

Deep Reinforcement Learning Hands On Apply Modern RL Methods With Deep Q Networks Value Iteration Policy Gradients Trpo Alphago Zero And More

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

All the key deep learning methods built step-by-step in PyTorch Key Features Understand the internals and principles of PyTorch Implement key deep learning methods in PyTorch: CNNs, GANs, RNNs, reinforcement learning, and more Build deep learning workflows and take deep learning models from prototyping to production Book Description PyTorch is a new, lightweight, and Python-first tool for deep learning. Built by Facebook to offer flexibility and speed, it has quickly become the preferred tool for deep learning experts. PyTorch helps you release deep learning models faster than ever before. PyTorch Deep Learning Hands-On shows how to implement every major deep learning architecture in PyTorch. Starting with simple neural networks, it covers PyTorch for computer vision (CNN), natural language processing (RNN), GANs, and reinforcement learning. You will also build deep learning workflows with the PyTorch framework, migrate models built in Python to highly efficient TorchScript, and deploy to production using the most sophisticated available tools. Each chapter focuses on a different area of deep learning. Chapters start with a refresher on the core principles, before sharing the code you need to implement them in PyTorch. If you want to become a deep learning expert this book is for you. What you will learn Use PyTorch to build: Simple Neural Networks - build neural networks the PyTorch way, with high-level functions, optimizers, and more Convolutional Neural Networks - create advanced computer vision systems Recurrent Neural Networks - work with sequential data such as natural language and audio Generative Adversarial Networks - create new content with models including SimpleGAN and CycleGAN Reinforcement Learning - develop systems that can solve complex problems such as driving or game playing Deep Learning workflows - move effectively from ideation to production with proper deep learning workflow using PyTorch and its utility packages Production-ready models - package your models for high-performance production environments Who this book is for Machine learning professionals and enthusiasts who know Python and want to build efficient and powerful deep learning systems in PyTorch.

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. Walks through the hands-on process of building intelligent agents from the basics and all the way up to solving complex problems including playing Atari games and driving a car autonomously in the CARLA simulator. Discusses various learning environments and how to transform real-world problems into learning environments and solve using the agents. 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.

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.

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.

Explore a diverse set of meta-learning algorithms and techniques to enable human-like cognition for your machine learning models using various Python frameworks Key Features Understand the foundations of meta learning algorithms Explore practical examples to explore various one-shot learning algorithms with its applications in TensorFlow Master state of the art meta learning algorithms like MAML, reptile, meta SGD Book Description Meta learning is an exciting research trend in machine learning, which enables a model to understand the learning process. Unlike other ML paradigms, with meta learning you can learn from small datasets faster. Hands-On Meta Learning with Python starts by explaining the fundamentals of meta learning and helps you understand the concept of learning to learn. You will delve into various one-shot learning algorithms, like siamese, prototypical, relation and memory-augmented networks by implementing them in TensorFlow and Keras. As you make your way through the book, you will dive into state-of-the-art meta learning algorithms such as MAML, Reptile, and CAML. You will then explore how to learn quickly with Meta-SGD and discover how you can perform unsupervised learning using meta learning with CACTUs. In the concluding chapters, you will work through recent trends in meta learning such as adversarial meta learning, task agnostic meta learning, and meta imitation learning. By the end of this book, you will be familiar with state-of-the-art meta learning algorithms and

able to enable human-like cognition for your machine learning models. What you will learn Understand the basics of meta learning methods, algorithms, and types Build voice and face recognition models using a siamese network Learn the prototypical network along with its variants Build relation networks and matching networks from scratch Implement MAML and Reptile algorithms from scratch in Python Work through imitation learning and adversarial meta learning Explore task agnostic meta learning and deep meta learning Who this book is for Hands-On Meta Learning with Python is for machine learning enthusiasts, AI researchers, and data scientists who want to explore meta learning as an advanced approach for training machine learning models. Working knowledge of machine learning concepts and Python programming is necessary.

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 algori

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

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.

Implement key reinforcement learning algorithms and techniques using different R packages such as the Markov chain, MDP toolbox, contextual, and OpenAI Gym Key Features Explore the design principles of reinforcement learning and deep reinforcement learning models Use dynamic programming to solve design issues related to building a self-learning system Learn how to systematically implement reinforcement learning algorithms Book Description Reinforcement learning (RL) is an integral part of machine learning (ML), and is used to train algorithms. With this book, you'll learn how to implement reinforcement learning with R, exploring practical examples such as using tabular Q-learning to control robots. You'll begin by learning the basic RL concepts, covering the agent-environment interface, Markov Decision Processes (MDPs), and policy gradient methods. You'll then use R's libraries to develop a model based on Markov chains. You will also learn how to solve a multi-armed bandit problem using various R packages. By applying dynamic programming and Monte Carlo methods, you will also find the best policy to make predictions. As you progress, you'll use Temporal Difference (TD) learning for vehicle routing problem applications. Gradually, you'll apply the concepts you've learned to real-world problems, including fraud detection in finance, and TD learning for planning activities in the healthcare sector. You'll explore deep reinforcement learning using Keras, which uses the power of neural networks to increase RL's potential. Finally, you'll discover the scope of RL and explore the challenges in building and deploying machine learning models. By the end of this book, you'll be well-versed with RL and have the skills you need to efficiently implement it with R. What you will learn Understand how to use MDP to manage complex scenarios Solve classic reinforcement learning problems such as the multi-armed bandit model Use dynamic programming for optimal policy searching Adopt Monte Carlo methods for prediction Apply TD learning to search for the best path Use tabular Q-learning to control robots Handle environments using the OpenAI library to simulate real-world applications Develop deep Q-learning algorithms to improve model performance Who this book is for This book is for anyone who wants to learn about reinforcement learning with R from scratch. A solid understanding of R and basic knowledge of machine learning are necessary to grasp the topics covered in the book.

Apply modern reinforcement learning and deep reinforcement learning methods using Python and its powerful libraries Key Features Your entry point into the world of artificial intelligence using the power of Python An example-rich guide to master various RL and DRL algorithms Explore the power of modern Python libraries to gain confidence in building self-trained applications Book Description Reinforcement Learning (RL) is the trending and most promising branch of artificial intelligence. This Learning Path will

help you master not only the basic reinforcement learning algorithms but also the advanced deep reinforcement learning algorithms. The Learning Path starts with an introduction to RL followed by OpenAI Gym, and TensorFlow. You will then explore various RL algorithms, such as Markov Decision Process, Monte Carlo methods, and dynamic programming, including value and policy iteration. You'll also work on various datasets including image, text, and video. This example-rich guide will introduce you to deep RL algorithms, such as Dueling DQN, DRQN, A3C, PPO, and TRPO. You will gain experience in several domains, including gaming, image processing, and physical simulations. You'll explore TensorFlow and OpenAI Gym to implement algorithms that also predict stock prices, generate natural language, and even build other neural networks. You will also learn about imagination-augmented agents, learning from human preference, DQfD, HER, and many of the recent advancements in RL. By the end of the Learning Path, you will have all the knowledge and experience needed to implement RL and deep RL in your projects, and you enter the world of artificial intelligence to solve various real-life problems. This Learning Path includes content from the following Packt products: Hands-On Reinforcement Learning with Python by Sudharsan Ravichandiran Python Reinforcement Learning Projects by Sean Saito, Yang Wenzhuo, and Rajalingappaa Shanmugamani What you will learn Train an agent to walk using OpenAI Gym and TensorFlow Solve multi-armed-bandit problems using various algorithms Build intelligent agents using the DRQN algorithm to play the Doom game Teach your agent to play Connect4 using AlphaGo Zero Defeat Atari arcade games using the value iteration method Discover how to deal with discrete and continuous action spaces in various environments Who this book is for If you're an ML/DL enthusiast interested in AI and want to explore RL and deep RL from scratch, this Learning Path is for you. Prior knowledge of linear algebra is expected.

Deep Reinforcement Learning with Python - Second Edition will help you learn reinforcement learning algorithms, techniques and architectures – including deep reinforcement learning – from scratch. This new edition is an extensive update of the original, reflecting the state-of-the-art latest thinking in reinforcement learning.

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.

Hands-on projects cover all the key deep learning methods built step-by-step in PyTorch Key Features Internals and principles of PyTorch Implement key deep learning methods in PyTorch: CNNs, GANs, RNNs, reinforcement learning, and more Build deep learning workflows and take deep learning models from prototyping to production Book Description PyTorch Deep Learning Hands-On is a book for engineers who want a fast-paced guide to doing deep learning work with Pytorch. It is not an academic textbook and does not try to teach deep learning principles. The book will help you most if you want to get your hands dirty and put PyTorch to work quickly. PyTorch Deep Learning Hands-On shows how to implement the major deep learning architectures in PyTorch. It covers neural networks, computer vision, CNNs, natural language processing (RNN), GANs, and reinforcement learning. You will also build deep learning workflows with the PyTorch framework, migrate models built in Python to highly efficient TorchScript, and deploy to production using the most sophisticated available tools. Each chapter focuses on a different area of deep learning. Chapters start with a refresher on how the model works, before sharing the code you need to implement them in PyTorch. This book is ideal if you want to rapidly add PyTorch to your deep learning toolset. What you will learn Use PyTorch to build: Simple Neural Networks – build neural networks the PyTorch way, with high-level functions, optimizers, and more Convolutional Neural Networks – create advanced computer vision systems Recurrent Neural Networks – work with sequential data such as natural language and audio Generative Adversarial Networks – create new content with models including SimpleGAN and CycleGAN Reinforcement Learning – develop systems that can solve complex problems such as driving or game playing Deep Learning workflows – move effectively from ideation to production with proper deep learning workflow using PyTorch and its utility packages Production-ready models – package your models for high-performance production environments Who this book is for Machine learning engineers who want to put PyTorch to work.

Loosely based on the Odyssey, this landmark of modern literature follows ordinary Dubliners through an entire day in 1904. Captivating experimental techniques range from interior monologues to exuberant wordplay and earthy humor.

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.

This book will give you an in-depth view of the potential of deep learning and neural networks in game development. You will also learn to use neural nets combined with reinforcement learning for new types of game AI.

Defining Racist and Racism -- The Slippery Nature of Racial Microaggressions -- The Inability of Whites to See Themselves as Racial Beings -- Using Narrative Disclosure to Set a Tone for Examining Race -- Colleagues as Critical Lenses on Race -- What Students' Eyes Tell Us about Examining Race in the Classroom -- We Need to Prep Students -- Modeling by Leaders Is Crucial -- Conversations about Race Will Not Produce Solutions -- Normalizing Racism -- Conclusion -- Chapter 13: Negotiating the Risks of Critical Reflection -- Impostorship -- Dealing with Impostorship -- Cultural Suicide -- Avoiding Cultural Suicide -- Lost Innocence -- Marginalization -- Avoiding Political Marginalization -- Conclusion -- Chapter 14: Practicing Critically Reflective Leadership -- What Is Critically Reflective Leadership? -- Followers' Eyes -- Colleagues' Perceptions -- Theory -- Personal Experience -- Embedding Critical Reflection in Meetings -- The Circle of

Voices -- The Critical Incident Questionnaire (CIQ) -- Clearness Committee -- Appreciative Pause -- Modeling Critically Reflective Leadership -- Conclusion -- Bibliography -- Index -- EULA

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

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

"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 Beylveld 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.

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

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.

This book focuses on expert-level explanations and implementations of scalable reinforcement learning algorithms and approaches. Starting with the fundamentals, the book covers state-of-the-art methods from bandit problems to meta-reinforcement learning. You'll also explore practical examples inspired by real-life problems from the industry.

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 ESBS. 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 (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.

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 Autogenerate 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.

Deep Reinforcement Learning with Python - Second Edition will help you learn reinforcement learning algorithms, techniques and architectures - including deep reinforcement learning - from scratch. This new edition is an extensive update of the original, reflecting the state-of-the-art latest thinking in reinforcement learning.

Master reinforcement learning, a popular area of machine learning, starting with the basics: discover how agents and the environment evolve and then gain a clear picture of how they are inter-related. You'll then work with theories related to reinforcement learning and see the concepts that build up the reinforcement learning process. Reinforcement Learning discusses algorithm implementations important for reinforcement learning, including Markov's Decision process and Semi Markov Decision process. The next section shows you how to get started with Open AI before looking at Open AI Gym. You'll then learn about Swarm Intelligence with Python in terms of reinforcement learning. The last part of the book starts with the TensorFlow environment and gives an outline of how reinforcement learning can be applied to TensorFlow. There's also coverage of Keras, a framework that can be used with reinforcement learning. Finally, you'll delve into Google's Deep Mind and see scenarios where reinforcement learning can be used. What You'll Learn Absorb the core concepts of the reinforcement learning process Use advanced topics of deep learning and AI Work with Open AI Gym, Open AI, and Python Harness reinforcement learning with TensorFlow and Keras using Python Who This Book Is For Data scientists, machine learning and deep learning professionals, developers who want to adapt and learn reinforcement learning.

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

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.

This book introduces basic-to-advanced deep learning algorithms used in a production environment by AI researchers and principal data scientists; it explains algorithms intuitively, including the underlying math, and shows how to implement them using popular Python-based deep learning libraries such as TensorFlow.

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 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.

This friendly and accessible guide to AI theory and programming in Python requires no maths or data science background. Key Features Roll up your sleeves and start programming AI models No math, data science, or machine learning background required Packed with hands-on examples, illustrations, and clear step-by-step instructions 5 hands-on working projects put ideas into action and show step-by-step how to build intelligent software Book Description AI is changing the world – and with this book, anyone can start building intelligent software!

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. Taking a graduated approach that starts with the basics before easing readers into more complicated formulas and notation, Hadelin helps you understand what you really 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: Google Colab Python TensorFlow Keras 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 key skills of deep learning, reinforcement learning, and deep reinforcement learning Understand Q-learning and deep Q-learning Learn from friendly, plain English explanations and practical activities Build fun projects, including a virtual-self-driving car Use AI to solve real-world business problems and win classic video games Build an intelligent, virtual robot warehouse worker 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).

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Explore reinforcement learning (RL) techniques to build cutting-edge games using Python libraries such as PyTorch, OpenAI Gym, and TensorFlow Key Features Get to grips with the different reinforcement and DRL algorithms for game development Learn how to implement components such as artificial agents, map and level generation, and audio generation Gain insights into cutting-edge RL research and understand how it is similar to artificial general research Book Description With the increased presence of AI in the gaming industry, developers are challenged to create highly responsive and adaptive games by integrating artificial intelligence into their projects. This book is your guide to learning how various reinforcement learning techniques and algorithms play an important role in game development with Python. Starting with the basics, this book will help you build a strong foundation in reinforcement learning for game development. Each

Get Free Deep Reinforcement Learning Hands On Apply Modern RL Methods With Deep Q Networks Value Iteration Policy Gradients Trpo Alphago Zero And More

chapter will assist you in implementing different reinforcement learning techniques, such as Markov decision processes (MDPs), Q-learning, actor-critic methods, SARSA, and deterministic policy gradient algorithms, to build logical self-learning agents. Learning these techniques will enhance your game development skills and add a variety of features to improve your game agent's productivity. As you advance, you'll understand how deep reinforcement learning (DRL) techniques can be used to devise strategies to help agents learn from their actions and build engaging games. By the end of this book, you'll be ready to apply reinforcement learning techniques to build a variety of projects and contribute to open source applications. What you will learn

- Understand how deep learning can be integrated into an RL agent
- Explore basic to advanced algorithms commonly used in game development
- Build agents that can learn and solve problems in all types of environments
- Train a Deep Q-Network (DQN) agent to solve the CartPole balancing problem
- Develop game AI agents by understanding the mechanism behind complex AI
- Integrate all the concepts learned into new projects or gaming agents

Who this book is for If you're a game developer looking to implement AI techniques to build next-generation games from scratch, this book is for you. Machine learning and deep learning practitioners, and RL researchers who want to understand how to use self-learning agents in the game domain will also find this book useful. Knowledge of game development and Python programming experience are required.

[Copyright: 1304d5cb1e3c8b1c3fa190ab89f54591](#)