendously aid them in delving into the interesting field of reinforcement learning. WHAT YOU WILL LEARN ? Familiarize yourself with the fundamentals of Reinforcement Learning and Deep Reinforcement Learning. ? Make use of Python and Gym framework to model an external environment. ? Apply classical Q-learning, Monte Carlo, Policy Gradient, and Thompson sampling techniques. ? Explore TensorFlow and PyTorch to practice the fundamentals of deep reinforcement learning. ? Design a smart agent for a particular problem using a specific technique. WHO THIS BOOK IS FOR This book is for machine learning engineers, deep learning fanatics, AI software developers, data scientists, and other data professionals eager to learn and apply Reinforcement Learning to ongoing projects. No specialized knowledge of machine learning is necessary; however, proficiency in Python is desired. TABLE OF CONTENTS Part I 1. Introducing Reinforcement Learning 2. Playing Monopoly and Markov Decision Process 3. Training in Gym 4. Struggling With Multi-Armed Bandits 5. Blackjack in Monte Carlo 6. Escaping Maze With Q-Learning 7. Discretization Part II. Deep Reinforcement Learning 8. TensorFlow, PyTorch, and Your First Neural Network 9. Deep Q-Network and Lunar Lander 10. Defending Atlantis With Double Deep Q-Network 11. From Q-Learning to Policy-Gradient 12. Stock Trading With Actor-Critic 13. What Is Next?

digitaltutorials.jrn.columbia.edu